

Design and Implementation of Parameter Tuning on Oracle DBMS

تصميم وتنفيذ تطابق المتغيرات في نظم إدارة قواعد البيانات Oracle

By

Mohammad Jamel Jaber

Supervisor

Prof. Dr. Alaa Hussain Al-Hamami

**A Thesis Submitted in Partial Fulfillment of the
Requirement for the Degree of Master of Science
In Computer Science**

Department of Computer Science

Graduate College of Computing Studies

Amman Arab University for Graduate Studies

December, 2008

Acknowledgment

I would like to thank my supervisor Prof. Dr. Alaa Al-Hamami for his support and important suggestions he gave during my research period.

Mohammad Jameel



Resolution of the Examining Committee

This Dissertation titled "Design and Implantation of Parameter Tuning on Oracle DBMS". Has been defended and Approved on :28/3/2009




Examining Committee	Title	Signature
Dr. Ahmad Al-Jaber	Chair	
Prof. Dr. Alaa Al-Hamamy	Member and Supervisor	
Dr. Moayad Abd-Alrazaq	Member	

Table of Contents

Acknowledgment.....	II
Table of Contents.....	IV
List of Figures	VII
List of tables.....	VIII
List of Abbreviations.....	IX
Abstract.....	XI
Arabic Summary.....	XIII
Chapter 1 Introduction	1
1.1.Introduction to SQL	2
1.1.1 Parsing.....	2
1.1.2 Sharing SQL:.....	3
1.1.3 Execute SQL Statement.....	3
1.1.4 SQL Classifications	4
1.2 Oracle Database Architecture	6
1.2.1 Database Structure	6
1.2.2 Oracle Memory Structure	7
1.2.3 Process Structure.....	9
1.2.4 Physical Database Structure	9
1.2.4.1 Important files for successful running of database	10
1.2.5 Tablespaces and Data Files Construction:.....	11
1.2.5.1 Data Blocks	12
1.2.5.2 Extents	13
1.2.5.3 Segments	13
1.2.5.4 SYSTEM and SYSAUX Tablespaces.....	16
1.3 Oracle Files	17
1.3.1 The Control Files	17
1.3.2 The Parameter File (init_<SID>.ora)	17
1.3.3 The Alert Log (alert_<SID>.log).....	18
1.3.4 The Redo Log Files	18
1.4 Database Datafiles	19
1.5 Rollback Segments Construction	20

1.6 The Shared Global Area and its Contents	21
1.6.1 Oracle Initialization file init<SID>.ora.....	21
1.6.2 Dynamically Changeable Parameters	21
1.7 Introduction to Linux Operating System	23
1.7.1 Introduction to Unix	23
1.7.2 Introduction to Linux	25
1.7.3 Current application of Linux systems.....	26
1.8 Statement of the problem	27
1.9 Aim of the thesis:.....	28
1.10 Suggested solution.....	28
1.11 Research Hypothesis	28
1.12 Research Methodology	28
1.13 Contributions:	29
1.14 Thesis organization	29
Chapter 2 Linux Operating System, Oracle DBMS Literature Review	31
2.1 Fundamentals:.....	32
2.1.1 Linux Operating system:.....	33
2.1.1.1 Tunable Kernel Parameters.....	33
2.1.1.2 Undocumented Kernel Parameters	33
2.1.2 Oracle DBMS	34
2.1.2.1 Measurable Tuning Goals	34
2.1.2.2 Overview of an Instance and Instance Management.....	35
2.1.2.3 The System Global Area (SGA)	36
2.1.2.5 Managing Memory Parameters	39
2.1.2.6 .inux and Oracle Database management system parameters.....	40
2.1.6.7 SGA Configuration in Oracle	40
Chapter 3 Parameters' Design and Configuration in Linux operating system with Oracle Database Management System	42
3.1 Configuring Kernel Parameters on Linux.....	42
3.1.1 Definitions.....	45
3.1.2 Configuring Semaphores.....	47
3.1.2.2 SEMMNI	48
3.1.2.3 SEMMNS.....	49
3.1.2.4 SEMOPM	49
3.1.2.5 SHMMAX.....	50

3.1.2.6 SHMMNI.....	51
3.1.2.7 SHMALL.....	52
3.2. Oracle Initialization file init<SID>.ora.....	53
3.3 Oracle Initialization file init<SID>.ora Contents Description	54
3.4 Oracle Instance Management	65
Chapter 4 Implementation.....	67
4.1. Measuring Time in Oracle Language	67
4.1.1 Of Timer (Package spec).....	67
4.1.2 of timer (Package body)	68
4.1.3 Sptimer (package spec).....	70
4.1.4 Sptimer (package body)	70
4.2 Computer Specification: (Test Environment).....	75
4.2.1 First Environment Specification	75
4.2.2- Second Environment Specification.....	77
Chapter 5 Conclusions and Future Works	80
5.1 Conclusions.....	80
5.2 Future Works.....	80
References	82
Web Sites	86
Appendixes	88

List of Figures

Figure No	Description	Page No
Figure 4.1	Efficiency after setting parameters	54
Figure 4.2	Comparison SQL statements in different operating systems-256 RAM	55
Figure 4.3	Comparison SQL statements in different operating systems-256 RAM	56

List of tables

Table No	Description	Page No
Table 1.1	DDL and DML comparison	5
Table 1.2	SQL and SQL* Plus Comparison	5
Table 2.1	SGA Components	28
Table 2.2	Oracle and Linux parameter	30
Table 2.3	SGA Parameter in oracle after configuration)	30
Table 3.1	Settings Linux paramètres	32
Table 3.2	View Linux paramètres	33
Table 3.3	Linux Parameter Definitions	34

List of Abbreviations

ARC	Archive
CKPT	Checkpoint
DB	Data base
DBA	Database Administrator
DBMS	Database Management System
DBW	Database Writer
DCL	Data Control Language
DDL	Data Definition Language
DML	Data Manipulation Language
GB	Giga Byte
JVM	Java Virtual Machine
LGWR	Log Writer
OS	Operating System
PGA	Program Global Area
PMON	Process Monitor
RDBMS	Relational Database Management System
SGA	Shared Global Area
SMON	System Monitor
SQL	Structured Query Language
TB	Tera Byte
ORDBMS	Object relational database management system
ODBC	Open database connectivity

RECO	Recovery
HP_UX	Hewlett-Packard company – Unix operating system
KCtune	Command to reset Kernel tunable parameter
KCweb	Kernel Configuration Web-based system-Admin tools

Design and Implementation of Parameter Tuning on Oracle DBMS

Prepaid By:

Mohammad Jameel Jaber

Supervised by:

Prof. Dr. Alaa Hussain Al-Hamami

Abstract

Retrieving data in the huge database resources needs a lot of time since hardware resources can not be modified in windows operating system. This delay may lead to serious problems in large institutions such as airlines, banks, security or military institution; accordingly, the quality of services and the time of offering these services to their clients will be affected negatively.

This delay will be very clear in some operating systems, databases and some reports especially intelligent reports that are retrieved from more than one table. This delay will cause many problems; the most important one is hanging of operating systems. In order to solve this problem, and with the same resources (Memory, CPU, etc...), and without changing hardware or any computer specifications, some parameters from Linux operating system with some Oracle parameter are modified (tuning Parameter) to reduce the delay time of data retrieved from database tables, and thus reducing the hang time on available hardware resources.

This research aims to study the relationship between different Operating Systems and Database application, to increase the performance of the database using database parameter tuning with operating system parameters, as a concept can be applied on any Operating System with many differences. For instance, Windows Operating System is not as clear as Unix or Linux OSs because we can not control the hardware resources of the computer to serve operating system requirements. In the contrast with windows which is dependant on the registry editor construction, since it's difficult to be applied in such operating system. So Linux operating system was chosen to implement this study.

As a result of this thesis, and when parameters from operating system and Oracle Database are connected , the execution time is reduced in a certain amount depending on database size .

Design and Implementation of Parameter Tuning on Oracle DBMS

Prepaid By:

Mohammad Jameel Jaber

Supervised by:

Prof. Dr. Alaa Hussain Al-Hamami

Arabic Summary

المخلص

يحتاج استرجاع البيانات في مصادر قواعد البيانات الضخمة وقتاً طويلاً لأنه من غير الممكن تعديل مصادر المعدات في نظام تشغيل ويندوز ، هذا التأخير يمكن أن يؤدي إلى مشاكل في المؤسسات الكبرى (الخطوط الجوية، البنوك ، الأمن ، المؤسسات العسكرية) وتبعاً لذلك ستتأثر جودة الخدمة المقدمة لعملاء هذه المؤسسات.

يحدث هذا التأخير بوضوح في بعض أنظمة التشغيل وقواعد البيانات وبعض التقارير خصوصاً التقارير التي تحتاج لاسترجاع البيانات من أكثر من جدول مما يسبب مشاكل كثيرة أهمها فشل أنظمة التشغيل في مهامها ، ولحل هذه المشكلة وبدون تغيير أو إضافة أي معدات (ذاكرة ، وحدة معالجة مركزية...الخ) تم تعديل بعض المتغيرات من نظام تشغيل لينوكس مع متغيرات قواعد البيانات وأراكل للحد من الوقت اللازم لاسترجاع المعلومات من جداول قواعد البيانات وبالتالي تقليل فشل أنظمة التشغيل بثبات المصادر.

وقد تم إعداد هذا البحث لدراسة العلاقة بين نظم التشغيل المختلفة وتطبيقات قواعد البيانات لتحسين أداء قواعد البيانات بتعديل متغيراتها مع متغيرات نظام التشغيل، وكمفهوم يمكن تطبيقه على أي نظام تشغيل مع اختلاف هذه الأنظمة . على سبيل المثال ، يعتبر نظام التشغيل ويندوز غير واضح بالمقارنة

مع أنظمة تشغيل يونيكس أو لينوكس وذلك لأنه لا يمكن التحكم في متغيرات معدات الحاسوب (الذاكرة ، وحدة المعالجة المركزية...) وذلك لاعتماده على (Registry Editor) ، أما بالنسبة لأنظمة التشغيل يونيكس و لينوكس فانه يمكن التحكم بها. ولهذا السبب تم إختيار نظام التشغيل لينوكس لعمل هذه الدراسة.

Chapter 1

Introduction

It is well known that Database Management System (DBMS) becomes popular, and increases dramatically in various fields. So, different studies are introduced to improve the performance of application systems that depend on database management system.

Tuning Application becomes an important area to database administrators and application developers. So, both of them provide a considerable amount of time to tune the various functions of the database. A poorly tuned business application can potentially affect not only the end-users but also the entire business operation. For these reasons, companies invest significant resources to ensure smooth running of vital tuning application for their businesses.

It has to be mentioned that performing tuning of any software application is a complex job, since it requires an understanding of various components and subsystems of the software. Database tuning is the activity of making a database application run more quickly, or in other words, it may mean lower response time for time-critical applications.

Therefore Database Tuning becomes more relevant than ever before. As a consequence of this trend, Database tuning becomes an important research problem. So, the aim of this thesis is to focus more particularly on Database tuning.

1.1.Introduction to SQL

It is a standard method used for accessing any relational database, define a complete set of commands that let any person access, define, manipulate and control the database. We use the SQL because of its Market acceptance, easiness to use and its power compared to other tools.

SQL allows you to perform a wide variety of actions such as:

- Creating all types of database objects.
- Storing, retrieving and deleting data in database objects.
- Handling all operations over database objects.
- Combining data and performing calculations on it.
- Providing data security.

1.1.1 Parsing

Parsing is the process of preparing the statements for execution, this process analogous to the process of any language compiler or interpreter must be undertaken in order to translate high level language statements into machine code representation.

The Parse Process will include the following phases:

1- Check that SQL statements are syntactically valid :

The statement confirms the rules of the SQL language and all keywords and operators are valid and correctly used.

2- Check that SQL statements are semantically valid :

This means that the references to existed database objects (Tables, Indexes, etc...) are correct.

3- Check that SQL statement owner (user) has the permissions to run or to access a specific object through it.

4- **Determine the execution plan for the SQL statements :**

That describes a series of steps, which database will perform in order to access “retrieve” or update data involved in.

1.1.2 Sharing SQL:

In order to avoid unnecessary parsing process, Oracle maintains a cache “Buffer” of recently executed SQL statement together with these execution plans, this buffer area resides in SQL Shared pool area.

Whenever a request to execute SQL statement is issued, Oracle compares it with the stored SQL statement in the buffer and if matched, Oracle avoids most of the time-overhead process involved in Parsing.

1.1.3 Execute SQL Statement

Once the SQL statement is parsed and all variables had been bounded, Oracle executes the statement depending on:

1. **DML executing the statement:** results in the SQL being acted on immediately, but in case of “select” statement the execute call readies the cursor for fetch operation.
2. **Fetch operation:** the fetch call retrieves one or more rows from the database storage and store results in host variables which can be manipulated by the program.
3. **Result Set:** the output from an SQL query referred to as result set, which contain rows and columns and may be thought as a temporary table containing the query result.

1.1.4 SQL Classifications

SQL consists of several statements that manage the data and its storage units, so those statements are divided into three categories according to their work mechanism, those categories are:

1. DML: (Data Manipulation Language); this part is aware of adding, modifying and deleting data, some clauses in this category are:

- SELECT.
- INSERT.
- UPDATE.
- DELETE.

2. DDL: (Data Definition Language); this part is aware of managing data storage units and objects such as tables and indexes like creating, modifying or dropping them, some clauses in this category are :

- CREATE.
- ALTER.
- DROP.
- TRUNCATE.

3. DCL: (Data Control Language). This part is aware of controlling user's access and distributing the privileges among them, the statements in this category are:

- GRANT.
- REVOKE.

Now we need to differentiate between SQL Classes statements in DML and DDL as in Table 1.1:

Table 1.1 (DDL and DML comparison)

Data Definition Languages (DDL)	Data manipulation Languages (DML)
Affect the Structure of the object	Affect the contents of the object
Not able to roll back	Able to rollback or commit
Alter the structure of the object	Update the contents of an object
Drop the structure of an object	Delete the contents of an object

The optimality in using SQL through Oracle database appears by the enhancements supported by Oracle through creating a full environment to be used in. What we call SQLPLUS, that's used to execute any SQL statements and offers the ability to connect the database server and affects it. In other applications like MS-Access and SQL-Server we can call ODBC.

SQLPLUS is not only a connection tool, but also a container that handles any access to the database, controls and executes it. Here we can set the differences between SQL as a language and SQLPLUS as an environment for Oracle applications, see table 1.2:

Table 1.2 (SQL and SQL* Plus Comparison)

SQL	SQL*PLUS
Database Language.	Environment to work with SQL in Oracle
ANSI Standard, not product of any Company	Product of ORACLE.
Manipulates Data and Tables in the D.B	Doesn't allow changing of Values.
Consist of Statements.	Consists of Commands.
Could not be abbreviated.	Could be abbreviated.
Ends with semicolon (;).	Ends by the end of line.
Needs a buffer to be run in.	Line Interpreter. (No Need for buffer)

1.2 Oracle Database Architecture

The Oracle server is the key to information management. In general, an Oracle server must reliably manage a large amount of data in a multi-user environment so that many users can concurrently access the same data. All this must be accomplished while delivering high performance. An Oracle server must also prevent unauthorized access and provide efficient solutions for failure recovery [34].

1.2.1 Database Structure

Each running Oracle database is associated with an Oracle instance. When a database is started on a database server, the Oracle software allocates a shared memory area called the System Global

Area (SGA) and starts several Oracle background processes.

This combination of the SGA and the Oracle processes is called an Oracle instance.

After starting an instance, the Oracle software associates the instance with a specific database. This is called mounting the database. The database is then ready to be opened, which makes it accessible to authorized users. Multiple instances can execute concurrently on the same computer, each accessing its own physical database.

An Oracle database uses memory structures and processes to manage and access the database. All memory structures exist in the main memory of the computers that constitute the database server. Processes are jobs that work in the memory of these computers. A process is defined as a "thread of control" or a mechanism in an operating system that can run a series of steps[20,23].

1.2.2 Oracle Memory Structure

The basic memory structures associated with an Oracle instance include the following:

- **System Global Area (SGA):** shared by all servers and background processes.
- **Program Global Area (PGA):** Private to each server and background process. There is one PGA for each process.

The SGA is a memory area that contains data and controls information for the instance

The SGA includes the following data structures:

- **Database buffer cache:** Caches blocks of data retrieved from the database.
- **Redo log buffer:** Caches redo information (used for instance recovery) until it can be written to the physical redo log files stored on the disk.
- **Shared pool:** Caches various constructs that can be shared among users.
- **Large pool:** Is an optional area that provides large memory allocations for certain large processes, such as Oracle backup and recovery operation, and I/O server processes.
- **Java pool:** Is used for all session-specific Java code and data within the Java Virtual Machine (JVM).
- **Streams pool:** Is used by Oracle Streams[12].

A Program Global Area (PGA) is a memory region that contains data and controls information for each server process. An Oracle server process services a client's requests. Each server process has its own private PGA that is created when the server process is started. Access to the PGA is exclusive to that server process, and the PGA is read and written only by the Oracle code acting on its behalf.

With the dynamic SGA infrastructure, the size of the database buffer cache, the shared pool, the large pool, the Java pool, and the Streams pool change without shutting down the instance.

The Oracle database uses initialization parameters to create and configure memory structures.

For example, the `SGA_TARGET` parameter specifies the total amount of space available to the SGA. If you set `SGA_TARGET` to 0, Automatic Shared Memory Management is disabled.

1.2.3 Process Structure

When you invoke an application program or an Oracle tool, such as Enterprise Manager, the Oracle server creates a server process to execute the commands issued by the application. The Oracle server also creates a set of background processes for an instance that interact with each other and with the operating system to manage the memory structure, asynchronously perform I/O to write data to disk, and perform other required tasks. Which background processes are present depends on the features that are being used in the database.

1.2.4 Physical Database Structure

The files that constitute an Oracle database are organized into the following:

- **Control files:** Contain data about the database itself (that is, physical database structure information). These files are critical to the database. Without them, you cannot open data files to access the data within the database.
- **Data files:** Contain the user or application data of the database.
- **Online redo log files:** Allow for instance recovery of the

database. If the database crashes and does not lose any data files, then the instance can recover the database with the information in these files.

1.2.4.1 Important files for successful running of database

The following additional files are important to the successful running of the database:

- **Parameter file:** Is used to define how the instance is configured when it starts up.
- **Password file:** Allows users to connect remotely to the database and perform administrative tasks.
- **Backup files:** Are used for database recovery. You typically restore a backup file when a media failure or user error has damaged or deleted the original file.
- **Archive log files:** Contain an ongoing history of the data changes (redo) that are generated by the instance. Using these files and a backup of the database, you can recover a lost data file. That is, archive logs enable the recovery of restored data files.

- **Trace files:** Each server and background process can write to an associated trace file. When an internal error is detected by a process, the process dumps information about the error to its trace file. Some of the information written to a trace file is intended for the database administrator, whereas other information is for Oracle Support Services.
- **Alert log files:** These are special trace files. They are also known as alert logs. The alert log of a database is a chronological log of messages and errors. Oracle recommends that you review these files.

1.2.5 Tablespaces and Data Files Construction:

A tablespace is the primary level of logical storage in an Oracle database. All "hard" database objects such as tables, indexes, sequences, and clusters reside in tablespaces.

An Oracle database can have numerous tablespaces. The amount depends on the limit imposed by the specific operating system of open files for a given process (assuming one datafile per tablespace and subtracting for redo logs). A tablespace must have at least one datafile. Datafiles are the physical storage for an Oracle database. With the exception of redo logs and control files all data in an Oracle database is kept in datafiles.

Objects that perform similar functions and behave in similar ways should be placed in the same tablespace. For example, all lookup tables should be placed in a separate tablespace from interactive tables which should be in a separate tablespace from indexes which should be in a separate tablespace from rollback segments.

Oracle performs best if it is widely distributed. The physical datafiles should be spread across as many disks as possible. The actual number of datafiles that can be used by an Oracle instance is set at database creation with the MAXDATAFILES parameter in the CREATE DATABASE command. The number of datafiles allowed is also dependent upon the size of the operating system blocks.

The number of potential applications will drive the number and size of database tablespaces above and beyond the five base tablespaces[12].

1.2.5.1 Data Blocks

At the finest level of granularity, an Oracle database's data is stored in data blocks. One data block corresponds to a specific number of bytes of physical database space on the disc. A data block size is specified for each tablespace when it is created. A database uses and allocates free database space in Oracle database blocks [12].

1.2.5.2 Extents

The next level of logical database space is called an extent. An extent is a specific number of contiguous data blocks (obtained a single allocation) that are used to store a specific type of information.

1.2.5.3 Segments

The level of logical database storage above an extent is called a segment. A segment is a set of extents allocated for a certain logical structure. For example, the different types of segments include:

- **Data segments:** Each non-clustered, non-indexed-organized table has a data segment. All of the table's data is stored in the extents of its data segment. For a partitioned table, each partition has a data segment. Each cluster has a data segment. The data of every table in the cluster is stored in the cluster's data segment.
- **Index segments:** Each index has index segment that stored all of its data. For a partitioned index, each partition has an index segment.
- **Undo segment:** One UNDO tablespace is created by the database administrator to temporarily store undo information. The information in an undo segment is used to generate read-consistent database information and, during database recovery, to roll back uncommitted transactions for users.

- **Temporary segment:** Temporary segments are created by the Oracle database when a SQL statement needs a temporary work area to complete execution. When the statement finishes execution, the temporary segment's extents are returned to the instance for future tablespace, which is used database wide [19].

The Oracle database dynamically allocates space. When the existing extents of a segment are full, additional extents are added. Because extents are allocated as needed, the extents of a segment may or may not be contiguous on the disk.

Oracle is an Object-Relational Database Management System or, ORDBMS for short. A traditional RDBMS (Oracle Version 7 and below) stores data in tables called relations. These relations are two dimensional representations of data where the rows, called tuples in relational jargon, represent records, and the columns, called attributes, are the pieces of information contained in the record.

Oracle provides object oriented extensions to the Oracle RDBMS forming the new entity, the Object-Relational database. In an Object-Relational database, columns can represent a single value (as in standard relational databases), an array (a fixed number of additional records) or a reference to a second table where a variable amount of data can be stored. This takes the two dimensional and relational views and adds a third dimension.

In addition, in an object relational database, procedures known as methods can be tied to the tables.

Oracle files are either used for storage of data, transaction information or parameter values. Oracle data is stored in files known as datafiles. The tablespace, a logical unit of storage in Oracle, maps to one or more datafiles. Each Oracle instance must have at least a single datafile, for the SYSTEM tablespace, a control file, to maintain transaction control numbers and datafile information, at least two redo logs to contain transaction redo information and a parameter files to specify constants used to initialize the Oracle system. Therefore the minimal Oracle installation will consist of a set of 5 files and the Oracle executables.

Once Oracle starts up, more files are created. These files fall into either event or trace logging categories. A single alert log that tracks overall instance status is opened and, trace files for all baseline Oracle processes are started. There will be a minimum of four baseline Oracle processes, PMON, SMON, DBWR, LGWR with a possible fifth (RECO) if the distributed option is loaded.

A database is divided into logical storage units called tablespaces, which can be used to group related logical structures together. Each database is logically divided into one or more tablespaces. One or more data files are explicitly created for each tablespaces to physically store the data of all logical structures in a tablespaces.

Note: You can also create the bigfile tablespaces, which are tablespaces with a single but very large (up to 4 billion data blocks) data file. The traditional smallfile tablespaces (which are the default) can contain multiple data files, but the files cannot be as large.

1.2.5.4 SYSTEM and SYSAUX Tablespaces

Each Oracle database contains a SYSTEM tablespace and a SYSAUX tablespace. They are automatically created when the database is created. The system default is to create a smallfile tablespace. You can also create bigfile tablespaces, which enable the Oracle database to manage ultralarge files (up to 8 exabytes).

A tablespace can be on line (accessible) or offline (not accessible). The SYSTEM tablespace is always on line when the database is open. It stores tables that support the core functionality of the database, such as the data dictionary tables.

The SYSAUX tablespace is an auxiliary tablespace to the SYSTEM tablespace. The SYSAUX tablespace stores many database components, and it must be online for the correct functioning of all database components [20,23].

1.3 Oracle Files

1.3.1 The Control Files

The control files contain information on all physical database files (the database physical structure) and their current transaction state. The control files are read to mount and open the database and transaction numbers are recorded for each datafile. If the control files and datafiles are out of sync, the database will not start up and will report that either recovery is needed or the datafiles are out of sync with the control files. Control files are required for database startup and recovery. The database is required to have one control file, however a minimum of two on separate physical disks or on separate disk farms in a raid configuration, are recommended.

1.3.2 The Parameter File (init_<SID>.ora)

While not considered to be a part of the database since the Oracle processes do not write to it, the parameter file contains the initialization parameters that instruct the Oracle instance how to set itself up with buffers, processes, caches and the like. The parameter file is read while starting the instance during the mount and open phases of startup. Thus, the parameter file sizes all SGA components either through direct parameter values or indirectly from calculations based on parameter values specified in the parameter file. The DBA is responsible for configuration of the database using the initialization parameters. Depending on the version of Oracle, there will be more or less parameters available.

1.3.3 The Alert Log (alert_<SID>.log)

The alert log contains informational, warning and error messages dealing with the Oracle core processes and the SGA. Additionally, the alert log contains a history of any physical changes to the database such as the addition or status change of datafiles, redo logs or rollback segments. Using optional initialization parameters, information concerning checkpoints can also be recorded in the alert log.

The alert log is the only location where errors such as detection of a corrupted member of a mirrored redo log or the filling of the archive log destination (archiver stuck) are reported. Other informational messages useful for configuration (such as excessive archive waits for a checkpoint or waits due to the redo logs writing to archive) are also reported in the alert log. Oracle recommends that the log be checked at least once per day.

1.3.4 The Redo Log Files

The redo log files are set up at instance creation. A minimum of two one-member groups in a single thread is required for instance startup. We recommend running three or more groups of two mirrored members with each member being placed on a separate disk or separately controlled disk farm if possible to reduce contention. In most cases we have found five groups optimal to prevent checkpoint and archiver wait errors. The redo logs hold records used for recovery purposes. The redo logs contain information on all data modifying transactions that occur in the database unless these transactions have been run as non-recoverable. The LGWR process writes data on

changes from the redo log buffers to the redo logs. A COMMIT is not considered complete until the LGWR signals that all redo log entries have been written to disk. Remember, the redo log buffer is a part of the SGA.

1.4 Database Datafiles

The Oracle system uses logical and physical storage. Logical storage uses the concept of tablespaces. A tablespace is physically implemented through one or more datafiles. Datafiles are subdivided into segments which are subdivided into extents which may be of several types depending on what they store or their usage:

- The table segment contains data that corresponds to a table object in the database. These will usually be the most common type of segments since they store the actual data contained in the database.
- The index segment contains table index information. The index segments will contain data and ROWID pointers and will be stored as either B*-tree (the most common), cluster (either B-tree or hash) or a bitmapped format, depending on the type of index object the segment is assigned to.
- The rollback segment contains records of changes for multiple transactions. Each transaction gets assigned to a single rollback segment extent. The most costly statement in terms of rollback segment space usage is an update because it must capture both the before and after image; the least expensive is a delete because the rollback only captures the deleted ROWIDs.

Hint:

Oracle defines the Oracle instance as the Shared Global Area (SGA) and the required background processes. The base processes are SMON, PMON, DBWR and LGWR. The combination of the SGA and processes is the instance.

1.5 Rollback Segments Construction

Rollback segments are one of the least focused areas in database tuning for many DBAs. Most of the DBAs pay very little attention to tuning the rollback segments. Tuning rollback segments need a greater understanding of internal workings of Oracle. It is difficult to tune by normal means. This part gives a brief introduction to transactions and rollback segments. It goes beyond the basics of 'rollback segment is used for transaction rollback', so the reader is familiar with what is going on inside a rollback segment.

A database's read, modify, and write life cycle is called a transaction. A rollback entry is made for all transactions unless specified. A rollback entry consists of a pre update image value, a block address, a data file number, a transaction ID and the status of the transaction (Active or Committed). If the transaction fails for any reason, the old image is taken from the rollback segment. This is what is called a transaction rollback. Rollback segments are owned by SYS irrespective of who creates it. They are accessible only to Oracle, never to a user.

1.6 The Shared Global Area and its Contents

The SGA is made up of several memory components; the largest of these is usually the database buffer cache, followed by the shared pool area, and the redo log buffers. The shared pool consists of the shared SQL area (also known as the library cache) where SQL and PL/SQL statements are kept, the data dictionary cache where data dictionary related entries are kept,

1.6.1 Oracle Initialization file `init<SID>.ora`

The INIT.ORA file stores the initialization parameters for Oracle startup. By default, the file is named INIT.ORA; however, it can be renamed as long as it is used in the STARTUP command.

The control memory in Oracle database is set by assigning values to the memory parameters in the INIT.ORA file for your system. The location of this file varies, depending on the operating system you're using. Consult your Oracle Installation and Users Guide for information.

In some systems, there will be multiple INIT.ORA files; this allows different databases to have their own parameter settings. For example, INITDEV.ORA may control the developer database, INITQA.ORA the QA database, and INITPROD.ORA the production database.

1.6.2 Dynamically Changeable Parameters

- You can modify certain parameters dynamically using the command:

ALTER SYSTEM SET parameter = value;

- You can also set the parameters for a user session using the command:

ALTER SESSION SET parameter = value;

NOTE: Never assume the settings in the init.ora correctly represent the settings of the database. Many of them can be modified dynamically, and can take on new settings. These new settings do not reset when the database is restarted.

You can refer to the **V\$PARAMETER** table to determine if a parameter is changeable at either the session or database wide level. The columns of primary interest are **isses_modifiable** and **issys_modifiable**.

The **isses_modifiable** column indicates if a user who has the ALTER SESSION privilege can change the init.ora parameter for his or her session. If this column is TRUE, then the user can do so; otherwise, the column is FALSE.

The **issys_modifiable** column indicates if someone with ALTER SYSTEM privileges can change a parameter.

It has three statuses:

- **IMMEDIATE:** This indicates that the parameter is changeable and that the change takes effect immediately

- **FALSE:** This means that the parameter is not changeable dynamically
- **DEFERRED:** which indicates that the parameter is changeable but that it will not take effect until your next session

In order to take advantage of dynamically settable parameters, the user must have the ALTER SESSION privilege (which changes session-level settings) and the ALTER SYSTEM privilege (which changes system-level settings). Take care before you issue these privileges; some of the settings allow the user to change memory allocations for the user's current session. Having multiple users changing their memory settings can significantly impact performance and how much memory is used on the database. In addition, once you change a setting with the ALTER SYSTEM command, that setting remains constant even if you shut down and restart the database. Thus, the init<SID>.ora settings for a parameter may be totally incorrect.

1.7 Introduction to Linux Operating System

1.7.1 Introduction to Unix

In order to understand the popularity of Linux, we need to travel back in time, about 30 years ago...

Imagine computers as big as houses, even stadiums. While the sizes of those computers posed substantial problems, there was one thing that made this even worse: every computer had a different operating system.

Software was always customized to serve a specific purpose, and software for one given system didn't run on another system. Being able to work with one system didn't automatically mean that you could work with another. It was difficult, both for the users and the system administrators.

Computers were extremely expensive then, and sacrifices had to be made even after the original purchase just to get the users to understand how they worked. The total cost per unit of computing power was enormous [25].

Technologically the world was not quite that advanced, so they had to live with the size for another decade. In 1969, a team of developers in the Bell Labs laboratories started working on a solution for the software problem to address these compatibility issues. They developed a new operating system, which was:

1. Simple and elegant.
2. Written in the C programming language instead of assembly code.
3. Able to recycle code.

The Bell Labs developers named their project "UNIX." The code recycling features were very important. Until then, all commercially available computer systems were written in a code specifically developed for one system. UNIX on the other hand needed only a small piece of that special code, which is now commonly named the

kernel. This kernel is the only piece of code that needs to be adapted for every specific system and forms the base of the UNIX system. The operating system and all other functions were built around this kernel and written in a higher programming language, C.

1.7.2 Introduction to Linux

This language was especially developed for creating the UNIX system. Using this new technique, it was much easier to develop an operating system that could run on many different types of hardware.

The software vendors were quick to adapt, since they could sell ten times more software almost effortlessly.

Weird new situations came in existence: imagine for instance computers from different vendors communicating in the same network, or users working on different systems without the need for extra education to use another computer. UNIX did a great deal to help users become compatible with different systems[35,36].

Throughout the next couple of decades the development of UNIX continued. More things became possible to do and more hardware and software vendors added support for UNIX to their products.

UNIX was initially found only in very large environments with mainframes and minicomputers (note that a PC is a "micro" computer). You had to work at a university, for the government or for large financial corporations in order to get your hands on a UNIX system[24].

But smaller computers were being developed, and by the end of the 80's, many people had home computers.

By that time, there were several versions of UNIX available for the PC architecture, but none of them was truly free and more important: they were all terribly slow, so most people ran MS DOS or Windows 3.1 on their home PCs.

1.7.3 Current application of Linux systems

Today Linux has joined the desktop market. Linux developers concentrated on networking and services in the beginning, and office applications have been the last barrier to be taken down.

Microsoft is ruling this market, so plenty of alternatives have been started over the last couple of years to make Linux an acceptable choice as a workstation, providing an easy user interface and MS compatible office applications like word processors, spreadsheets, presentations and the like.

On the server side, Linux is well-known as a stable and reliable platform, providing database and trading services for companies like Amazon, the well-known online bookshop, US Post Office, the German army and many others. Especially Internet providers and Internet service providers have grown fond of Linux as firewall, proxy- and web server, and you will find a Linux box within reach of every UNIX system administrator who appreciates a comfortable management station. Clusters of Linux machines are used in the creation of movies such as "Titanic", "Shrek" and others. In post offices, they are the nerve centers

that route mail and in large search engine, clusters are used to perform internet searches. These are only a few of the thousands of heavy-duty jobs that Linux is performing day-to-day across the world.

It is also worth to note that modern Linux not only runs on workstations, mid- and high-end servers, but also on "gadgets" like PDA's, mobiles, a shipload of embedded applications and even on experimental wristwatches. This makes Linux the only operating system in the world covering such a wide range of hardware.[29]

1.8 Statement of the problem

In recent years, the database size and user interaction with database are increased because of the increased use of different applications by different institutions. These applications become slow because the response time is increased and throughput rates are decreased due to the large size of database from which the data are retrieved. Since there is constraints on resources of different institutions because of the high costs of modifying the hardware devices from time and then, so why don't we try to find a way to reduce the cost of regular modification of hardware devices by making use of the available resources of the hardware with smart management of Oracle database and Linux operating system. By such management, the retrieval of data will become faster in addition to that; the cost of modifying hardware is minimized.

1.9 Aim of the thesis:

The goal is not troubleshooting for the database performance problems because of its high cost. This research will encourage the system analyst and database designer to do well from the start and so prevent possible troubleshooting and bottlenecks for the system design. This will offer more time to enhance the system instead of solving high-aged problems. So the goal in this research is to provide a way to optimize Oracle-based systems, not only just tuning a database but by offering common proven tools to support system performance and tuning Oracle under different Operating Systems.

1.10 Suggested solution

The suggestion is using Oracle with Linux or Unix operating system because all hardware Parameters can be modified. Also applying such approach will be beneficial in systems using windows operating system with large database records to minimize the Execution Time for SQL statements. This approach can be applied in institution dealing with huge database such as Banks, Airlines Companies and Universities.

1.11 Research Hypothesis

Hypothesis 1: the Execution time will be reduced if we update some parameters from Linux operating system with some parameters in Oracle DBMS.

1.12 Research Methodology

The method used in this research will be empirical, contrastive and descriptive. Comparison and analysis will be a major feature in

this study. An Attempt will be made to improve the Hypothesis presented under section 1.11.

1.13 Contributions:

The main contribution of this thesis is a new effective method for high speed of retrieving information with the same hardware and without any additional costs.

1.14 Thesis organization

The thesis is divided into five chapters.

Chapter one gives a brief overview for the Oracle DBMS and different operating systems. It also describes the statement of the problem, suggests solutions and outlines the research hypothesis.

Chapter two describes a lot of researches conducted for improving the performance of database systems. Linux overviews are also outlined in this chapter.

In chapter three, we present details about Oracle files and parameters and their contents, and describe the relationship between parameters in Oracle DBMS with the same parameters in Linux operating system.

In chapter four, we design and implement a program to evaluate execution time. This program is connected in Oracle downloaded on linux, and Oracle downloaded on windows to improve that Oracle with linux operating system is better than Oracle with windows operating system with the same hardware resources.

In chapter five, we present the conclusion, the limitation of the newly developed methods, and suggest directions for future work.

Chapter 2

Linux Operating System, Oracle DBMS Literature Review

There is a lot of research conducted for improving the performance of database systems. IBM Redbook [2] is a complete system design to help system designers, system administrators, and database administrators design, size, implement, maintain, monitor, and tune a Relational Database Management System (RDBMS) for optimal performance on AIX. An RDBMS is usually a significant factor in building the profit line of a company. It represents an important investment and its performance is vital to the success of the company. This Redbook contains hints and tips from experts that work on RDBMS performance every day. It also provides introductions to general database layout concepts from a performance point of view, design and sizing guidelines, tuning recommendations, and performance and tuning information for DB2, Oracle, and IBM Informix databases. The performance tips provided here relate to the things that a system administrator or database administrator can change.

Oracle Database Performance Tuning Guide [9] is an aid for people responsible for the operation, maintenance, and performance of Oracle. This book describes detailed ways to enhance Oracle performance by writing and tuning SQL properly, using performance tools, and optimizing instance performance. It also explains how to create an initial database for good performance and includes performance-related reference information. This book could be useful for database administrators, application designers, and programmers

. Some papers have concentrated mainly on Oracle database. One study has described how Oracle overcomes these challenges and provides a way to perform automatic performance [5] diagnosis and tuning. This Paper has defined a new measure which is called 'Database Time' that provides a common currency to gauge the performance impact of any resource or activity in database.

In [3] emphasis, "Tuning rests on a foundation of informed common sense". It does not present the reader with a set of rules which, when followed, will yield an optimally functioning database. It does not present the reader with a set of formulas into which he can plug his database statistics in the hopes of generating a recommended set of database parameters. Instead, this book focuses on the foundation a rationally thinking database administrator needs to analyze and improve his system's performance.

One Guide has described how to tune Novell eDirectory on Linux and Unix Platforms[8] to improve its Performance, providing both eDirectory tuneable and Operating System specific tuneables to optimize Performance when Deployment eDirectory. Other papers have described the relationship between database and the Operating System community [26].

2.1 Fundamentals:

The following subsections will introduce some basic concepts about Linux Operating system and Oracle DBMS. This information is very important for understanding the research hypothesis.

2.1.1 Linux Operating system:

2.1.1.1 Tunable Kernel Parameters

Tunables (or configurable) kernel parameters are kernel variables that allow the operating system to be configured to fit specific system needs, resulting in better performance and/or more efficient allocation of resources. The ideal value for each parameter is often determined by the system's particular hardware configuration, the specific mix of applications the system runs, and the trustworthiness of system users; factors that vary widely from system to system.

HP attempts to provide reasonable default parameter settings, but it may be necessary or beneficial to modify these settings to better suit the needs of your particular system's users [28].

2.1.1.2 Undocumented Kernel Parameters

Some of the configurable parameters that appear in the kernel master file are not documented, and some are not known to or are not supported by kcweb and kmtune for any of several possible reasons:

- The parameter is obsolete. It is no longer used in current HP-UX releases, but might appear in an existing kernel configuration. If kcweb and kmtune encounter an obsolete parameter in the current kernel configuration, they do not display it in the list of configurable parameters that can be changed. They also remove that parameter when creating the new configuration file used to build the pending kernel.

- The parameter is not supported by kweb and kmtune. As with obsolete parameters, it is not displayed in the list of configurable parameters, but it is retained in the new kernel configuration file to ensure that no malfunctions are introduced due to a missing parameter.
- The parameter is assigned a value by kernel configuration software, which is frequently based on external factors. The assigned value might be used when calculating values for one or more other parameters.
- The parameter is for HP factory or support use only. No change from the default value should be made unless specifically directed otherwise by official HP support personnel.
- The parameter supports obsolete or obsolescent software. For information about how to select a nondefault value, consult the documentation furnished with the software that the parameter supports.
- The parameter supports independent software. For information about how to select a nondefault value, consult the documentation furnished with the software that the parameter supports[33].

2.1.2 Oracle DBMS

2.1.2.1 Measurable Tuning Goals

When tuning an Oracle10g database environment, the DBA should establish measurable tuning goals. Without them, it is difficult to determine when you have performed enough tuning.

- Response time addresses how long it takes for a user to receive data from a request table, or to generate a report.
- Database availability is also a good measure for tuning goals. Availability can be backup and recovery procedures, or by shutting down and starting the instance to tune parameters.
- Database hit percentages provide a good baseline from which to determine if performance is increasing or decreasing over time.
- Memory utilization is also a valid measure for tuning, because excessive paging and swapping can impact database and operating system performance. Memory utilization can also impact database hit percentages.

2.1.2.2 Overview of an Instance and Instance Management

Oracle Database is comprised of a set of operating system files containing data entered by users or applications and structural information about the database itself called database metadata. Information is stored persistently in these files.

In order to view or update the data contained in the database, Oracle needs to start a set of processes, called background processes, and needs to allocate some memory to be used during database operation. The background processes and memory allocated by Oracle are together known as an instance. Therefore before the database can be used, the database instance must be started. When the database instance is not available, your data is safe in the database but it cannot be accessed by any user or application[12,33].

The properties of a database instance are specified using instance initialization parameters. When the instance is started, an initialization parameter file is read and the instance is configured accordingly.

The Oracle instance and the Oracle database are separate entities, although the term instance is often used to mean one or the other or both. Technically, they are distinguished as follows:

- An Oracle instance consists of the shared memory structures and background processes that run the Oracle database. DBA can have an instance without a database (for example, when DBA have not yet created a database), and if a database exists, it can be open or not.
- A database instance refers to the physical and logical components of a specific database, and its operation.

2.1.2.3 The System Global Area (SGA)

The SGA is a shared memory area that contains data and control information for the instance. Multiple users can share data within this memory area (controlled by Oracle) and information stored in the SGA can avoid repeated access from physical disk, a time consuming operation.

For optimal performance, the SGA should be large enough to avoid frequent disk reads and writes[20,23,24].

The SGA has several subcomponents as listed in table 2.1:

Table 2.1 (SGA Components)

Component	Description
Buffer Cache	Before any data stored in the database can be queried or modified, it must be read from disk and stored in memory. The buffer cache is the component of the SGA that acts as the buffer to store any data being queried or modified. All user processes connected to the database share access to the buffer cache.
Shared Pool	The shared pool caches information that can be shared among users. Some examples: <ul data-bbox="516 1016 1377 1562" style="list-style-type: none">• SQL statements are cached so that they can be reused• Information from the data dictionary such as user account data, table and index descriptions, and privileges is cached for quick access and reusability• Stored procedures, which are executable code that is stored in the database, can be cached for faster access

Redo Log Buffer	This buffer improves performance by caching redo information (used for instance recovery) until it can be written at once and at a more opportune time to the physical redo log files that are stored on disk. Redo information and redo log files are discussed in "Redo Log Files".
Large Pool	This is an optional area that is used for buffering large I/O requests for various server processes.
Java Pool	The Java pool memory is used for all session-specific Java code and data within the Java Virtual Machine (JVM)
Streams Pool	The Streams pool is used by the Oracle Streams product.

2.1.2.4 Program Global Area (PGA)

A program global area (PGA) is a memory area used by a single Oracle server process. A server process is a process that services a client's requests. Each server process has its own private PGA area that is a nonshared area of memory created by Oracle when a server process is started [7].

The PGA is used to process SQL statements and to hold logon and other session information.

The amount of PGA memory used and its content depends on the instance configuration, that is, whether the instance is running in dedicated server or shared server mode.

2.1.2.5 Managing Memory Parameters

Some initialization parameters, referred to here as memory parameters, determine the total size of the system global area (SGA) and the program global area (PGA), and of the subcomponents of the SGA. The settings of memory parameters can affect the performance of database. When you install the database, these parameters are tuned to meet the requirements of the environment that specify[4].

If you enabled Automatic Shared Memory Configuration when you configured your database, Oracle automatically sizes the subcomponents of the SGA, which include the shared pool and buffer cache. Oracle recommends that you enable memory auto tuning on windows operating system.

2.1.2.6 .inux and Oracle Database management system parameters

Linux and Oracle Database management system parameters that are used in this research are listed in table 2.2:

Table2.2 (Oracle and Linux parameter)

Init.ora Parameter in Oracle	Kernel Parameter in Linux
db_cache_size	Shmmax
db_files (maxdatafiles)	nfile, maxfiles
large_pool_size	Shmmax
log_buffer	Shmmax
Processes	nproc, semmsl, semmns
shared_pool_size	Shmmax

2.1.6.7 SGA Configuration in Oracle

In this section oracle parameters which are configured to meet operating system parameters are determined to support research hypothesis, that Oracle parameters variables are approximately gotten by experiments and the default value depend on the operating system size.

Oracle parameters are studied and it's found that those parameters which will be configured are the same that have a direct effect, so

that configuration is done on them.(the parameters in table 2.3 after configuration) knowing that these parameters are SGA components.

Table 2.3 (SGA Parameter in oracle after configuration)

Parameter Name	Value
Fixed Size	778520 Bytes
Variables Size	128720616 Bytes
Database Buffers	159383552 Bytes
Redo Buffers	524288 Bytes
Total SGA	289406976 Bytes

Chapter 3

Parameters' Design and Configuration in Linux operating system with Oracle Database Management System

3.1 Configuring Kernel Parameters on Linux

In this chapter Linux parameters which are configured to meet oracle parameters are determined to support research hypothesis knowing that Linux parameters variables are approximately gotten by experiments.

Linux operating system parameters are studied and it is found that those parameters, which will be configured, are the same that have a direct effect, so that configuration is done on them.

Verify that the kernel parameters shown in the table 3.1 are set to values greater than or equal to the recommended value shown.

The procedure following the table describes how to verify and set the values.

Table 3.1 (Settings Linux parameters)

Parameter	Default	Recommended	File
semmsl	100	250	/proc/sys/kernel/sem
semmns	16000	32000	
semopm	32	100	
semmni	64	128	
shmall	1097152	2097152	/proc/sys/kernel/shmall
shmmax	33554432	Half the size of physical memory	/proc/sys/kernel/shmmax
shmmni	2048	4096	/proc/sys/kernel/shmmni
file-max	25280	65536	/proc/sys/fs/file-max
ip_local_port_range	32768 61000	1024 65000	/proc/sys/net/ipv4/ip_local_port_range

To view the current value specified for these kernel parameters, and to change them if necessary, follow these steps:

Enter the commands shown in table 3.2 to view the current values of the kernel parameters:

Table3.2 (View Linux Parameter)

Parameter	Command
semmsl, semmns, semopm, and semmni	# /sbin/sysctl -a grep sem This command displays the value of the semaphore parameters in the order listed.
shmall, shmmax, and shmmni	# /sbin/sysctl -a grep shm
file-max	# /sbin/sysctl -a grep file-max
ip_local_port_range	# /sbin/sysctl -a grep ip_local_port_range This command displays a range of port numbers.

If the value of any kernel parameter is different from the recommended value, complete the following steps

Using any text editor, create or edit the /etc/sysctl.conf file and add or edit lines similar to the following:

kernel.shmall = 2,097,152

kernel.shmmax = 2,147,483,648

kernel.shmmni = 4096

kernel.sem = 250 32000 100 128

fs.file-max = 65,536

net.ipv4.ip_local_port_range = 1024 65000

By specifying the values in the `/etc/sysctl.conf` file, they persist when you reboot the system.

Enter the following command to change the current values of the kernel parameters:

- a. `# /sbin/sysctl -p`
- b. Review the output from this command to verify that the values are correct. If the values are incorrect, edit the `/etc/sysctl.conf` file, then enter this command again.
- c. On UnitedLinux only, enter the following command to cause the system to read the `/etc/sysctl.conf` file when it reboots:
- d. `# /sbin/chkconfig boot.sysctl on`

3.1.1 Definitions

The following table (Table 3.3) shows previous Linux parameters' definitions briefly .

Table3.3 (Linux Parameter Definitions)

maxfiles	Soft file limit per process.
maxuprc	Maximum number of simultaneous user processes per userid.
nfile	Maximum number of simultaneously open files systemwide at any given time.
nproc	Maximum number of processes that can exist simultaneously in the system.
shmmax	The maximum size(in bytes) of a single shared memory segment.
shmmin	The minimum size(in bytes) of a single shared memory segment.
shmni	The number of shared memory identifiers.
shmseg	The maximum number of shared memory segments that can be attached by a process.
semmns	The number of semaphores in the system.

semnmi	The number of semaphore set identifiers in the system; determines the number of semaphore sets that can be created at any one time.
semmsl	The maximum number of semaphores that can be in one semaphore set. It should be same size as maximum number of Oracle processes

3.1.2 Configuring Semaphores

Now that we have configured our shared memory settings, it is time to take care of configuring our semaphores. A semaphore can be thought of as a counter that is used to control access to a shared resource. Semaphores provide low level synchronization between processes (*or threads within a process*) so that only one process (*or thread*) has access to the shared segment, thereby ensuring the integrity of that shared resource. When an application requests semaphores, it does so using "sets" [32].

To determine all semaphore limits, use the following values:

1. # ipcs -ls
2. max number of arrays = 128
3. max semaphores per array = 250
4. max semaphores system wide = 32000

5. max ops per semop call = 32
6. semaphore max value = 32767
7. You can also use the following command:

```
# cat /proc/sys/kernel/sem  
250 32000 32 128
```

In the following , there is a brief illustration about Linux parameters that are used and configured in this research:

3.1.2.1 SEMMSL

The SEMMSL kernel parameter is used to control the maximum number of semaphores per semaphore set [28].

Oracle recommends setting SEMMSL to the largest PROCESS instance parameter setting in the init.ora file for all databases hosted on the Linux system plus 10. Also, Oracle recommends setting the SEMMSL to a value of no less than 100.

3.1.2.2 SEMMNI

The SEMMNI kernel parameter is used to control the maximum number of semaphore sets on the entire Linux system.

Oracle recommends setting the SEMMNI to a value of no less than 100.

3.1.2.3 SEMMNS

The SEMMNS kernel parameter is used to control the maximum number of semaphores (not semaphore sets) on the entire Linux system.

Oracle recommends setting the SEMMNS to the sum of the PROCESSES instance parameter setting for each database on the system, adding the largest PROCESSES twice, and then finally adding 10 for each Oracle database on the system. To summarize:

SEMMNS = sum of PROCESSES setting for each database on the system
+ (2 * [largest PROCESSES setting])
+ (10 * [number of databases on system])

To determine the maximum number of semaphores that can be allocated on a Linux system, use the following calculation. It will be the lesser of:

SEMMNS -or- (SEMMSL * SEMMNI)

3.1.2.4 SEMOPM

The SEMOPM kernel parameter is used to control the number of semaphore operations that can be performed per SEMOPM system call [28].

The SEMOP system call (function) provides the ability to do operations for multiple semaphores with one SEMOP system call. A

semaphore set can have the maximum number of SEMMSL semaphores per semaphore set and is therefore recommended to set SEMOPM equal to SEMMSL.

3.1.2.5 SHMMAX

The SHMMAX parameter is used to define the maximum size (in bytes) for a shared memory segment and should be set large enough for the largest SGA size. If the SHMMAX is set incorrectly (*too low*), it is possible that the Oracle SGA (*which is held in shared segments*) may be limited in size. An inadequate SHMMAX setting would result in the following:

ORA-27123: unable to attach to shared memory segment

You can determine the value of SHMMAX by performing the following:

```
# cat /proc/sys/kernel/shmmax  
33,554,432
```

As you can see from the output above, the default value for SHMMAX is 32MB. This is often too small to configure the Oracle SGA. I generally set the SHMMAX parameter to 2GB.

To change the value SHMMAX, you can use either of the following three methods:

- This method is used most often. This method sets the SHMMAX on startup by inserting the following kernel parameter in the /etc/sysctl.conf startup file:

```
# Echo "kernel.shmmax=2,147,483,648" >> /etc/sysctl.conf
```

- If you wanted to dynamically alter the value of SHMMAX without rebooting the machine, you can make this change directly to the /proc file system. This command can be made permanent by putting it into the /etc/rc.local startup file:

```
# Echo "2147483648" > /proc/sys/kernel/shmmax
```

- You can also use the sysctl command to change the value of SHMMAX:

```
# sysctl -w kernel.shmmax=2147483648
```

3.1.2.6 SHMMNI

We now look at the SHMMNI parameter. This kernel parameter is used to set the maximum number of shared memory segments system wide. The default value for this parameter is 4096. This value is sufficient and typically does not need to be changed.

You can determine the value of SHMMNI by performing the following:

```
# cat /proc/sys/kernel/shmmni  
4096
```

3.1.2.7 SHMALL

Finally, we look at the SHMALL shared memory kernel parameter. This parameter controls the total amount of shared memory (in pages) that can be used at one time on the system. In short, the value of this parameter should always be at least:

$$\text{ceil}(\text{SHMMAX}/\text{PAGE_SIZE})$$

The default size of SHMALL is 2097152 and can be queried using the following command:

```
# cat /proc/sys/kernel/shmall  
2097152
```

From the above output, the total amount of shared memory (in bytes) that can be used at one time on the system is:

$$\begin{aligned}\text{SM} &= (\text{SHMALL} * \text{PAGE_SIZE}) \\ &= 2097152 * 4096 \\ &= 8,589,934,592 \text{ bytes}\end{aligned}$$

maxfiles - Soft file limit per process.

maxuprc - Maximum number of simultaneous user processes per userid.

nfile - Maximum number of simultaneously open files systemwide at any given time.

nproc - Maximum number of processes that can exist simultaneously in the system.

shmmax - The maximum size(in bytes) of a single shared memory segment.

shmmmin - The minimum size(in bytes) of a single shared memory segment.

shmmni - The number of shared memory identifiers.

shmseg - The maximum number of shared memory segments that can be attached by a process.

semmns - The number of semaphores in the system.

semmni - The number of semaphore set identifiers in the system; determines the number of semaphore sets that can be created at any one time.

semmsl - The maximum number of semaphores that can be in one semaphore set. It should be same size as maximum number of Oracle processes, See Appendix B for more abbreviations

3.2. Oracle Initialization file init<SID>.ora

The INIT.ORA file stores the initialization parameters for Oracle startup. By default, the file is named INIT.ORA; however, it can be renamed as long as it is used in the STARTUP command[11].

You control memory in your Oracle database by assigning values to the memory parameters in the INIT.ORA file for the system. The location of this file varies depending on the operating system you're using. Consult Oracle Installation and Users Guide for information.

In some systems, there will be multiple INIT.ORA files; this allows different databases to have their own parameter settings. For example, INITDEV.ORA may control the developer database, INITQA.ORA the QA database, and INITPROD.ORA the production database. See Appendix A.

3.3 Oracle Initialization file init<SID>.ora Contents Description

db_name = "project" The name of the database.

DB_NAME can specify a database identifier of up to eight characters. If specified, it must correspond to the name specified in the **CREATE DATABASE** statement. Although the use of **DB_NAME** is optional, it should generally be set before invoking **CREATE DATABASE** and then referenced in that statement.

If not specified, a database name must appear on either the **STARTUP** or the **ALTER DATABASE MOUNT** command line for each instance of the parallel server.

db_domain = world

This parameter specifies the extension components of a global database name, consisting of valid identifiers, separated by periods. Specifying **DB_DOMAIN** as a unique string for every database is highly recommended.

instance_name = project

INSTANCE_NAME is a string value representing the name of the instance and is used to uniquely identify a specific instance when multiple instances share common services names. **INSTANCE_NAME** should not be confused with the **SID**, which actually uniquely identifies the instances, shared memory on a host.

service_names = project.world

SERVICE_NAMES specifies the service names supported by the instance. , it is one or more strings which represent the names of the database on the network. It is possible to provide multiple services names so that different usages of a single database can be identified separately. Service names can also be used to identify a single service that is available from two different databases through the use of replication.

db_files = 1024

Number of database files that can be open when the database is running. Set this value lower than the default if you are not using

32 data files (to reduce the space used in the SGA). You can increase this value by shutting down your database, changing the parameter value, and restarting the database. All instances must be set to the same value if you are using the parallel server.

```
control_files = ("D:\Oracle\oradata\project\control01.ctl",  
                "D:\Oracle\oradata\project\control02.ctl",  
                "D:\Oracle\oradata\project\control03.ctl")
```

The names for the instance control files. CONTROL_FILES specifies one or more names of control files, separated by commas.

```
open_cursors = 300
```

Maximum number of cursors that a user session can have opens at any one time. If the number of cursors being held by users is frequently near the maximum, increase the OPEN_CURSORS value for a performance boost. If the setting is too high, you will be wasting memory. A typical setting for users of a large application is between 200 and 300.

max_enabled_roles = 30

Maximum number of roles per user. The actual number of roles a user can enable is 2 plus the value of MAX_ENABLED_ROLES, because each user has two additional roles, PUBLIC, and the user's own role.

db_file_multiblock_read_count = 8

Number of blocks read into the buffer cache at once when performing a sequential scan. The maximum number of blocks that can be read in one I/O request is defined by the system = (max_io_size/db_block_size).

Setting this value too low causes additional I/O on sequential reads. Setting db_file_multiblock_read_count as high as possible however can cause core dumps, Oracle errors, and performance degradation when using a 'where key in' expression. Since most users rarely perform full-table scans there would be no benefit anyway.

For best performance, it is usually best to set the parameter to 8 for OLTP application and either 16 or 32 for decision support databases and batch processing.

db_block_buffers = 2048

Sets the size of the database buffer cache in memory. You must set this parameter, as well as the SHARED_POOL_SIZE and LOG_BUFFER, to get optimal performance by caching data from the database in memory. The buffer cache stores tables, indexes, clusters, sort data, dictionary data, and rollback segments. Our testing shows that the higher the number of block buffers, the less I/O and the better your system will perform. If excessive paging and swapping activity occurs for user processes or if any paging or swapping of the SGA occurs, you will have to reduce DB_BLOCK_BUFFERS to free memory. You should enlarge this parameter only after you are certain that the SHARED_POOL_SIZE parameter (dictionary and library caches) has been adequately tuned.

For example, if the database block size is 2K (the default value for db_block_size), the default value of db_block_buffers is 24576.

shared_pool_size = 52428800

Size of the shared buffer pool in the SGA. This pool stores shared SQL and PL/SQL blocks, packages, procedures, functions, triggers, the data dictionary cache, and (if your site is using a multithreaded server architecture) some session information.

Make the shared pool large enough, but not too large. If your shared pool is too large, you are wasting memory that could otherwise be used to enlarge the buffer cache. If you set it too low, you'll need to do too many disk accesses as objects are reloaded into memory, and performance will suffer-sometimes as much as 50%.

large_pool_size = 614400

Size in bytes of the large allocation pool.

java_pool_size = 20971520

JAVA_POOL_SIZE specifies the size in bytes of the Java pool.

log_checkpoint_interval = 10000

Number of new redo log file blocks needed to trigger a checkpoint. Note that these blocks are operating system blocks, not Oracle blocks. If you set this parameter smaller than the sizes of the redo log files, it will cause a checkpoint to occur prior to a log switch. A checkpoint causes all modified database buffers to be written to disk and stores the location from the redo log where check pointing has occurred in the control file and database files.

log_checkpoint_timeout = 1800

Forces more frequent checkpoints. You'll get the best performance by leaving the LOG_CHECKPOINT_TIMEOUT parameter at its default value of 0, which causes a checkpoint only on a change of redo log. The exception to the rule may be to set it to every 10 minutes or so if you are experiencing very long delays when a log switch is occurring. This does not always solve the problem, because there are many potential causes. Make sure that you also set LOG_CHECKPOINT_INTERVAL to a size larger than that of the redo log.

Processes = 150

Max number of simultaneous connects to Oracle.

log_buffer = 32768

Number of bytes allocated to the redo log buffer in the SGA. The log buffer can affect the performance of the RDBMS by buffering information before writing to the redo logs. If you are experiencing I/O bottlenecks on the disks that contain the redo logs, increase this parameter. For high-volume, intensive-update applications, we have noticed response improvements of 50% by

increasing the buffer size from 8 kilobytes (the default value) to 1 megabyte. There were no advantages to increasing beyond 1 megabyte.

```
max_dump_file_size = 10240
```

Limits the physical size of the trace file to the specified number of operating system blocks (or UNLIMITED). If you enable the SQL_TRACE parameter for the entire database, this option helps control the amount of disk space used. To find out what size to specify, find out the number of operating system blocks available in your system. If SQL_TRACE runs out of space, it will truncate your output; you'll have to allocate more space and start again.

```
global_names = true
```

Enables db link name checking. Oracle recommends that if you use or will use distributed processing; set this parameter to TRUE to ensure a unique identifying name for your database in a networked environment.

```
background_dump_dest = D:\Oracle\admin\project\bdump
```

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

An ALERT file in the directory specified by **BACKGROUND_DUMP_DEST** logs significant database events and messages. Anything that affects the database instance-wide or globally is recorded here. This file records all instance start ups and shut downs, messages to the operator console, and errors that cause trace files to be written. It also records every CREATE, ALTER, or DROP operation on a database, tablespace, or rollback segment.

The ALERT file is a normal text file. Its filename is operating system-dependent. For platforms that support multiple instances, it takes the form **ALERT_sid.LOG**. This file grows slowly, but without limit, so the database administrator might want to delete it periodically. The file can be deleted even when the database is running.

db_block_size = 8192

Size of each block database buffer. Oracle recommends that you set the parameter to a minimum of 4 kilobytes. We performed several tests, creating a small database on a UNIX system with a DB_BLOCK_SIZE twice the default size. (The default was 2048, and we set it to 4096.)

job_queue_interval = 60

The job queue processes “wake up” periodically and check the job queue catalog to see if any jobs are due to execute.

open_links = 4

Max number of open database links per user.

Compatible = 8.1.0

COMPATIBLE allows you to use a new release, while at the same time guaranteeing backward compatibility with an earlier release. This is helpful if it becomes necessary to revert to the earlier release.

This parameter specifies the release with which the Oracle server must maintain compatibility. It allows you to take advantage of the maintenance improvements of a new release immediately in

your production systems without testing the new functionality in your environment. Some features of the current release may be restricted.

sort_area_size = 65536

Size in bytes that a user process has available for sorting. If the machine on which the database is being created has an abundance of memory, you can increase this parameter beyond the default (for example, 2,097,152).

sort_area_retained_size = 65536

All users are given the number of bytes specified in this parameter for sorting. If the user requires some more memory for sorting, the user's memory allocation can increase up to SORT_AREA_SIZE. This parameter is the size in bytes to which Oracle will reduce your sort area if sort data is not being referenced. Memory is reduced only after all of the rows have been fetched from the sort space. Sometimes, a number of concurrent sorts may be required, and each is given its own memory allocation of a size that is determined by this parameter.[36]

3.4 Oracle Instance Management

An Oracle database server consists of an Oracle database and Oracle instance. An Oracle instance is made up of memory structures, known as the System Global Area (SGA), and background processes that handle much of the behind-the-scenes work involved in running an instance[14]. The most common background processes are the following:

- **System Monitor (SMON):** Performs crash recovery when the instance is started following a failure.
- **Process Monitor (PMON):** performs process cleanup when a user process fails.
- **Database Writer (DBWn):** writes modified blocks from the database buffer cache to the data files on the disk.
- **Checkpoint (CKPT):** Updates all the data files and control files of the database to indicate the most recent checkpoint.
- **Log Writer (LGWR):** Writes redo log entries to the disk
- **Archiver (ARCn):** Copies redo log files to the archival storage when logs switch occurs [31].

SMON - System Monitor

The SMON (Server Monitor) process monitors the Oracle instance, recovers temporary segment space when it is no longer needed. It coalesces contiguous areas of free space in database tablespaces where the default storage parameter, PCTINCREASE, is not set to zero, and recovers dead transactions skipped during crash and instance recovery because of file-read or offline errors. These transactions are eventually recovered by SMON when the tablespace or file is brought back online. If SMON dies, the instance will crash requiring instance recovery on restart.

PMON - Process Monitor

The PMON (Process Monitor) process cleans up failed transactions. This clean up involves cleaning out cache entries, releasing locks and freeing other process resources. If MTS is being used, PMON watches for dead dispatcher and server processes and restarts them as needed. If PMON dies, the instance will crash requiring instance recovery on startup. In the situation where an instance crashed and is then restarted, PMON will rollback any uncommitted transactions using the information in the rollback segments.

Chapter 4

Implementation

4.1. Measuring Time in Oracle Language

In this research Linux parameters and Oracle parameters are configured to prove the research hypothesis which says that if Oracle and Linux parameters are configured in a certain values, then the results of performance are found in a more effective way than windows and oracle parameters when used, so a program is developed to measure the time needed to execute SQL statements on each operating system with showing the actual time needed to execute these statements.

The following program describes Of-Timer and SPTimer packages(procedures and functions) , each one of them consists of package specification and package body, and this is a standard way to write in oracle language, the package is called from oracle form to execute it and show the results.

4.1.1 Of Timer (Package spec)

```
PACKAGE of_timer IS  
  PROCEDURE capture (context_in IN VARCHAR2 := NULL);  
  PROCEDURE show_elapsed  
    (prefix_in IN VARCHAR2 := NULL,  
     reset_in IN VARCHAR2 := 'RESET',  
     context_in IN VARCHAR2 := NULL,  
     type_in IN VARCHAR2 := 'MESSAGE');
```

```
FUNCTION elapsed RETURN NUMBER;  
END;
```

4.1.2 of timer (Package body)

```
PACKAGE BODY of_timer IS  
  PROCEDURE capture (context_in IN VARCHAR2 := NULL) IS  
  BEGIN  
    sptimer.capture (context_in);  
  END;  
  
  PROCEDURE show_elapsed  
  (prefix_in IN VARCHAR2 := NULL,  
   reset_in IN VARCHAR2 := 'RESET',  
   context_in IN VARCHAR2 := NULL,  
   type_in IN VARCHAR2 := 'MESSAGE')  
  IS  
    alert_response INTEGER;  
  BEGIN  
    IF UPPER (type_in) = 'MESSAGE'  
    THEN  
      MESSAGE  
        (sptimer.elapsed_message (prefix_in,  
reset_in, context_in));  
    ELSE
```

```
SET_ALERT_PROPERTY
(type_in,
  ALERT_MESSAGE_TEXT,
  sptimer.elapsed_message      (prefix_in,      reset_in,
context_in));
  alert_response := SHOW_ALERT (type_in);
END IF;
END;

FUNCTION elapsed RETURN NUMBER IS
BEGIN
  RETURN sptimer.elapsed;
END;

END of_timer;
```

4.1.3 Sptimer (package spec)

```
PACKAGE sptimer IS
  PROCEDURE capture (context_in IN VARCHAR2 := NULL);
  FUNCTION elapsed RETURN NUMBER;
  FUNCTION elapsed_message
    (prefix_in IN VARCHAR2 := NULL,
     reset_in IN VARCHAR2 := 'RESET',
     reset_context_in IN VARCHAR2 := NULL ) RETURN
  VARCHAR2;
  PROCEDURE show_elapsed(prefix_in IN VARCHAR2 :=
  NULL,reset_in IN VARCHAR2 := 'RESET');
END;
```

4.1.4 Sptimer (package body)

```
PACKAGE BODY sptimer IS
  /* Package variable which stores the last timing made */
  last_timing NUMBER := NULL;

  /* Package variable which stores context of last timing */
  last_context VARCHAR2 (500) := NULL;
  PROCEDURE capture (context_in IN VARCHAR2 := NULL)
  /* Save current time and context to package variables. */
  IS
  BEGIN
```

```

        last_timing := DBMS_UTILITY.GET_TIME;
        last_context := context_in;
END;

FUNCTION elapsed_message
    (prefix_in IN VARCHAR2 := NULL,
     reset_in IN VARCHAR2 := 'RESET',
     reset_context_in IN VARCHAR2 := NULL)
RETURN VARCHAR2
/*
|| Construct message for display of elapsed time.
Programmer can
|| include a prefix to the message and also ask that the last
|| timing variable be reset/updated. This saves a separate
call
|| to elapsed.
*/
IS
    current_timing NUMBER;
    return_value VARCHAR2 (500);
BEGIN
    IF last_timing IS NULL
    THEN
        /* If there is no last_timing, cannot show
anything. */

```

```

return_value := NULL;

ELSIF last_context IS NOT NULL
THEN
    /* Construct message with context of last call to
elapsed */
    current_timing := elapsed;
    return_value :=
        (prefix_in || ' Elapsed since ' ||
         last_context || ': ' ||
         TO_CHAR (ROUND (current_timing/100, 2))
||
         ' seconds.');
```

```

last_context := NULL;

ELSE
    /* Construct message without the context. */
    current_timing := elapsed;
    return_value :=
        (prefix_in || ' Elapsed: ' || TO_CHAR
(current_timing));
    END IF;

IF UPPER (reset_in) = 'RESET'
```

```

        THEN
            capture (reset_context_in);
        END IF;

        RETURN return_value;
    END;

FUNCTION elapsed RETURN NUMBER IS
BEGIN
    IF last_timing IS NULL
    THEN
        RETURN NULL;
    ELSE
        RETURN      DBMS_UTILITY.GET_TIME      -
last_timing;
    END IF;
END;

PROCEDURE show_elapsed
    (prefix_in IN VARCHAR2 := NULL,
    reset_in IN VARCHAR2 := 'RESET')
/* Little more than a call to the elapsed_message function! */
IS
BEGIN

```

```

DBMS_OUTPUT.PUT_LINE      (elapsed_message
(prefix_in, reset_in));
END;
END;

```

This research mainly concentrates on the configuration of different Oracle and database parameter , because controlling these parameters will lead to dramatic reduction on time required for data retrieval more than any other steps that SQL statements pass through.




Figure 4.1 (efficiency after setting parameters)

It is well known that the size of data increases with time because new clients will be recruited to the database, As a result of that there will be “A” bottlenecks when executing SQL statements and this affects the efficiency of operating system and database. With configuration of Oracle parameters and operating system, we have reduced the time

of data retrieval. In other words we have shifted bottlenecks from “A” to “B” where larger bulk of data can be retrieved with minimal delay time given that hardware resources are not modified at all.

4.2 Computer Specification: (Test Environment)

4.2.1 First Environment Specification

Processor : 798 MHz

Memory : 256 MB

Operating systems: Windows XP Professional (Service Pack 2)

Red Hat Enterprise Linux ES (2.6.9-55 EL)

The following table contains SQL statements that are chosen randomly and executed by using the previous program on two different operating systems (Linux, Windows), the results are shown below and also the figure for clarification.

	SQL	Linux (Time in Seconds)	Windows (Time in Seconds)
1	Select max(crnt_bal) from transact	7.91	34.44
2	Select * from all_tables	11.52	16.55

3	Select cus_sho_name from address , account where address.bra_code = account.bra_code and address.cus_num = account.cus_num and address.cur_code = 0 and address.led_code = 0 and address.sub_acct_code = 0	150	186.1
4	select cur_code ,deb_cre_ind,count(*) as count_cur ,sum(tra_amt) as sum_cur from transact where cur_code in (1,2,3)group by cur_code ,deb_cre_ind	7.77	10.11
5	select distinct bra_code from transact	9.88	10.68




Figure 4.2(Comparison SQL statements in different operating systems-256 RAM)

4.2.2- Second Environment Specification

Processor : 1500 MHz

Memory : 512 MB

Operating systems: Windows XP Professional (Service Pack 2)

Red Hat Enterprise Linux ES (2.6.9-55 EL)

The following table contains SQL statements that are chosen randomly and executed by using the previous program on two different operating systems (Linux, Windows), the results are shown below and also the figure for clarification.

	SQL	Linux (Time in Seconds)	Windows (Time in Seconds)
1	Select max(crnt_bal) from transact	1.2	15.5
2	Select * from all_tables	0.8	7.4
3	Select cus_sho_name from address , account where address.bra_code = account.bra_code and address.cus_num = account.cus_num and address.cur_code = 0 and address.led_code = 0 and address.sub_acct_code = 0	54.1	80.7
4	select cur_code ,deb_cre_ind,count(*) as count_cur ,sum(tra_amt) as sum_cur from transact where cur_code in (1,2,3) group by cur_code ,deb_cre_ind	1.0	7.1
5	select distinct bra_code from transact	2.2	6.2




Figure 4.3 (Comparison SQL statements in different operating systems – 512 RAM)

After executing SQL statements on different operating systems , results show that using 512 RAM memory is more efficient than using 256 RAM memory , since SGA components consumes more memory , so the setting of parameters is done on 512 RAM memory.

The above figures show the Executing time in Linux operating system is better than Windows operating system because the red line could reach to hanging point.

Chapter 5

Conclusions and Future Works

5.1 Conclusions

In this research we provide an overview of kernel parameters that are related to Linux operating system and the way of configuring them. On the other hand, we illustrate Oracle Architectures, parameters and configuration.

- Desire to reduce the execution and retrieval time in the same resources without additional cost.
- We connect the parameters from operating systems with Oracle parameters. We find that if we only tune the operating system parameters, time will be reduced in a certain amount, also if we only tune Oracle parameters, time will also be reduced in another certain amount, so we combine all parameters to get the least possible execution time. That means the SQL statements when applied on Linux operating system are executed successfully, in contrast with windows operating system which hanged when applied the same SQL statements , although the memory space is the same.

5.2 Future Works

This thesis suggested to use Oracle with Linux or Unix operating system because it is possible to modify all hardware Parameters, and in applying such approach will be beneficial in systems using windows operating system with large database records to minimize the Execution Time for SQL Statements. This approach can be applied in institutions dealing with huge database such as banks, airlines companies and universities.

The future vision of this subject is to thoroughly investigate it in different dimensions; the following ideas may be applicable in the near future:

- The database administrator designs a utility to reach the hardware parameters that are located in windows operating system where this utility can modify these parameters automatically to make them suitable for oracle engine.
- By utilizing the same hardware the database administrator may concentrate on memory management to enhance the efficiency and hasten the performance of query executions, thus reducing the time required for data retrieval.

To achieve the best results, database administrators should build an architecture server that contains Unix or Linux operating system along with oracle engine. In this server the database administrator configures the parameters where the end users (Clients) will have access to the server by using (TNSlistener) regardless of the operating system the client is using.

References

- [1] A. Aboulnaga , “Building Histograms Without Looking at Data.” , Computer Science Department , university of Wisconsin, 1999 .
- [2] B. Darmawan, G. Groenewald, A Irving, S. Henriqu Soares Monteiro,K. M. Snedeker , "Database Performance Tuning on AIX" , January 2003.
- [3] D. Shasha & P. Bonnet , M. Kaufmann Publishers; "Database Tuning Principles, Experiments, and Troubleshooting" , <http://www.mkp.com/dbtune> Techniques , 2002.
- [4] Jeffrey D. Ullman , "Improving the Efficiency of Database System" , 2002.
- [5] K. Dias , M. Ramacher , U. Shaft , V. Venkataramani , G. Wood , "Automatic Performance Diagnosis and Tuning in Oracle" , Oracle Corporation , 2001.
- [6] McGraw-Hill Companies, "Oracle Performance Tuning 101 E-Books", OraclePressBook.com , 2001.
- [7] N. Bruno , S. Chaudhuri , "Automatic Physical Database Tuning" , Microsoft Researches, 2005 .

[8] Novel Company , "Performance Tuning for Linux and UNIX" , Novell E-Directory 8.7 , www.novell.com , 2002.

[9] Oracle Company , "Oracle Database Performance Tuning Guide 10g" Part No. B10752-01 , December 2003.

[10] Oracle Company , "Enterprise DBA-Architecture and Administration" 30049GC10 , August 1999.

[11] Oracle Company , "Oracle Database 10g – Administration workshop 1 , Volume 1 , D17090GC30 , November 2005" .

[12] Oracle Company , "Oracle Database 10g: Administration Workshop 1" Volume 2 , D17090GC30 , November 2005.

[13] Oracle Company , "Oracle Database 10g: Performance Tuning" , Volume 2 , D19165GC10 , February 2006.

[14] Oracle Company , "Oracle Database Administrator's Guide 10g" B14231 , May 2006.

[15] Oracle Company , “Oracle database installation guide 10g Release 1 (10.1) for UNIX Systems” Part No. B10811-01 , 2003.

[16] Oracle Company , “Oracle forms developer 10g- Build Internet Application” , D17251GC10 , June 2004.

[17] Oracle Company , Data Warehousing Fundamentals 50102GC20 , May 1999.

[18] Oracle Company , Oracle Database 10g SQL Fundamentals , D17108GC10 , March 2004 .

[19] Oracle Company , Oracle Database 10g, Administrator Workshop2 , D17092GC30 , January 2006.

[20] Oracle Company , Oracle Database 10g-Develop PL/SQL program units , Volume 1 , D17169GC11 , August 2004.

[21] Oracle Company , Oracle Reports Developer 10g: Build Reports , D17075GC10 , June 2004.

[22] Oracle Corporation , “Oracle Database 10g: Develop PL-SQL Program Units” , Volume 2 , D17169GC11 , August 2004.

[23] Oracle Forms Developer 10g: “Build Internet Application” Volume 2 , D39559 , June 2004.

[24] Michael Kifer, Scott A. Smolka , Introduction to Operating System Design and Implementation , 2007.

[25] Maurice J. Bach, The design of the UNIX operating system , 1986.

[26] S. Harizopoulos, A. Ailamaki , "A Case for Staged Database Systems" , In Proc. Of 1st Conference on Innovative Data Systems Research, 2003.

[27] S. Kounev , B. Weis and A. Buchmann , "Performance Tuning and Optimization of J2EE Application on the JBOSS Platform" , Department of Computer science , Darmstadt University of Technology , Germany , October 2004.

Web Sites

[28]

<http://docs.hp.com/en/939/KCParams/KCparams.OverviewAll.html>.
kernel parameters , Last Accessed on September 2008.

[29]<http://publib.boulder.ibm.com/infocenter/db2luw/v8/index.jsp?topic=/com.ibm.db2.udb.doc/start/t0008238.htm> Last accessed on
September 2008

[30] <http://www.adp-gmbh.ch/> , Oracle Tuning , Last Accessed On
January 2008.

[31]<http://www.csee.umbc.edu/help/Oracle8/server.815/a67781/c07procs.html> , oracle Tuning.

[32]http://www.idevelopment.info/data/Oracle/DBA_tips/Linux/LINUX_8.s_html#Configuring%20Shared%20Memory, Linux
Parameters, last Accessed on February 2008.

[33] <http://www.kernel.org/doc/Documentation/kernel-parameters.txt>
kernel parameters, Last Accessed on September 2008.

[34]

<http://www.Oracle.com/technology/tech/windows/rdbms/index.html> ,
Oracle RDBMS on windows last accessed on March 2009.

[35] <http://www.otn.Oracle.com> , Performance Tuning Guide,
Accessed on September 2007.

[36] www.adp-gmbh.ch/ora/misc/init_params.html , Initialization
parameters, Last Accessed on September 2007.

Appendixes

Appendixes A

Initializations parameters that can be set in the init.ora file

Initialization Parameters

The **initialization parameters** can be set in the init.ora file. There are two common ways to find out to what value an initialization parameter is set:

- show parameter <param_name> in sql*plus, or
- select value from v\$parameter where name = lower('param_name')

The parameters can be changed for the currently connected session with a alter session set ... command.

If a parameter should be set in another session, dbms system.set bool_param in session or dbms system.set int_param in session can be used.

Audit related parameters

audit file dest

audit syslog level

audit sys operations

audit trail

transaction auditing

NLS related parameters

nls_calendar

nls_comp

nls_characterset

nls_currency

nls_date_format

nls_date_language

nls_dual_currency

nls_iso_currency

nls_language

nls_length_semantics

nls_nchar_conv_excp

nls_numeric_characters

nls_sort

nls_territory

- nls_time_format
- nls_time_tz_format
- nls_timestamp_format
- nls_timestamp_tz_format

Optimizer related parameters

optimizer_dynamic_sampling

optimizer_features_enable

optimizer_index_caching

optimizer_index_cost_adj

optimizer_max_permutations

optimizer_mode

optimizer_percent_parallel

optimizer_secure_view_merging

Parameters affecting SGA

The following initialization parameters affect the size of the SGA:

db_block_buffers,

db_block_size,

db_cache_size,

db_keep_cache_size,

db_recycle_cache_size,

java_pool_size.

large_pool_size,

log_buffer,

shared_pool_size

streams_pool_size

Other parameters

ALLOW_FREELIST_GROUPS

This parameter was obsoleted after Oracle V6 and default to true since. At that time, it needed to be set in order to specify the freelist groups parameter in the storage clause.

ALWAYS_ANTI_JOIN

This parameter became obsolete in 9i.

BACKGROUND_DUMP_DEST

background_dump_dest specifies the directory (folder) where trace files of background processes are being written. It also specifies the location for the alert.log file.

It is also used for ORA-00600 errors.

BITMAP_MERGE_AREA_SIZE

BUFFER_POOL_KEEP

Deprecated in favour of db_keep_cache_size

BUFFER_POOL_RECYCLE

Deprecated in favour of db_recycle_cache_size

COMMIT_WRITE

Comes new with Oracle 10g R2, see also commit (sql) and On setting commit_write.

COMPATIBLE

Lowered afterwards.

The value of this parameter can be determined with dbms_utility.db_version.

The value of this parameter specifies the version that the database must adhere to.

With Oracle 10g, the value of this parameter must be set at least to **9.2**; and once it was set to **10** it cannot be

CONTROL_FILES

Every database must have at least one control file that describes important characteristics of the database. This parameter specifies their location.

CONTROL_FILE_RECORD_KEEP_TIME

This parameter controls the minimum number of days that a reusable record is kept in the control file.

Its range is 0... 365 (=1 year)

control_file_record_keep_time also governs the size of controlfiles

CURSOR_SHARING

This parameter influences hard parses and soft parses and is, according to metalink note 223299.1, one of the top parameters affecting performance.

The parameter can be set to either *exact*, *similar* or *force*.

DB_BLOCK_CHECKSUM

Specifies if integrity checking is enabled as block level.

DB_BLOCK_LRU_LATCHES

This parameter became obsolete in 9i.

DB_BLOCK_MAX_DIRTY_TARGET

This parameter became obsolete in 9i.

DB_CACHE_ADVICE

According to metalink note 223299.1, this is one of the top parameters affecting performance.

DB_CREATE_FILE_DEST

DB_CREATE_FILE_DEST sets the default location for Oracle-managed datafiles. This location is also used as the default for Oracle-managed control files and online redo logs if DB_CREATE_ONLINE_LOG_DEST_n is not specified. You can specify a file system directory as the default location for the creation of datafiles, control files, and online redo logs. However, the directory must already exist; Oracle does not create it. The directory must have appropriate permissions that allow Oracle to create files in it. Oracle generates unique names for the files, and a file thus created is an Oracle-managed file.

This parameter can be useful while creating a database.

DB_CREATE_ONLINE_LOG_DEST_n

DB_CREATE_ONLINE_LOG_DEST_n (where n= 1, 2, 3, ... 5) sets the default location for Oracle-managed control files and online redo logs. You should specify at least two parameters: DB_CREATE_ONLINE_LOG_DEST_1 and DB_CREATE_ONLINE_LOG_DEST_2. This provides greater fault tolerance for the logs if one of the destinations should fail.

If more than one directory is specified, then the control file or online redo log is multiplexed across the directories. One member of each online redo log is created in each directory, and one control file is created in each directory.

The directory must already exist; Oracle does not create it. The directory must have appropriate permissions that allow Oracle to create files in it. Oracle generates unique names for the files, and a file thus created is an Oracle-managed file.

This parameter can be useful while creating a database

DB_DOMAIN

DB_FILE_DIRECT_IO_COUNT

This parameter became obsolete in 9i.

DB_FILE_MULTIBLOCK_READ_COUNT

This parameter specifies how many blocks will be read at once when Oracle performs a full table scan or an *index range scan*. It doesn't affect reads on blocks that are indexed (in which case only one block is read).

The value for this parameter should be chosen carefully. The OS on which Oracle is running should be capable of reading $db_file_multiblock_read_count * db_block_size$ in one I/O request. If it is set too high, the optimizer will think that full table scan are cheap and will prefer them to the usage of indexes. On the other hand, setting it to low makes the optimizer choose indexes more often than necessary. By the way, the preference of indexes or full table scans is also influenced by optimizer_index_cost_adj.

DB_FILE_NAME_CONVERT

This parameter is needed if a standby database does not have the same layout on the disk for its files as the primary database.

DB_FILES

The maximum number of database files that can be opened for a database.

DB_FLASHBACK_RETENTION_TARGET

This is one of the relevant parameters for Flashback DB.

DB_NAME

This parameter must have the same value as the database name.

DB_RECOVERY_FILE_DEST

This is one of the relevant parameters for Flashback DB.

DB_RECOVERY_FILE_DEST_SIZE

This is one of the relevant parameters for Flashback DB.

DB_WRITER_IO_SLAVES

db_writer_io_slaves simulates asynchronous IO, but they do not perform asynchronous IO, and thus, they're only meaningful if the OS does not support asynchronous IO.

If the OS supports asynchronous, multiple dbwr processes should be used and disk_async_io be set to true.

DB_16K_CACHE_SIZE

According to metalink note **223299.1**, this is one of the top parameters affecting performance.

DB_2K_CACHE_SIZE

According to metalink note 223299.1, this is one of the top parameters affecting performance.

DB_32K_CACHE_SIZE

According to [metalink](#) note 223299.1, this is one of the top parameters affecting performance.

DB_4K_CACHE_SIZE

According to [metalink](#) note 223299.1, this is one of the top parameters affecting performance.

DB_8K_CACHE_SIZE

According to [metalink](#) note 223299.1, this is one of the top parameters affecting performance.

EVENT

, anchor=>'event')

name action

This parameter allows to set a diagnostic event.

Multiple events must be separated by colons:

event="<event 1>:<event 2>: <event 3>: <event n>"

GC_DEFER_TIME

This parameter became obsolete in 9i.

GC_RELEASABLE_LOCKS

This parameter became obsolete in 9i.

GC_ROLLBACK_LOCKS

This parameter became obsolete in 9i

GLOBAL_NAMES

HASH_AREA_SIZE

The default is 64K, which is far too small for most cases. A range of 512KB to 1MB should be considered.

The memory for a hash join (up to the value specified with hash_area_size) is allocated from the cursor work heap (in the uga.)

HASH_MULTIBLOCK_IO_COUNT

This parameter became obsolete in 9i.

INSTANCE_NAME

INSTANCE_NUMBER

INSTANCE_NODESET

This parameter became obsolete in 9i.

JOB_QUEUE_INTERVAL

This parameter became obsolete in 9i.

JOB_QUEUE_PROCESSES

Controls how many jobs can run.

LM_LOCK

This parameter became obsolete in 9i.

LM_RESS

This parameter became obsolete in 9i.

LOCK_NAME_SPACE

LOCK_SGA

On platform that support it, this parameter can be set to true which will lock the entire SGA into physical memory.

LOG_ARCHIVE_DEST

Deprectated in Enterprise Edition in favour of log_archive_dest_n.

LOG_ARCHIVE_DEST_n

LOG_ARCHIVE_DEST_n (as well as log_archive_dest) can only be used if the database is running in archive log mode.

A common mistake when moving from the (deprecated) log_archive_dest to log_archive_dest_n is to forget one of the attributes such as SERVICE= or LOCATION= which causes a ORA-16179: incremental changes to "log_archive_dest_1" not allowed with SPFILE when it altered with the alter system command.

Attributes:

SERVICE

A standby destination.

Use the lgwr option to specify LGWR transmission or the arch option to specify ARCH transmission.

- LOCATION
A local file system path, must be defined at least once.
- DELAY=minutes
delays applying of the redo log at the standby site.

varchive_dest_status allows to query the status (and possibly the errors) for each of the defined archive destinations.

LOG_ARCHIVE_DEST_STATE_n

log_archive_dest_state_N specifies the state for log_archive_dest N.

LOG_ARCHIVE_FORMAT

The following *expandables* can be used:

- %s: log sequence number
- %S: log sequence number, zero filled
- %t: thread number
- %T: thread number, zero filled
- %d: DBID

LOG_ARCHIVE_START

This parameter is deprecated in Oracle 10g

This parameter determines if the background process ARCH is started. It can be set to either true or false.

Of course, it makes no sense, if this parameter is set to true if the database is running in noarchive log mode. If ARCH is started with the database being in noarchive log mode, messages like media recovery disabled will be written into the alert.log file.

LOG_CHECKPOINT_INTERVAL

The unit of this parameter is measured in **physical operating system blocks**, not DB blocks. The operating system block size is (obviously) OS dependent. It can be retrieved through x\$kccle.

LOG_FILE_NAME_CONVERT

This parameter is needed if a standby database does not have the same layout on the disk for its files as the primary database.

FIXED_DATE

Fixed date can be set to a date in the following format:
YYYY-MM-DD HH24-MI-SS

If set, sysdate returns this date instead of the current date.

```
alter session set nls_date_format = 'dd.mon.yyyy hh24:mi:ss';
```

```
alter system set fixed_date=2004-03-02 22:23:24;
```

```
Select sysdate from dual;
```

```
SYSDATE
```

```
02.MAR.2004 22:23:24
```

MAX_DUMP_FILE_SIZE

This parameter specifies the maximum size for dump files such as trace files.

The unit of this parameter is measured in **physical operating system blocks** unless it has a suffix **M** or **K**, in which case the unit is Megabyte and Kilobyte, respectively. Note, the size of physical

operating system blocks is not equal to the size of DB blocks.
The operating system block size is (obviously) OS dependent. It can be retrieved through x\$kccl.

MAX_IDLE_TIME

O7_DICTIONARY_ACCESSIBILITY

Default was true until 8i, and is false since 9i.

false: only privileged users can access the data dictionary.

true: any user who has been granted *select any table* can select from tables owned by sys. Alternatively, select_catalog_role can be granted.

The parameter should (probably) be set to false. Users that need access to sys owned table should then be granted the select any dictionary privilege.

The setting of this parameter influences grant ... ANY .. to ... statements.

OPEN_CURSORS

This parameter defines how many cursors a *session* (not the cumulative sum of all sessions) can open at most.

PGA_AGGREGATE_TARGET

According to metalink note **223299.1**, this is one of the top parameters affecting performance.

PROCESSES

The value of processes affects the value that the kernel parameter SEMMSL (Maximum number of semaphores in a semaphore set): it should be equal to the value of processes + 10.

If there are more than one instance on a box, the value of the instance with the greatest processes must be taken.

It affects also the optimal setting for SEMMNS (Number of semaphores in the system): $2 * \text{highets process value} + 1 * \text{other process values} + 10 * \text{count of instances}$.

QUERY_REWRITE_ENABLED

This parameter must be set to *true* to make use of function based indexes. Additionally query_rewrite_integrity must be set to *trusted*.

QUERY_REWRITE_INTEGRITY

This parameter can be set to either

- enforced
- trusted
- stale_tolerated

This parameter must be set to *trusted* to make use of function based indexes. Additionally query_rewrite_enabled must be set to *true*.

REMOTE_ARCHIVE_ENABLE

REMOTE_LISTENER

REMOTE_LOGIN_PASSWORDFILE

`remote_login_passwordfile` specifies if Oracle checks for a password file and if this password file is shared among databases.

The following values are possible:

- **none**
Oracle ignores the password file if it exists.
- **exclusive**
Password file is exclusively used by one database. Any user can be added to the password file.
- **internal**
Used for Oracle Parallel Server
- **shared**
The password file is shared among databases. However, the only users that can be authenticated are sys (and obsoletly: internal).
If the password file is shared, only SYS can be added to the password file.

RESOURCE_LIMIT

This parameter must be set to true to enforce resource limits assigned to a user through profiles

RESOURCE_MANAGER_PLAN

Setting this parameter activates the resource manager.

If the parameter is set with a prepending FORCE:, the plan can only be changed by the database administrator.

ROLLBACK_SEGMENTS

Defines the rollback segments that the instance will acquire at startup

On startup, Oracle divides transactions by

transactions_per_rollback_segment. If the result is greater than the number of rollback segments actually brought online by the `rollback_segments` init param, additional rollback segments will be brought online.

SESSIONS

SESSION_CACHED_CURSORS

SHARED_SERVERS

SORT_AREA_SIZE

The default is 64K, which is far too small for most cases. A range of 512KB to 1MB should be considered.

The memory for a sort (up to the value specified with `sort_area_size`) is allocated from the cursor work heap (in the uga).

SORT_AREA_RETAINED_SIZE

SORT_MULTIBLOCK_READ_COUNT

This parameter became obsolete in 9i.

SPFILE

Specifies the spfile to be used.

STANDBY_ARCHIVE_DEST

This parameter specifies the location of archived redo logs that come from a primary database.

The value of this parameter is displayed in the v\$archive_dest view.

STAR_TRANSFORMATION_ENABLED

STANDBY_FILE_MANAGEMENT

If not set to auto, newly created tablespaces in a standby environment must be recreated manually on the standby servers as well. Similarly, newly added datafiles must be copied to the standby servers as well.

STATISTICS_LEVEL

According to metalink note 223299.1, this is one of the top parameters affecting performance.

It can be set to one of

- ALL
- TYPICAL
- BASIC

SQL_TRACE

Setting `sql_trace=true` is a prerequisite for using tkprof. It can also be set for a single session with alter session set sql_trace.

After setting `sql_trace` to true, a trace file will be written.

There is also dbms_support that should allow to trace sessions with more information.

`sql_trace` seems to be deprecated since 10.2, but not removed. It still behaves as in earlier versions of Oracle.

TEXT_ENABLE

This parameter became obsolete in 9i.

TIMED_STATISTICS

This parameter must be true in order to gather timing information in v\$system_event

It is also useful when using tk_prof.

TRANSACTIONS

TRANSACTIONS_PER_ROLLBACK_SEGMENT

UNDO_MANAGEMENT

Set to AUTO to use Oracle 9i's new automatic undo management.

UNDO_RETENTION

Specifies for how many seconds undo information is kept.

UNDO_SUPPRESS_ERRORS

This parameter is important if

An Oracle database is upgraded to version 9i,

the upgraded database uses Undo Tablespaces,

There are still applications that use SET TRANSACTION USE

ROLLBACK SEGMENT

If in such a case the parameter is set to TRUE (default is FALSE), there won't be any errors; although it gets written into the alert log.

UNDO_TABLESPACE

Specifies the undo tablespaces when using automatic undo management.

USER_DUMP_DEST

The value of user_dump_dest specifies the destination (path to a operating system directory) where user processes will write trace files.

It is also used for ORA-00600 errors.

USE_POST_WAIT_DRIVER

Setting this value to true makes Oracle use post-wait drivers instead of semaphores.

UTL_FILE_DIR

This parameter specifies one more more locations to where files can be written and from where files can be read using utl file.

Specifying multiple directories in the spfile:

alter system set

utl_file_dir='/foo/bar/dir1','/foo/baz/dir2','/tmp' scope=spfile

WORKAREA_SIZE_POLICY

According to metalink note 223299.1, this is one of the top parameters affecting performance.

Hidden parameters

Parameters whose name starts with an underscore are hidden.

Usually, they should not be touched! Oracle won't probably support the database if one of these parameters were changed.

_ALLOW_RESETLOGS_CORRUPTION

Allows resetlogs even if it will cause corruption.

_COLUMN_TRACKING_LEVEL

If set to 1 (the default), will cause SMON to update

sys.col_usage\$ with information regarding access patterns on table columns.

`_DB_AGING_COOL_COUNT`

Touch count set when buffer cooled.

`_DB_AGING_FREEZE_CR`

Make CR buffers always be too cold to keep in cache.

`DB_AGING_HOT_CRITERIA`

Touch count which sends a buffer to head of replacement list.

`DB_AGING_STAY_COUNT`

Touch count set when buffer moved to head of replacement list.

`_DB_AGING_TOUCH_TIME`

This parameter specifies a time period in which the touch count of a buffer within the buffer cache can at most be increased once.

`_DB_PERCENT_HOT_DEFAULT`

Percent of default buffer pool considered hot.

`_DB_PERCENT_HOT_KEEP`

Percent keep buffer pool considered hot.

`_DB_PERCENT_HOT_RECYCLE`

Percent recycle buffer pool considered hot.

`_INIT_SQL_FILE`

This parameter points to the file that is executed upon creation of the database (create database). As of 9i, the value is `*/rdbms/admin/sql.bsq`.

`_KGHDSIDX_COUNT`

Controls the number of shared area subpools.

`_LOG_IO_SIZE`

The unit of this parameter is measured in physical operating system blocks, not DB blocks. The operating system block size is (obviously) OS dependent. It can be retrieved through x\$kccl.

`_REALFREE_HEAP_PAGESIZE_HINT`

`_RECYCLEBIN`

This hidden parameter is available in 10g. If set to false, then tables are purged immediately at a *drop table*.

`_SMALL_TABLE_THRESHOLD`

`_SYSTEM_TRIG_ENABLED`

Defaults to true and Oracle recommends setting it to false only during database upgrade.

If this parameter is set to false, then **system triggers** won't be executed.

`_TRACE_FILES_PUBLIC`

Trace files (such as those created by a block dump are only readable by oracle, unless this parameter is set to true.

Setting this parameter to true causes a security risk as sensitive data might be written into the trace files!

`_USE_ISM`

If a system features ISM (intimate shared memory), Oracle uses it by default. This can be disabled by setting `_use_ism` to false.

As far as I can see, solaris is the only OS that has ISM.

`_USE_ISM_FOR_PGA`

`_WAIT_FOR_SYNC`

If set to false, a transaction that commits does not wait until the redo is flushed. However, a database can then not be restored if it crashes.

Appendix B

Popular Kernel Parameter

ACPI	ACPI support is enabled.
ALSA	ALSA sound support is enabled.
APIC	APIC support is enabled.
APM	Advanced Power Management support is enabled.
AX25	Appropriate AX.25 support is enabled.
CD	Appropriate CD support is enabled.
DEVFS	devfs support is enabled.
DRM	Direct Rendering Management support is enabled.
EDD	BIOS Enhanced Disk Drive Services (EDD) is enabled
EFI	EFI Partitioning (GPT) is enabled
EIDE	EIDE/ATAPI support is enabled.
FB	The frame buffer device is enabled.
HW	Appropriate hardware is enabled.
IA-32	IA-32 aka i386 architecture is enabled.
IA-64	IA-64 architecture is enabled.
IOSCHED	More than one I/O scheduler is enabled.
ISAPNP	ISA PnP code is enabled.
ISDN	Appropriate ISDN support is enabled.
JOY	Appropriate joystick support is enabled.
LP	Printer support is enabled.

LOOP	Loopback device support is enabled.
M68k	M68k architecture is enabled.
MCA	MCA bus support is enabled.
MDA	MDA console support is enabled.
MOUSE	Appropriate mouse support is enabled.
MSI	Message Signaled Interrupts (PCI).
MTD	MTD support is enabled.
NET	Appropriate network support is enabled.
NUMA	NUMA support is enabled.
NFS	Appropriate NFS support is enabled.
OSS	OSS sound support is enabled.
PARIDE	The ParIDE subsystem is enabled.
PARISC	The PA-RISC architecture is enabled.
PCI	PCI bus support is enabled.
PCMCIA	The PCMCIA subsystem is enabled.
PNP	Plug & Play support is enabled.
PPC	PowerPC architecture is enabled.
PPT	Parallel port support is enabled.
PS2	Appropriate PS/2 support is enabled.
RAM	RAM disk support is enabled.
S390	S390 architecture is enabled.
SCSI	Appropriate SCSI support is enabled.
SELINUX	SELinux support is enabled.

SERIAL	Serial support is enabled.
SMP	The kernel is an SMP kernel.
SPARC	Sparc architecture is enabled.
SWSUSP	Software suspend is enabled.
TS	Appropriate touchscreen support is enabled.
USB	USB support is enabled.
USBHID	USB Human Interface Device support is enabled.
V4L	Video For Linux support is enabled.
VGA	The VGA console has been enabled.
VT	Virtual terminal support is enabled.
WDT	Watchdog support is enabled.
XT	IBM PC/XT MFM hard disk support is enabled.
KNL	Is a kernel start-up parameter.
BOOT	Is a boot loader parameter.

