



Craig S. Mullins & Associates, Inc.

Database Performance Management

[Return to Home Page](#)

Summer 1995



An Introduction to the Architecture of Oracle

By Craig S. Mullins and Christopher T. Foot

Editor's note: As more sites establish multiple DBMS environments, developers and analysts need to understand the strengths and weaknesses of various client/server DBMS offerings. This article is third in a series on "Client/Server relational DBMS Architectures;" here, Christopher T. Foot and Craig S. Mullins examine the architectural "guts" of Oracle.

Introduction

To effectively administer the Oracle Server environment, you will find it worthwhile to examine the architecture of Oracle in more detail. Seeing "what makes Oracle tick" can best be achieved by

investigating the basic "architectural" blueprint of Oracle and examining its five basic components:

- File structures
- Memory structures
- Processes
- Rollback segments
- Redo logs

These core components operate in an integrated manner to provide support for the client tasks.

Instance vs. Database

An Oracle *instance* is the combination of all of the memory structures and background processes that are allocated when an Oracle *database* is started. Oracle users frequently confuse an Oracle *instance* with an Oracle *database*.

An Oracle *database* has both physical structures (data files) and logical structures (table, index). Oracle separates the physical from the logical structure so that the physical storage of data may be manipulated without affecting user access to the logical structures. The physical structure of an Oracle database is determined by the files created at the operating system level by the Oracle instance during database creation (controlfile, logfile) or by the DBA during normal operation (create table space, create control file).

The three physical file structures that comprise an Oracle database are:

Control files — a small administrative file that is used by the Oracle database

Redo log files — a record of changes made to data. The redo log files are used in recovery circumstances to ensure that no data is lost should a failure prevent changed data from being written to disk.

Database files — a file that contains the database information, including both system and application data.

Oracle parameter files

An Oracle *parameter file* contains all of the configuration parameters for an Oracle instance. These parameters can be set to different values to tailor the Oracle instance to the system's operating configuration as well as to the needs of the individual application or applications contained in the Oracle database. The parameter file specifications affect both memory and process settings of the associated Oracle *instance*. All parameters contained in the parameter file are optional and will default if omitted (see Oracle's *Oracle 7 Server for UNIX Administrator's Reference*, Oracle Corporation part no. A10324-1, for default values.)

Oracle database files

Oracle *database files* contain the data associated with a particular database. All of the files discussed are not absolutely required for normal database operations, but this configuration is highly recommended for a well-designed, efficient environment. Oracle database files can be grouped according to specific categories:

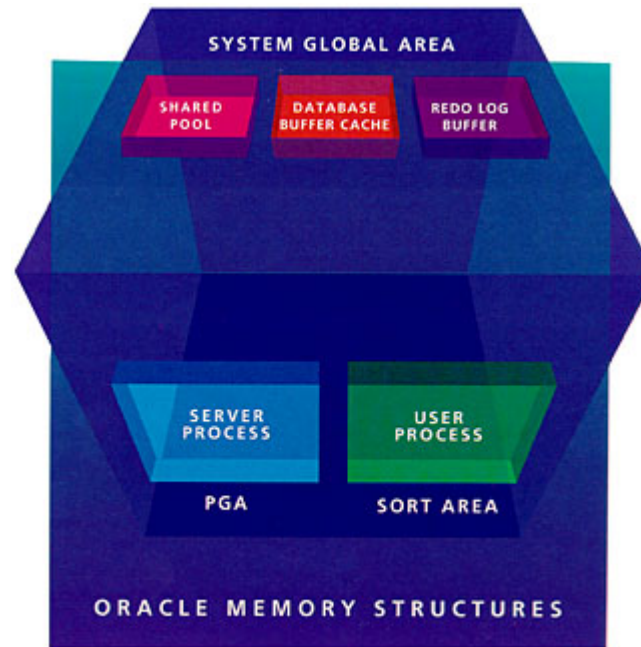
- *Control files* record the physical structure of the Oracle database. The control file is read during every database startup. It is stored in binary format and cannot be edited by normal means, but it will be automatically modified when database changes are made.
- *Data files* associated with oracle table spaces include:
 - *System data files* — these contain the contents of the system table space. The system table space consists of the data dictionary, and the names and locations of all the tablespaces, tables, indexes, triggers, procedures, and clusters for the database.
 - *Application data files and default data files* are physical database structures. These files hold *user* data that is accessed by clients in the course of normal database operation. The difference between default and application data files is that the default table space/data file combinations used when an explicit specification is not made.

- *Temporary files* — these contain temporary data that is used by queries and sorts. Regardless of their name, temporary data files are actually permanent files created to provide a transient work area for Oracle to sort and join tables and to execute complex SQL.
- Oracle *rollback files* — these contain information related to the rolling back of uncommitted data during error processing.
- Redo log files that record changes made to data. Every Oracle database has at least one set of two or more redo log files. These files record changes made to data and are used in recovery circumstances to prevent lost or damaged data.
- The config.ora file is associated with the Oracle client, not the server. It specifies certain defaults, file and directory names for the client portion of the Oracle environment.

Oracle memory structures

Oracle utilizes specific memory structures to perform DBMS-related tasks. These memory structures are contained in the main memory (RAM) of the computer running the Oracle instance.

The basic memory structures for an Oracle instance are the SGA (System Global Area), PGA (Program Global Area), and Oracle Sort Area.



The SGA is a group of shared memory structures allocated by the Oracle *instance* to contain data and control information. The SGA is automatically allocated when the instance is started and deallocated when the instance is shut down. If more than one user is connected to a multiple-process database, the information in the SGA is shared among the different users.

The SGA contains the following components:

Database buffer cache — a staging area for database reads and writes, this cache holds the most recently used data blocks.

Redo log buffer — a staging area for redo log entries, similar to the database buffer cache.

Shared pool — an area of memory set aside for the parsing and processing of SQL statements (shared SQL area), staging of Data Dictionary reads and writes (Data Dictionary cache), and, if the Oracle configuration is multithreaded, binding information and run-time buffers (*private SQL area*).

The PGA is a memory structure containing information for both user and background processes. Each process has its own PGA. Various processes can use the PGA as an internal "scratch pad" or work area. The contents of the PGA vary depending on the type of process (user or background) and Oracle configuration (multithreaded server or conventional server).

Every time a user process requests a sort to be performed a special memory structure is allocated called an *Oracle sort area*. The sort area exists in the memory of the user process that is requesting the sort. User processes typically perform sorting when query result sets are ordered or grouped.

Oracle processes

Oracle processes are the true workhorses of the Oracle instance. Each process is composed of a series of tasks. Also, each process has its own internal memory allocated to it (the PGA). This

memory allocation allows the process to perform the internal processing required to carry out its designated function.

Oracle has two general types of processes: user processes and Oracle processes. A *user process* is created to execute the program code of an application program. An *Oracle process* is called by another process to perform certain, specific functions on behalf of the invoking process.

Oracle processes can be further broken down into server processes and background processes:

- *Server processes* are created to interact between the user processes and the Oracle instance. The server process communicates with the user process and acts as a 'relay' between the user process and SGA information.
- *Background processes* perform designated functions that would otherwise have to be performed by each user process that interacts with the Oracle instance. This specialization of functions is much more efficient than if each user process was required to perform a variety of tasks. Oracle uses the following background processes:
 - The *PMON* (Process Monitory) background process performs 'cleanup duties' when a user process fails with an error condition.

PMON cleans up the cache, releases locks, and other miscellaneous tasks.

- The *SMON* (System Monitor) background process provides instance recovery during startup. SMON also cleans up temporary segments that are no longer in use, compresses free space into contiguous extents and, in a parallel server environment, provides instance recovery for a failed CPU.
- The *DBWR* (Database Writer) writes database buffers contained in cache to disk.
- The *LGWR* (Log Writer) is a background process that manages the redo log buffer.
- If an Oracle instance takes more checkpoints than the LGWR process can efficiently handle, checkpoint duties can be turned over to the *CKPT* (Checkpoint) background process.
- The *ARCH* (Archiver) process performs log archival by copying online redo log files to auxiliary storage devices as logs become full.
- The *RECO* (Recover) automatically resolves failures involving distributed transactions.
- The final types of processes are *user* and *server* processes. The user process provides

the communication link between the client application and the Oracle instance. The user process does not execute SQL statements. Those duties fall to the server process. The server process parses and executes SQL and performs all the tasks required to send results back to the requesting application.

Oracle configurations

Two levels of software must be executed by users accessing the Oracle server. The application or software tool on the client must be executed to initiate the connection to the Oracle server and transport the desired SQL statement to be executed. The server process on the Oracle server receives the connection request and its associated SQL statement and begins the process of executing the desired statements.

The Oracle server can be configured as either a single or multiple process instance. The different configurations are based on the operating system on which the instance will run and the processing requirements of the applications accessing Oracle data. Some operating systems (such as MS-DOS) cannot run multiple background processes concurrently. All of the code for the various background tasks required by the instance must be run in a single background process. These types of databases can be accessed by only one user at a time. This requires a single process Oracle instance.

The most popular class of configuration, however, is

the multiple process instance. There are three possible multiple process configurations:

- Combined user/server process-some operating systems (such as VAX VMS) can simultaneously execute the server process and the *application's* code in a single executable user process. The communication between the two processes running as one is done through the *program interface*. Although configured as a multiple process instance; only one Oracle connection at a time is allowed. Multiple user, concurrent access is not available.
- Dedicated server process-in this setup, the background runs separately. This is the most popular of all Oracle configurations. It is called the dedicated server process because each server process act on behalf of only one user process.
- Multithreaded server process-this process differs from the dedicated server process because several user processes have the capability of sharing one server process. A dedicated and shared server process perform the same functions.

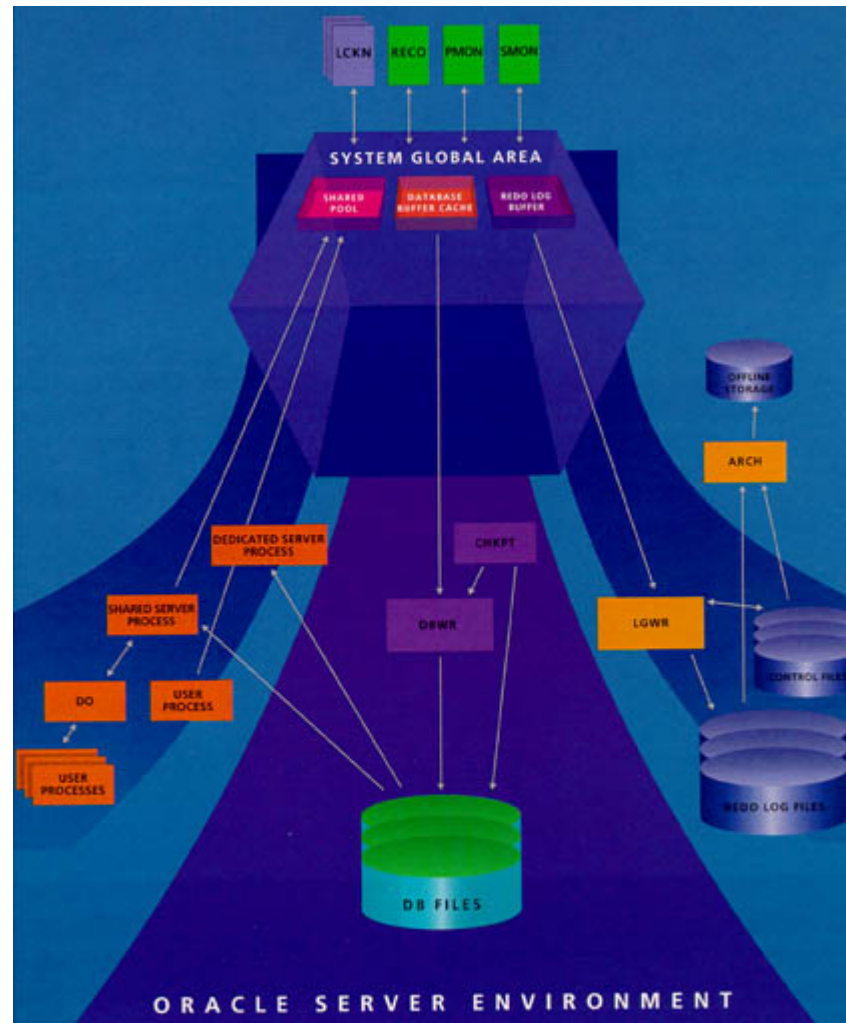


Figure 2 offers a complete picture of Oracle "under the covers." It contains all the components of Oracle that operate together to achieve an effective and useful relational database management system.

© 2005 Mullins Consulting, Inc. All rights reserved.

[Home](#). Phone: 281-494-6153 Fax: 281-491-0637