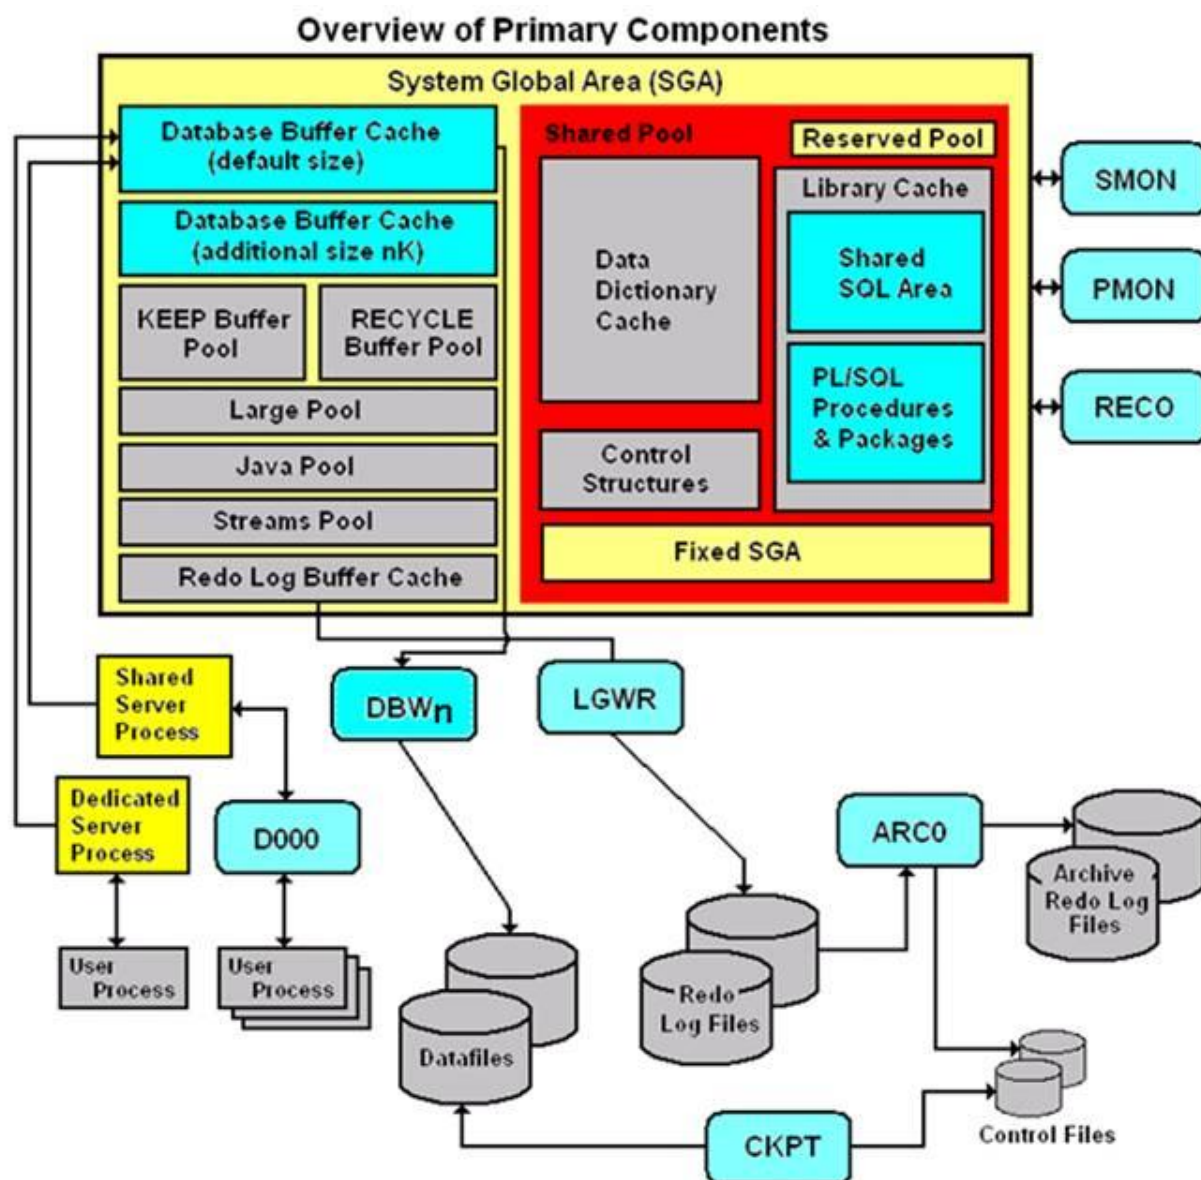


Oracle Architecture

Objectives

These notes introduce the Oracle server architecture. The architecture includes physical components, memory components, processes, and logical structures.

Primary Architecture Components



The figure shown above details the Oracle architecture.

Oracle server: An Oracle server includes an **Oracle Instance** and an **Oracle database**.

- An Oracle database includes several different types of files: datafiles, control files, redo log files and archive redo log

files. The Oracle server also accesses parameter files and password files.

- This set of files has several purposes.
 - One is to enable system users to process SQL statements.
 - Another is to improve system performance.
 - Still another is to ensure the database can be recovered if there is a software/hardware failure.
- The database server must manage large amounts of data in a multi-user environment.
- The server must manage concurrent access to the same data.
- The server must deliver high performance. This generally means fast response times.

Oracle instance: An Oracle Instance consists of **two** different sets of components:

- The first component set is the set of **background processes** (PMON, SMON, RECO, DBW0, LGWR, CKPT, D000 and others).
 - These will be covered later in detail – each background process is a computer program.
 - These processes perform input/output and monitor other Oracle processes to provide good performance and database reliability.
- The second component set includes the **memory structures** that comprise the Oracle instance.
 - When an instance starts up, a memory structure called the System Global Area (SGA) is allocated.
 - At this point the background processes also start.
- An Oracle Instance provides access to one and only one Oracle database.

Oracle database: An Oracle database consists of files.

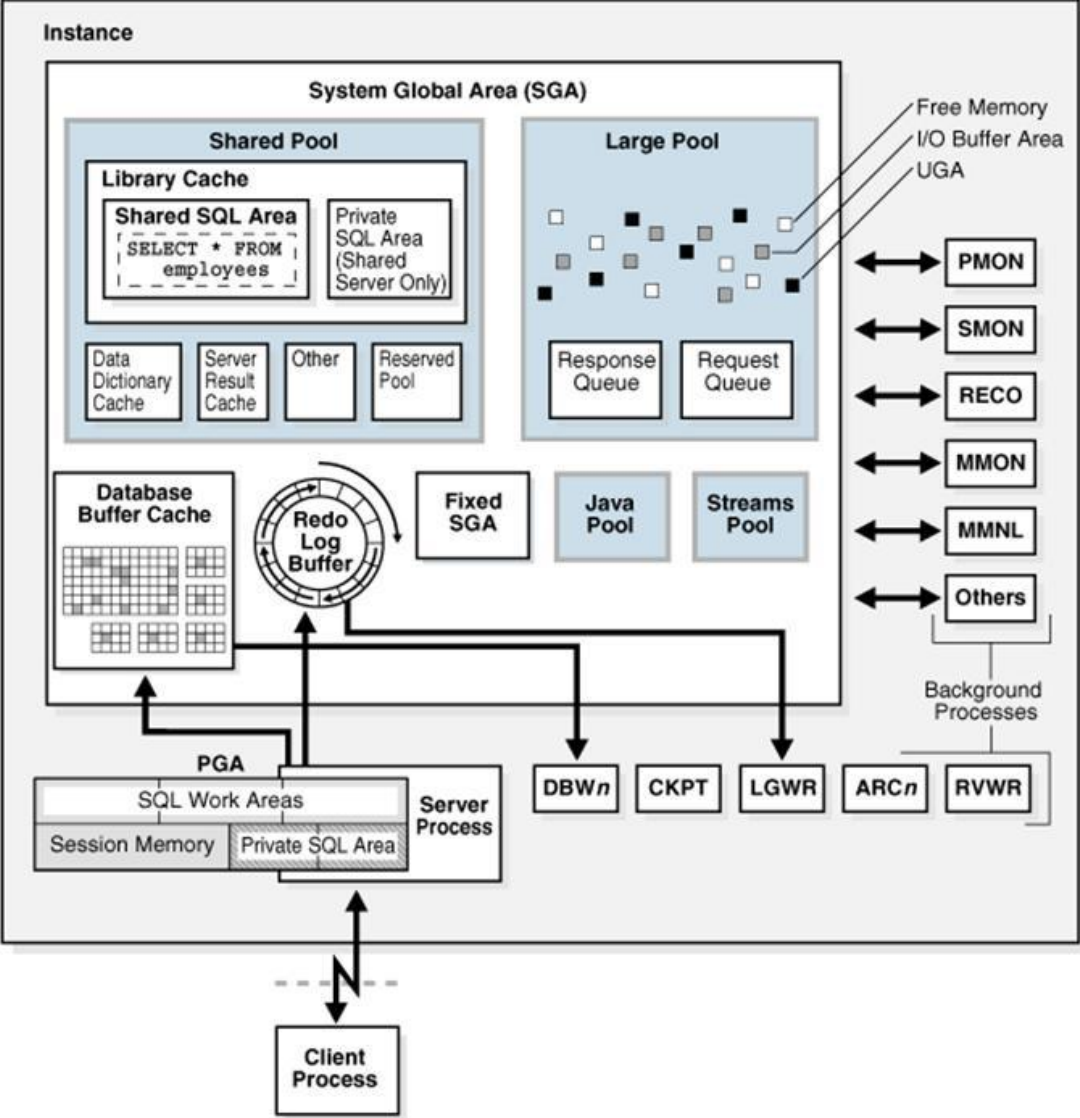
- Sometimes these are referred to as operating system files, but they are actually **database files** that store the database information that a firm or organization needs in order to operate.
- The **redo log files** are used to recover the database in the event of application program failures, instance failures and other minor failures.

- The **archived redo log files** are used to recover the database if a disk fails.
- Other files not shown in the figure include:
 - The required **parameter file** that is used to specify parameters for configuring an Oracle instance when it starts up.
 - The optional **password file** authenticates special users of the database – these are termed **privileged users** and include database administrators.
 - **Alert** and **Trace Log Files** – these files store information about errors and actions taken that affect the configuration of the database.

User and server processes: The processes shown in the figure are called **user** and **server processes**. These processes are used to manage the execution of SQL statements.

- A **Shared Server Process** can share memory and variable processing for multiple user processes.
- A **Dedicated Server Process** manages memory and variables for a single user process.


This figure from the *Oracle Database Administration Guide* provides another way of viewing the **SGA**.



Connecting to an Oracle Instance – Creating a Session

Connecting to an Oracle Instance:

- Establishing a user connection
- Creating a session



System users can connect to an Oracle database through SQLPlus or through an application program like the Internet Developer Suite (the program becomes the system user). This connection enables users to execute SQL statements.

The act of connecting creates a communication pathway between a user process and an Oracle Server. As is shown in the figure above, the User Process communicates with the Oracle Server through a Server Process. The User Process executes on the client computer. The Server Process executes on the server computer, and actually executes SQL statements submitted by the system user.

The figure shows a one-to-one correspondence between the User and Server Processes. This is called a **Dedicated Server** connection. An alternative configuration is to use a **Shared Server** where more than one User Process shares a Server Process.

Sessions: When a user connects to an Oracle server, this is termed a session. The **User Global Area** is session memory and these memory structures are described later in this document. The session starts when the Oracle server validates the user for connection. The session ends

when the user logs out (disconnects) or if the connection terminates abnormally (network failure or client computer failure).

A user can typically have more than one concurrent session, e.g., the user may connect using SQLPlus and also connect using Internet Developer Suite tools at the same time. The limit of concurrent session connections is controlled by the DBA.