



UTM
UNIVERSITI TEKNOLOGI MALAYSIA

**DATABASE ADMINISTRATION
SECP 3713**

**ASSIGNMENT 3
ORACLE ARCHITECTURE**

**SECTION 01
DR HASLINA HASHIM**

NAME	MADINA SURAYA BINTI ZHARIN
MATRIC NO.	A20EC0203

Oracle Database Memory Structure

1. System Global Area (SGA)

SGA is the memory area that contains data and control information for instance. The SGA includes the following data structures:

1) **Database buffer cache:** Caches blocks of data retrieved from the database

Before data stored in the database can be queried or modified, it must be read from a disk and stored in the buffer cache. All user processes connected to the database share access to the buffer cache. For optimal performance, the buffer cache should be large enough to avoid frequent disk I/O operations.

2) **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

The redo buffer is where data that needs to be written to the online redo logs will be cached temporarily before it is written to disk, this area is normally less than a couple of megabytes in size. These entries contain necessary information to reconstruct or redo changes by the INSERT, UPDATE, DELETE, CREATE, ALTER, and DROP commands.

The contents of this buffer are flushed:

- Every three seconds
- Whenever someone commits a transaction
- When it gets one third full or contains 1MB of cached redo log data.
- When LGWR is asked to switch logs
- Use the parameter LOG_BUFFER parameter to adjust but be careful increasing it too large as it will reduce your I/O but commits will take longer.

3) **Shared pool:** Caches various constructs that can be shared among users

The shared pool consists of the following areas:

- **Library cache** includes the shared SQL area, private SQL areas, PL/SQL procedures and packages the control structures such as locks and library cache handles. Oracle code is first parsed, then executed, this parsed code is stored in the library cache, oracle first checks the library cache to see if there is an already parsed and ready to execute form of the statement in there, if there is this will reduce CPU time considerably, this is called a soft parse, If Oracle has to parse it then this is called a hard parse. If there is not enough room in the cache oracle will remove older parsed code, obviously it is better to keep as much parsed code in the library cache as

possible. Keep an eye on missed cache hits which is an indication that a lot of hard parsing is going on.

- **Dictionary cache** is a collection of database tables and views containing information about the database, its structures, privileges, and users. When statements are issued the oracle will check permissions, access, etc. and will obtain this information from its dictionary cache, if the information is not in the cache, then it must be read in from the disk and placed into the cache. The more information held in the cache the less the oracle must access the slow disks.
- 4) **Large pool:** Optional area that provides large memory allocations for certain large processes, such as Oracle backup and recovery operations, and I/O server processes. The large pool is basically a non-cached version of the shared pool.

This is an optional memory area that provide large areas of memory for:

- **Shared Server** - to allocate the UGA region in the SGA
 - **Parallel execution of statements** - to allow for the allocation of inter-processing message buffers, used to coordinate the parallel query servers.
 - **Backup** - for RMAN disk I/O buffers
- 5) **Java pool:** Used for all session-specific Java code and data in the Java Virtual Machine (JVM). It is used to execute java code within the database.
- 6) **Streams pool:** Used by Oracle Streams to store information required by capture and apply. Streams are used for enabling data sharing between databases or application environments.

2. Program Global Areas (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 nonshared PGA when the process is started. 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.

Background Process	Purpose
Database Writer (DBWn)	The database writer writes modified blocks from the database buffer cache to the files on disk. Oracle allows a maximum of 20 database writer processes.
Log Writer (LGWR)	The log writer process writes redo log entries to disk. Redo log entries are generated in the redo log buffer of the SGA and the log writer process writes the redo log entries sequentially into an online redo log file.

Background Process	Purpose
System Monitor (SMON)	The system monitor performs crash recovery when a failed instance starts up again.
Process Monitor (PMON)	The process monitor performs recovery when a user process fails. It cleans up the cache and frees resources that the failed process was using.
Recoverer process (RECO)	The recoverer process (RECO) is a background process used with the distributed database configuration that automatically resolves failures involving distributed transactions. The RECO process of a node automatically connects to other databases involved in an in-doubt distributed transaction. When the RECO process re-establishes a connection between involved database servers, it automatically resolves all in-doubt transactions, removing from each database's pending transaction table any rows that correspond to the resolved in-doubt transactions.
Job Queue	Job queue processes are used for batch processing. They run user jobs. They can be viewed as a scheduler service that can be used to schedule jobs as PL/SQL statements or procedures on an Oracle instance. Given a start date and an interval, the job queue processes try to run the job at the next occurrence of the interval.
Archiver (ARCn)	Archiver processes copy the redo log files to archival storage when the log files are full, or a log switch occurs. The database must be in archive log mode to run archive process.
Queue monitor (QMNn)	The QMNn processes are optional background processes used by Oracle Streams Advanced Queueing (AQ), Streams and a variety of other Database products which monitor and maintain all the system and user owned AQ persistent and buffered objects. They provide the mechanism for message expiration, retry, and delay, maintaining queue statistics, removing processed messages from the queue table, and maintaining the dequeue IOT. They also handle all the supported buffered message operations.