



Exploring the Database Architecture


- Mari Kupatadze
- Senior Solutions Architect – Flashgrid
- Oracle Certified Master
- Oracle ACE ♠️



♠ About Me

- Senior Solutions Architect at Flashgrid
- Oracle Certified Master (OCM)
- Oracle ACE
- Co-writer, speaker, blogger

 mariam.kupa@gmail.com

 [mariam_kupa](https://twitter.com/mariam_kupa)

 dba010.com



500+ Technical Experts Helping Peers Globally



ORACLE
ACE Director



ORACLE
ACE



ORACLE
ACE Associate

3 Membership Tiers:

- Oracle ACE Director
- Oracle ACE
- Oracle ACE Associate

bit.ly/OracleACEProgram

Connect:

✉ oracle-ace_ww@oracle.com

f Facebook.com/oracleaces

🐦 [@oracleace](https://twitter.com/oracleace)

Objectives

- Major components of Oracle database
- Memory structures
- Background processes
- Logical and Physical structures

Disclaimer

“The postings on this document are my own and don’t necessarily represent my actual employer positions, strategies or opinions. The information here was edited to be useful for general purpose, specific data and identifications were removed to allow reach the generic audience and to be useful for the community.”

Major components of Oracle database

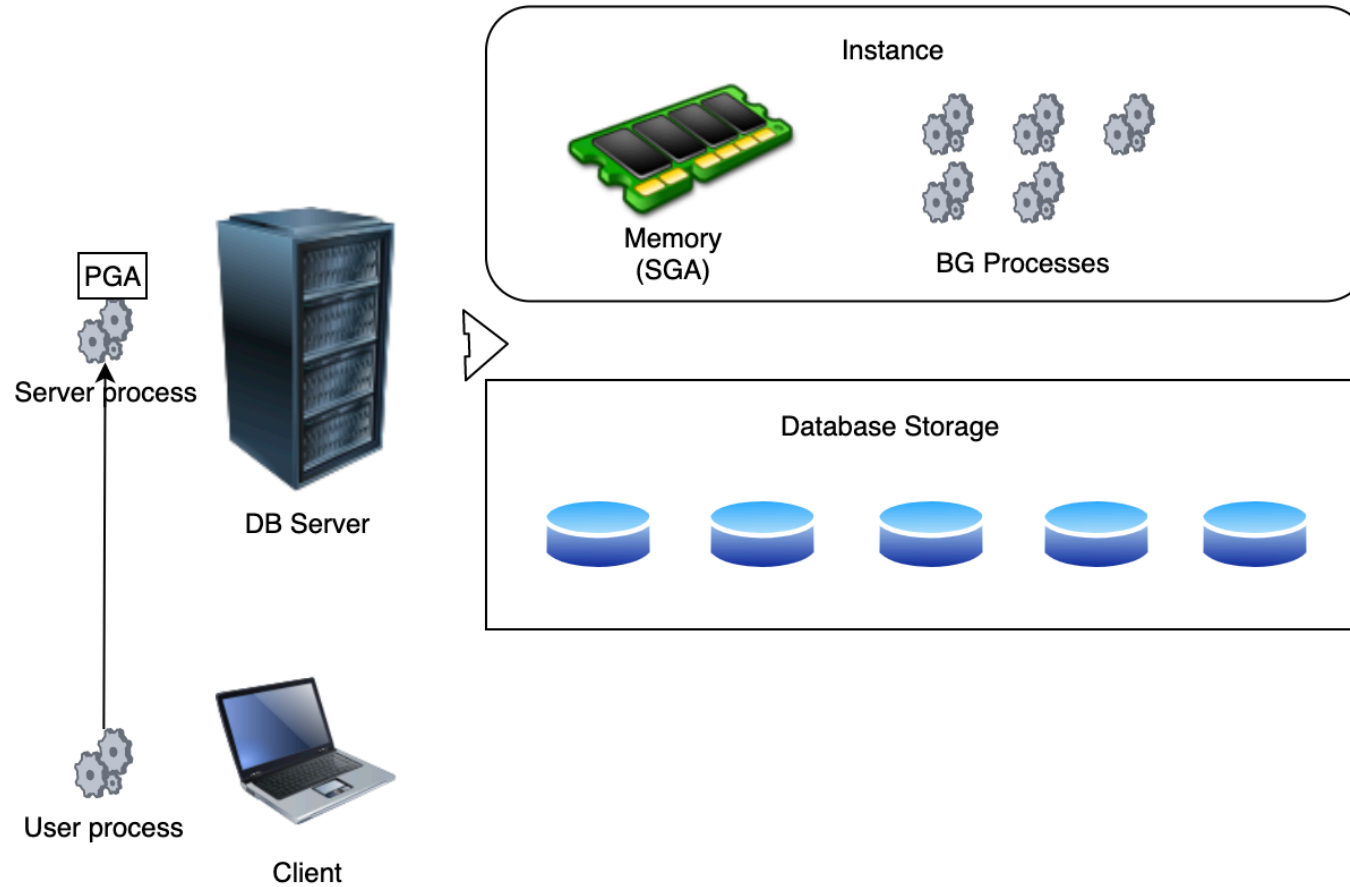
There are three major architectural structures in Oracle Database:

- Memory
- Processes
- Storage

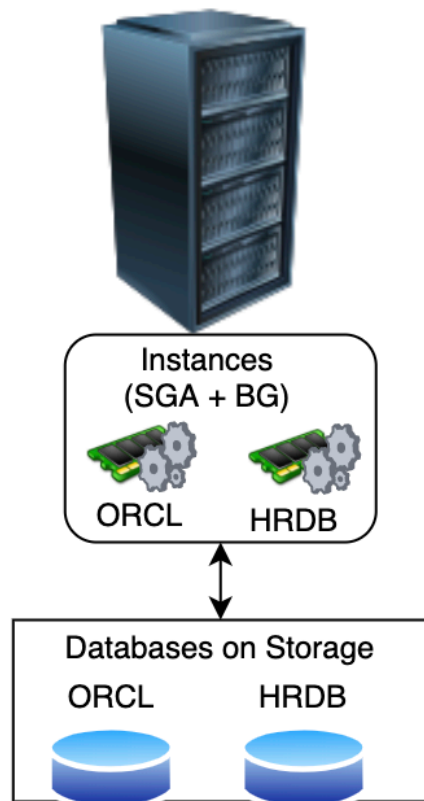
The instance consists of memory structures and background processes.

When instance starts SGA (System Global Area) is allocated and BG processes are started.

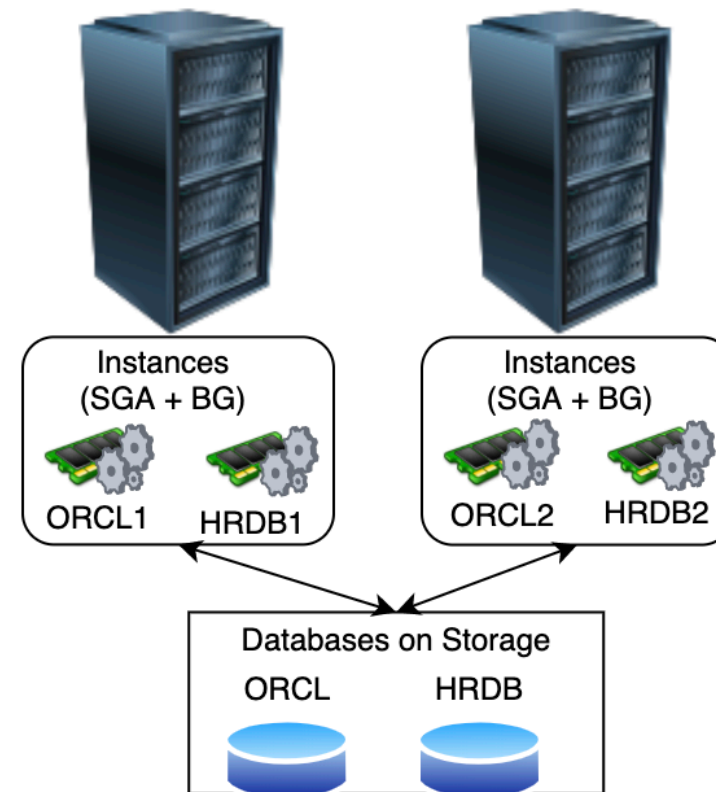
Database server architecture



Nonclustered system

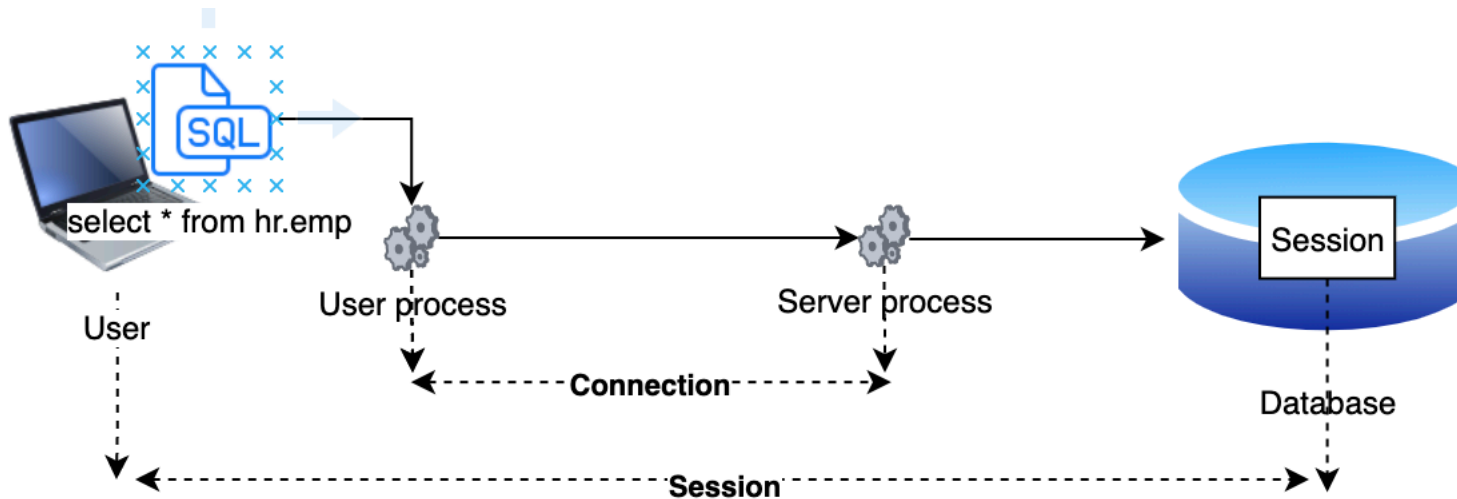


Clustered system

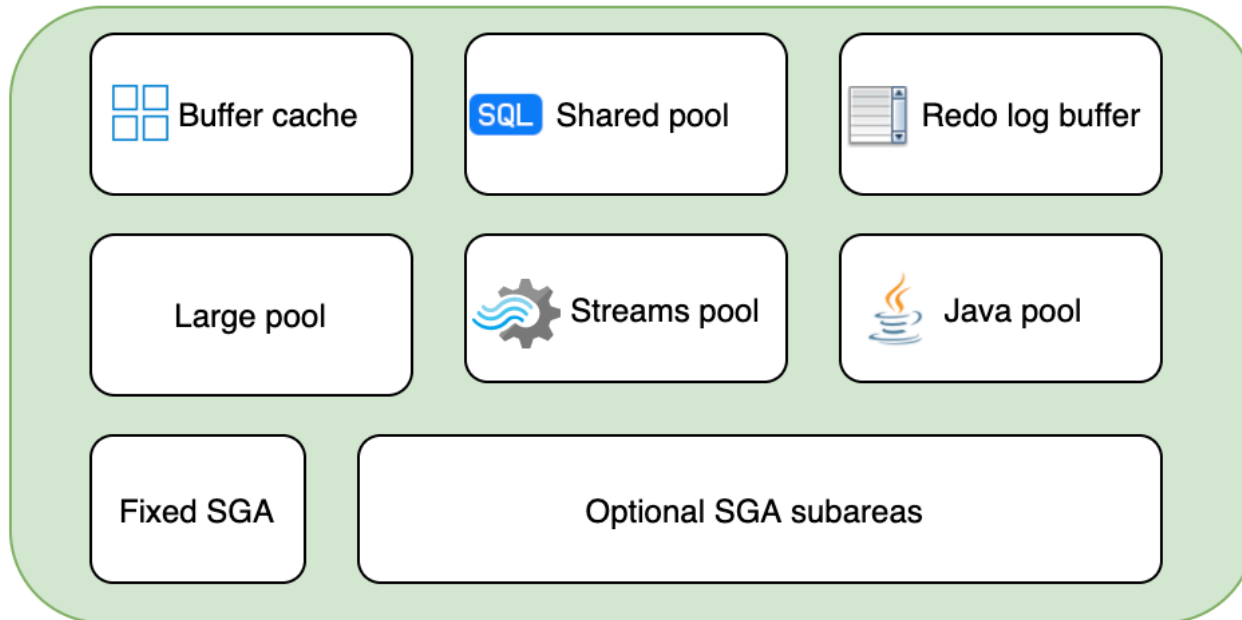


Connecting to the database


- Connection – communication link between user process and server process.
- Session – Specific connection between user and instance



Memory structures



System Global Area (SGA)



Instance Program Global Area (PGA)

For ASMM set:

`MEMORY_TARGET=0`

`MEMORY_MAX_TARGET=0`

`SGA_TARGET`

`SGA_MAX_SIZE`

`PGA_AGGREGATE_TARGET`

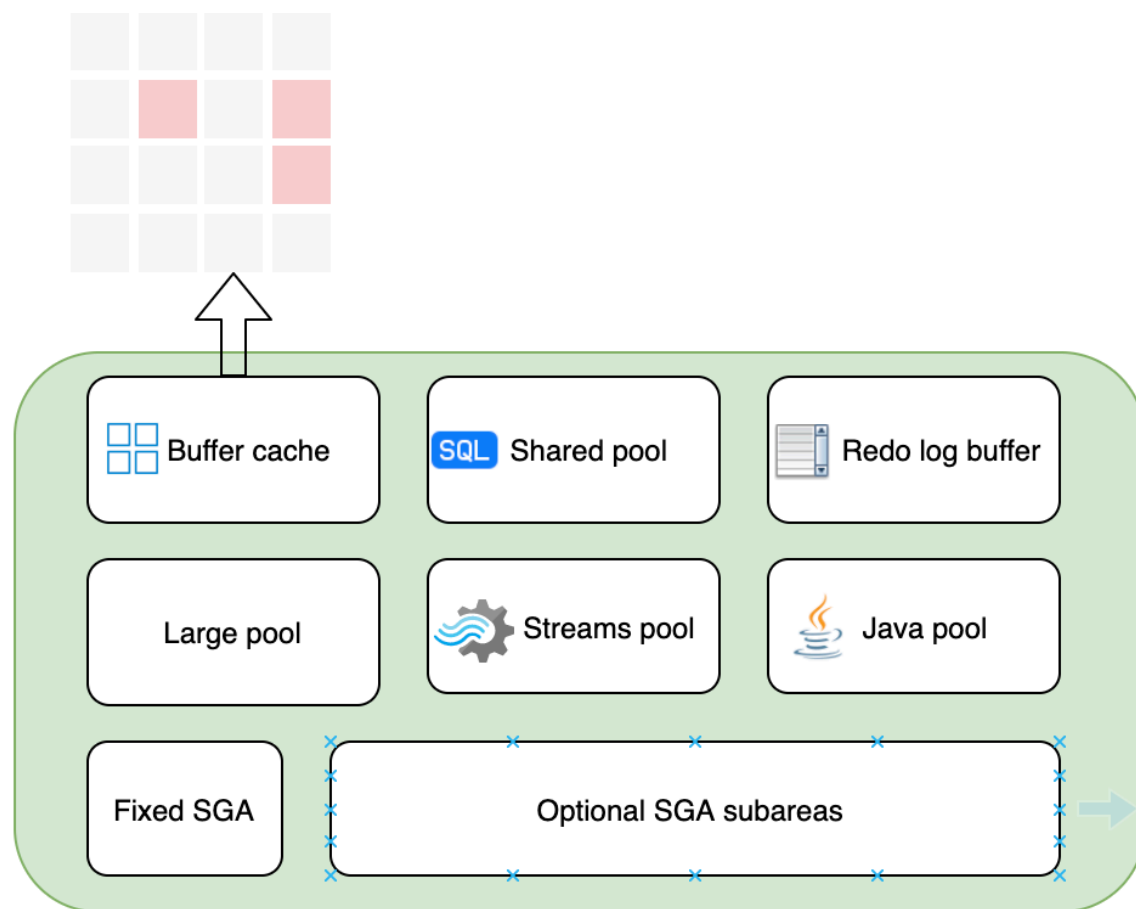
`PGA_AGGREGATE_LIMIT`

For AMM set:

`MEMORY_TARGET`

`MEMORY_MAX_TARGET`

Database buffer cache

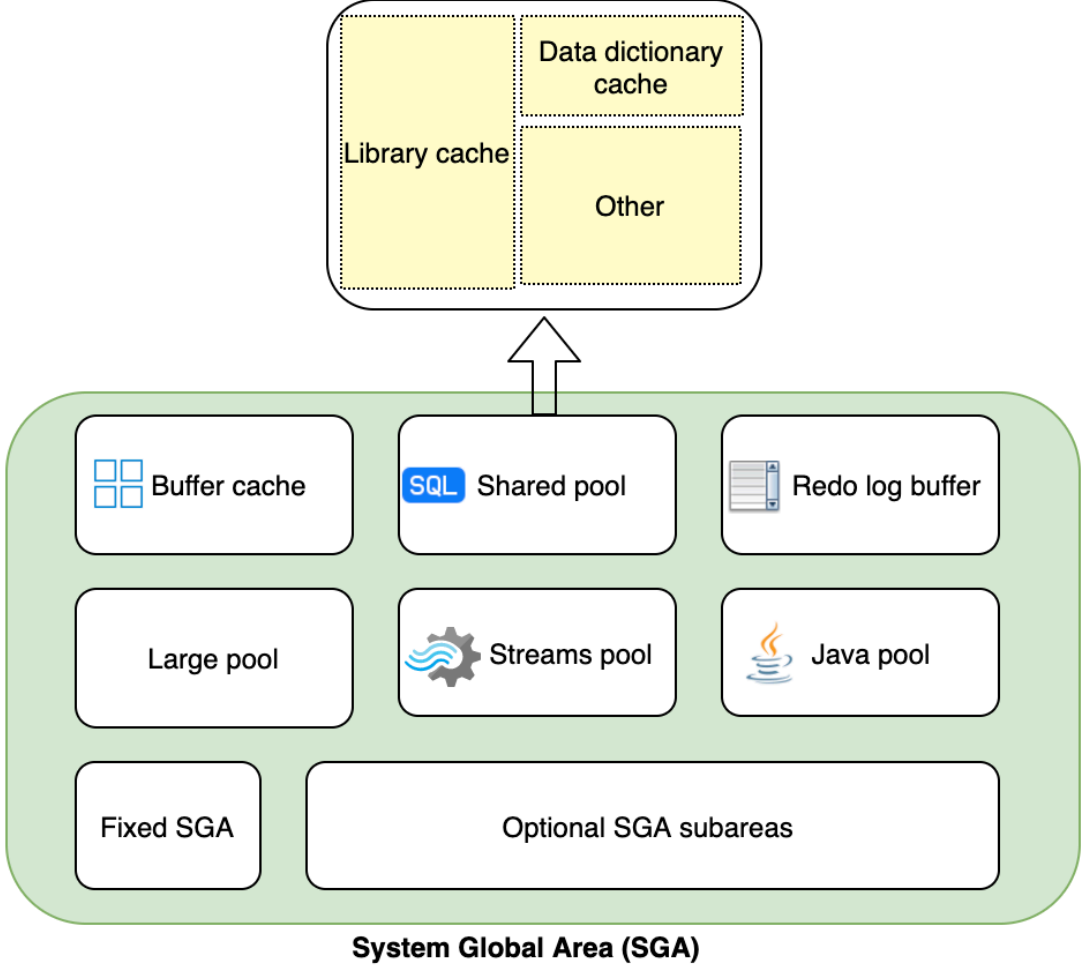


- Holds copies of data that are read from disk (data files)

■ Changed block
■ Unchanged block

System Global Area (SGA)

Shared pool

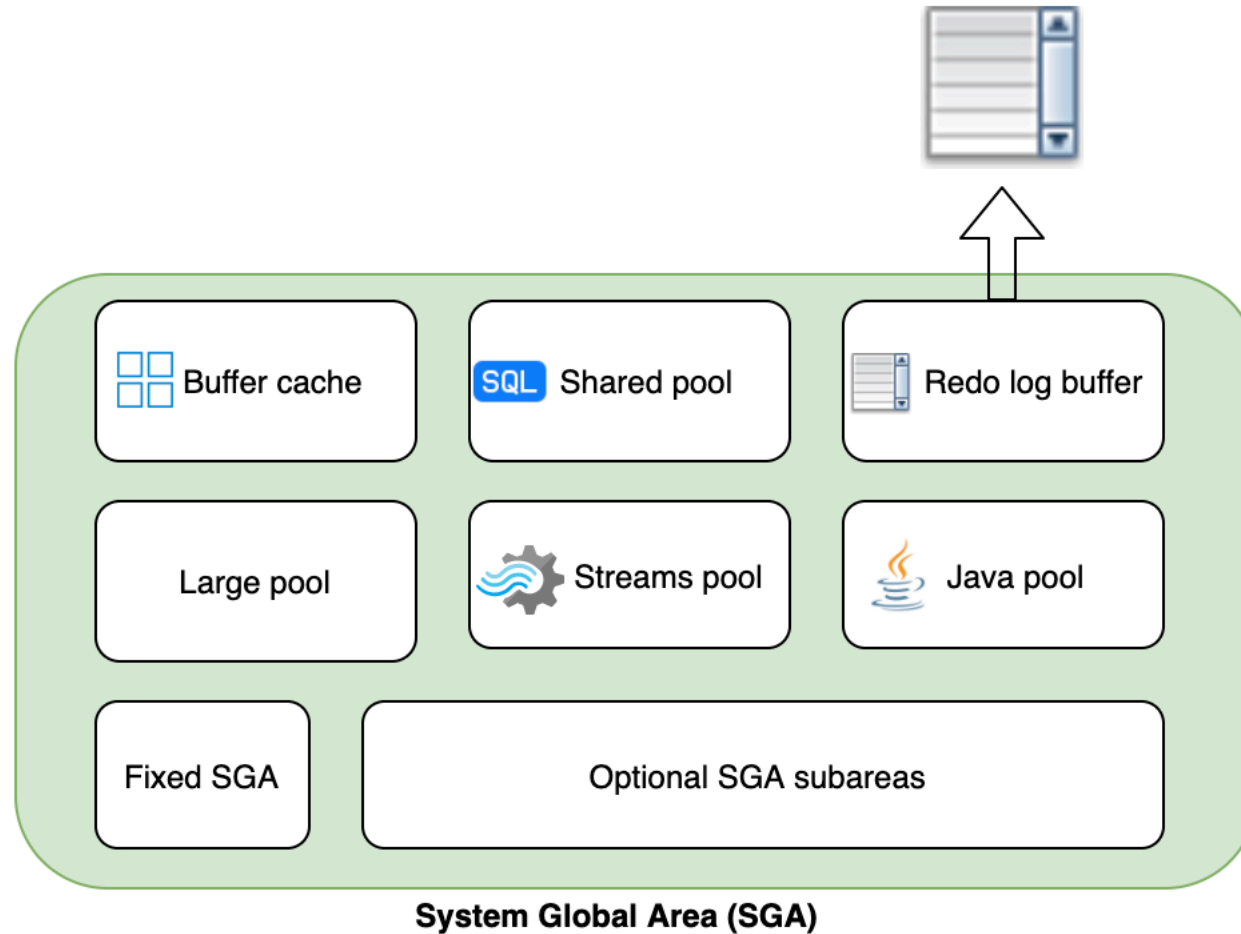


- **Data dictionary cache** – holds reference information about database, structures and users.
- **Library cache** - stores executable SQL and PL/SQL code:

```
e.g. select * from hr.emp
```

It also contains the statement parse tree and execution plan.

Redo log buffer

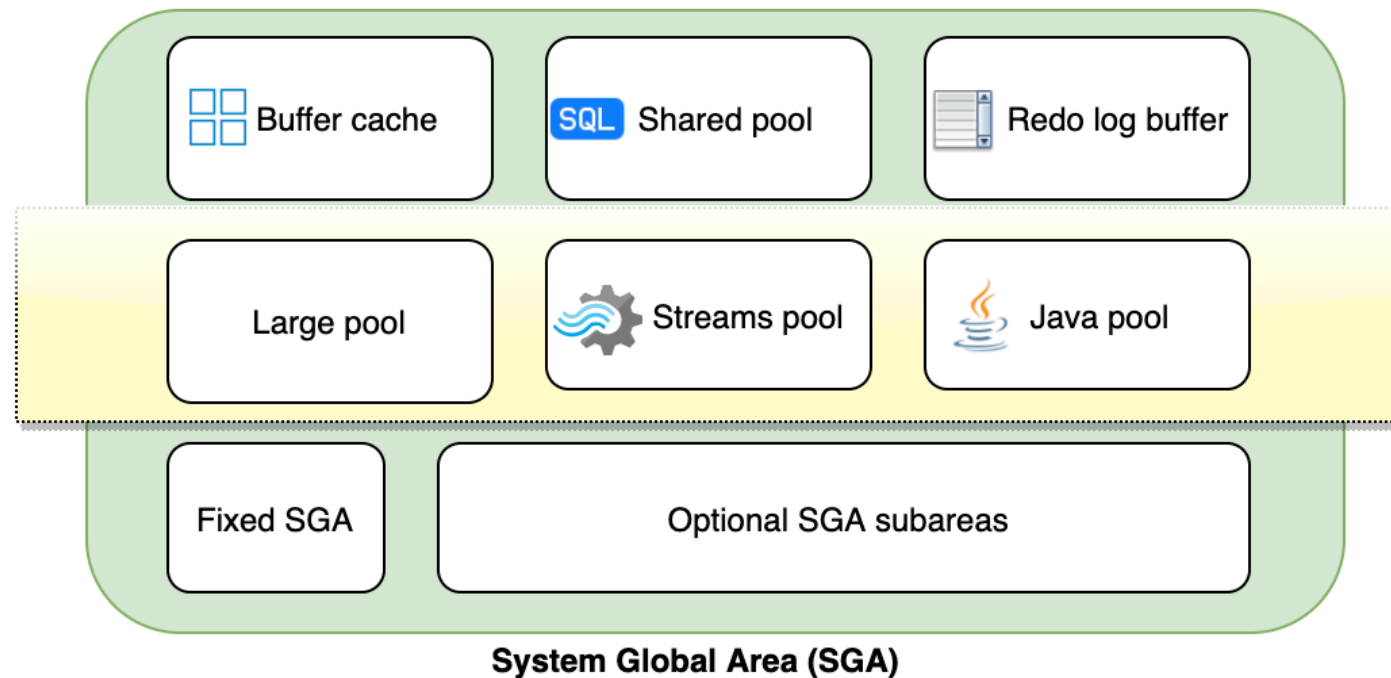


- Circular buffer
- Holds information about changes DML/DDDL and contains redo entries to redo changes made by these operations

Large pool - Used for allocation of big chunks of memory such as used by Parallel Query and RMAN

Java pool - stores all session-specific Java code and data within the Java Virtual Machine (JVM)

Streams pool – used by Oracle Streams to store information required capture and apply.

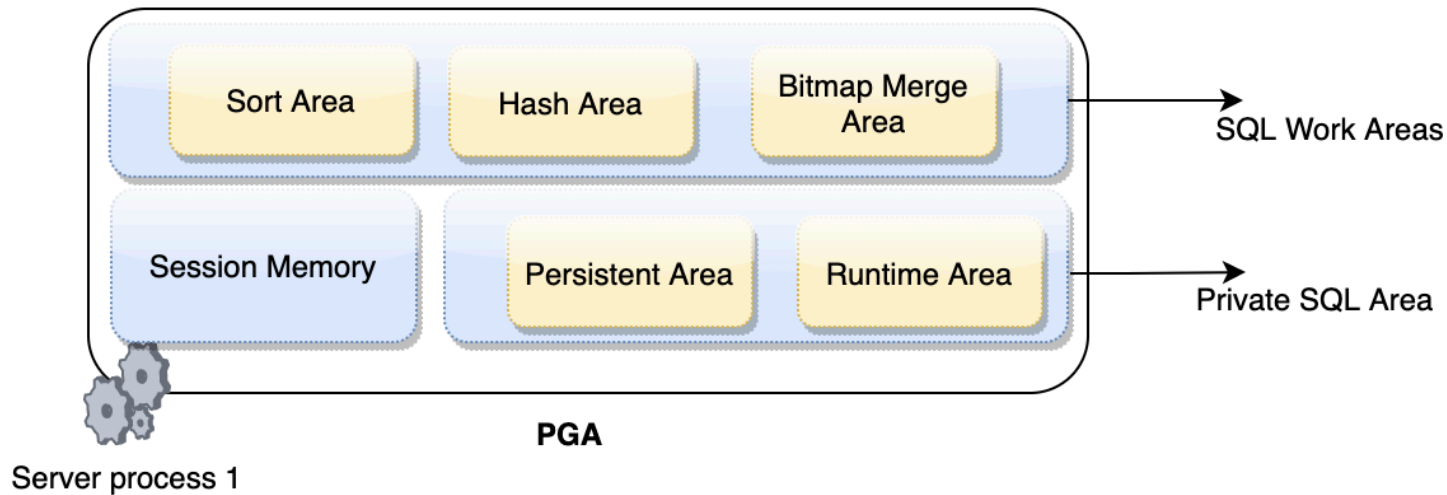


Fixed SGA

- Contains state of the database and the instance, background processes. It is bootstrap section of the SGA which is used internally by Oracle to find other pieces of SGA.

PGA

- Specific to an operating process or thread that is not shared by other processes



Sort area - used for sorting a set of rows

Hash area – used for hash join

Bitmap merge – used for bitmap merge operation, to merge data retrieved from multiple bitmap indexes.

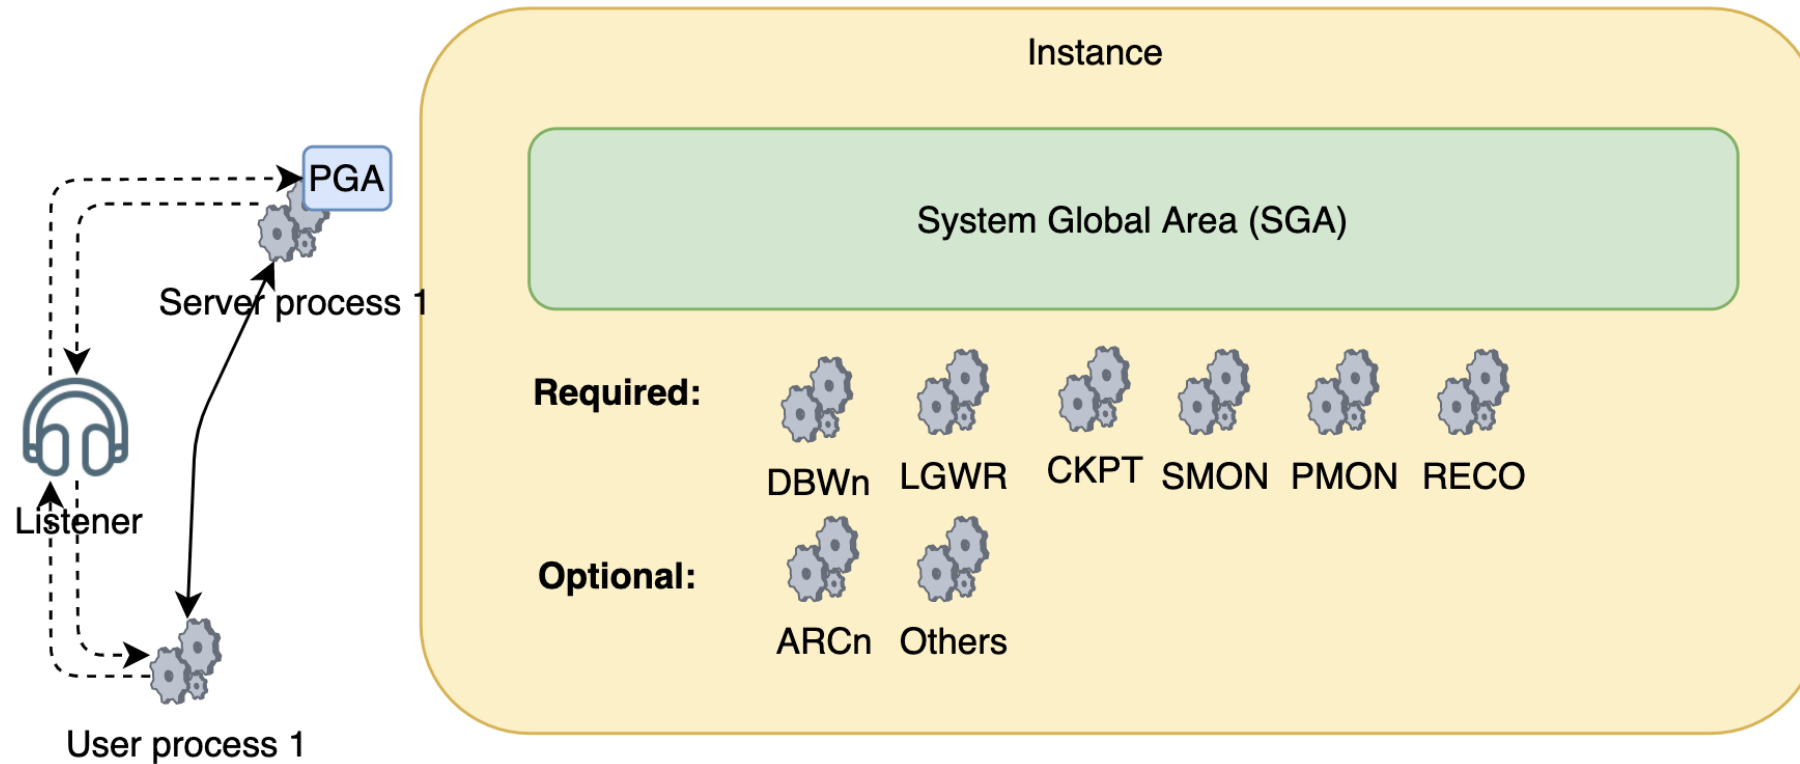
Persistent area - contains bind variable values.

Runtime Area – contains query execution state, tracks the number of rows retrieved so far.

Process architecture


- User process – process on the application/client side that connects to the database
- Database process
 - Server process – when user establishes a session, server process is started on the database side that connects to the instance
 - **Background processes** – are started when an Oracle instance is started
- Daemon/Application processes
 - Network listener
 - Grid infrastructure daemons

Process structure




Database Writer Process (DBWn)

- Writes dirty (changed) blocks from buffer cache to datafile (disk)




Log Writer Process (LGWR)

- Writes redo log buffer to redo log file on disk
 - When a user process commits a transaction
 - When log buffer is 1/3 full
 - Before DBWn writes modified blocks to disk
 - Every 3 sec



Checkpoint process (CKPT)

- Records checkpoint information into:
 - Controlfile
 - Datafiles headers



System Monitor Process (SMON)


- Performs recovery at instance startup
- Cleans unused temp segments

Process Monitor Process (PMON)

- Performs process recovery when a user process fails
 - Cleans db buffer cache
 - Frees resources used by failed user process, releases locks
- Monitors session idle timeout
- Dynamically registers database service with listeners (in 19c this is done by LREG process)


Recovery Processes (RECO)

- Used with distributed db configuration
- Automatically connects database and resolves/recovers all in-doubt/failed transactions



Archiver Process (ARCn)

- Copy redo logs to a designated storage after log switch
- Can collect redo data and transmit to standby destination
- Exists only when the database is in ARCHIVELOG mode



Database physical structure



Control file



Data file



Online redo log file



Parameter file




Backup file



Archive log

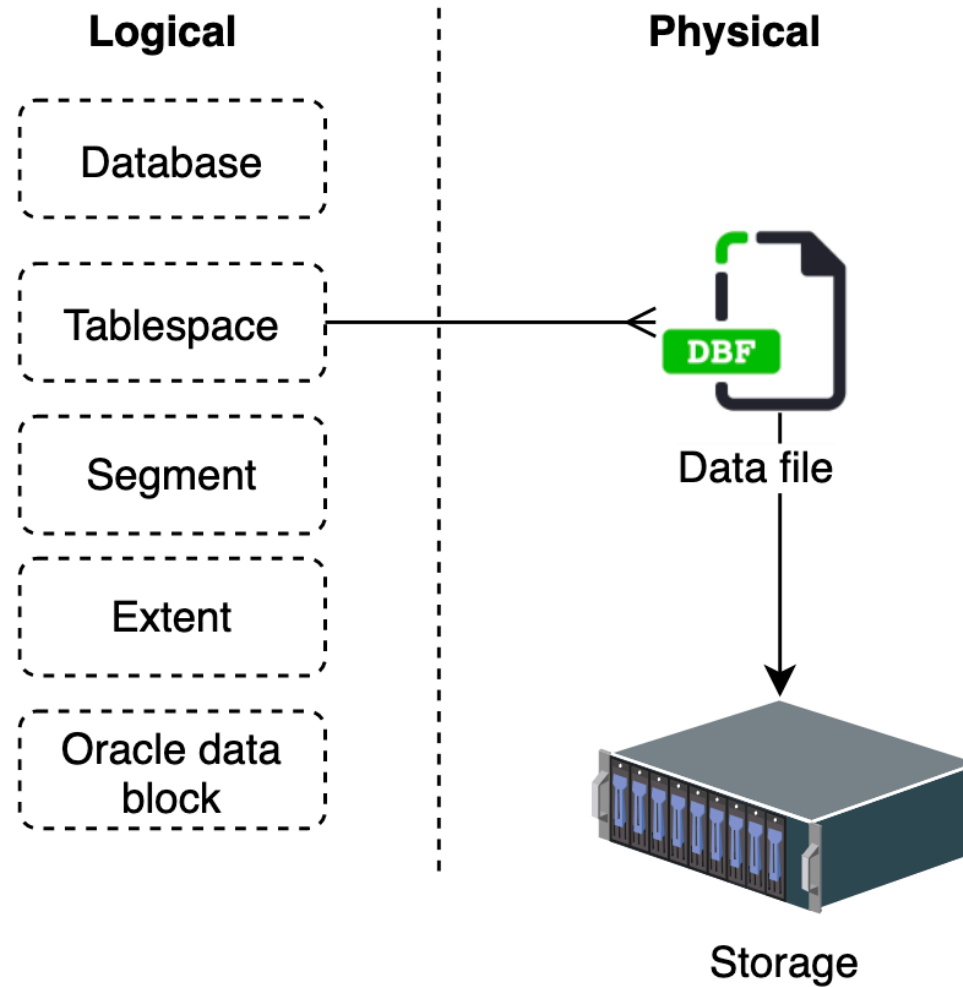


Password file




Alert log/trace files

Logical and Physical structures



Segment, Extent, Block

- Segment exist in a tablespace
- Segments are collection of extents
- Extent are collection of data blocks
- Data blocks are mapped to disk blocks



Segment



Extent

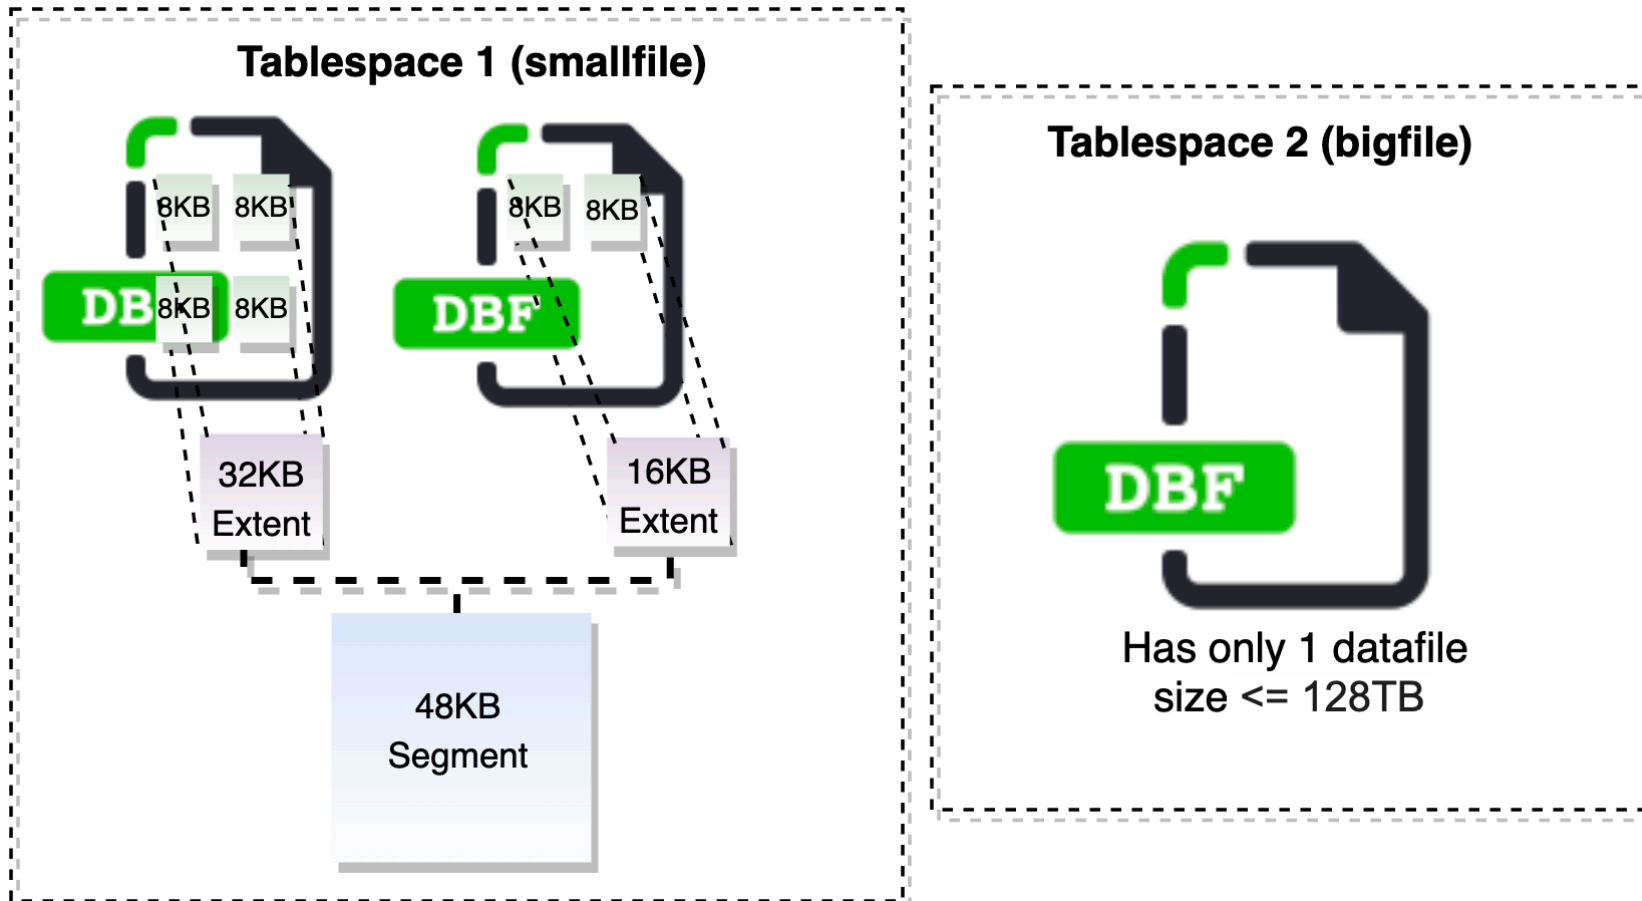


DB block




File system block

Tablespace and Data files



Important tablespaces

- **SYSTEM** – is always online when DB is open, contains the data dictionary tables for the entire database.
- **SYSAUX** – auxiliary tablespace to SYSTEM tablespace, centralized location for database metadata that does not reside in the SYSTEM. Contains additional data such as EM repository.
- **TEMP** – contains temporary data, may contain data for sorting operation which do not fit in memory
- **UNDO** - contains undo segments



 mariam.kupa@gmail.com

 mariam_kupa

1) How to check parameters in Oracle database?

```
SQL> show parameter sga_max_size
SQL> select value/1024/1024 from v$parameter where name= 'sga_max_size';
```

2) How to set size for the database buffer cache?

```
SQL> alter system set db_cache_size=300M scope=spfile;
```

The value will take affect after database restart.

3) How to check where my datafiles and redo log files are located?

```
SQL> select file_name from dba_data_files;
SQL> select member from v$logfile;
```

4) How to check where my controlfiles are located?

```
SQL> select name from v$controlfile;
SQL> show parameter control_files
```

5) How to check what is the maximum PGA usage after instance startup?

```
SQL> select VALUE/1024/1024 from v$pgastat where name='maximum PGA allocated';
```

6) How to find what is OS process ID of my DBW? Or LGWR?

```
$ ps -ef|grep dbw
$ ps -ef|grep lgwr
```