Introduction to Database storage

Kiryl Bucha

Overview

- Oracle Database 11g - Architecture
- Connection to Oracle
- Oracle Database Data Files - Storage Example
- The Memory Structures Oracle Database
- I/O Process Oracle Database
- Select, Update, Insert Operations

Oracle Database 11g Architecture

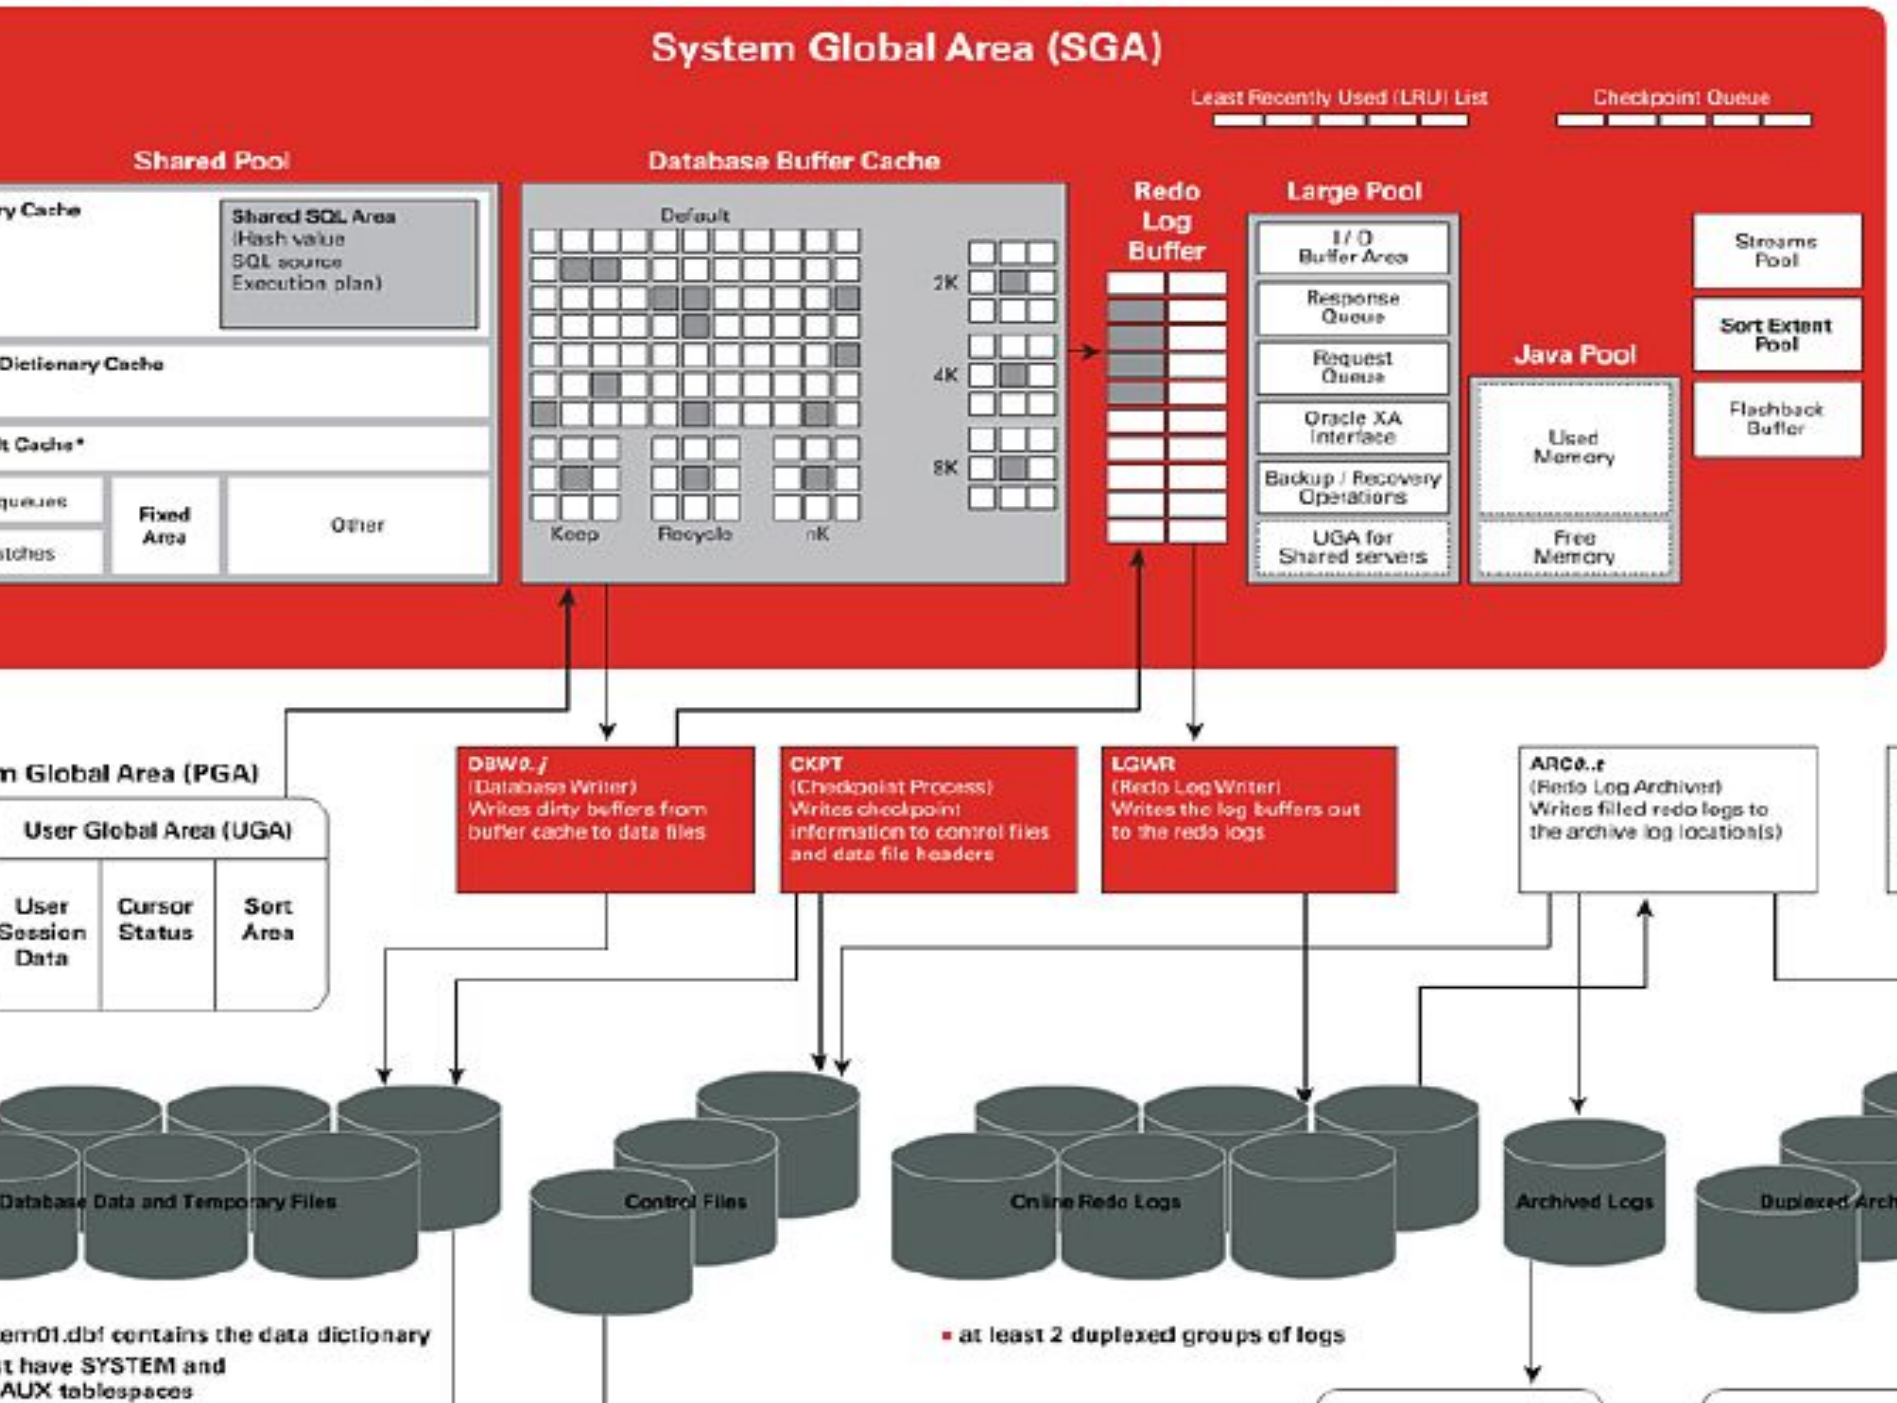
System Global Area (SGA)
Least Recently Used (LRU) List
Checkpoint Queue
Shared Pool
Shared SQL Area
(Hash value
SQL source
Execution plan)
Dictionary Cache
Fixed Area
Other
Database Buffer Cache
Default
2K
4K
8K
Keep
Recycle
nK
Redo Log Buffer
Large Pool
I/O Buffer Area
Response Queue
Request Queue
Oracle XA Interface
Backup / Recovery Operations
UGA for Shared servers
Java Pool
Used Memory
Free Memory
Streams Pool
Sort Extent Pool
Flashback Buffer
n Global Area (PGA)
User Global Area (UGA)
User Session Data
Cursor Status
Sort Area
DBW0..j
(Database Writer)
Writes dirty buffers from buffer cache to data files
CKPT
(Checkpoint Process)
Writes checkpoint information to control files and data file headers
LGWR
(Redo Log Writer)
Writes the log buffers out to the redo logs
ARC0..t
(Redo Log Archiver)
Writes filled redo logs to the archive log location(s)
Database Data and Temporary Files
Control Files
Online Redo Logs
Archived Logs
em01.dbf contains the data dictionary
t have SYSTEM and
AUX tablespaces
at least 2 duplexed groups of logs

Connection to Oracle

Connection to Oracle – Dedicated Server

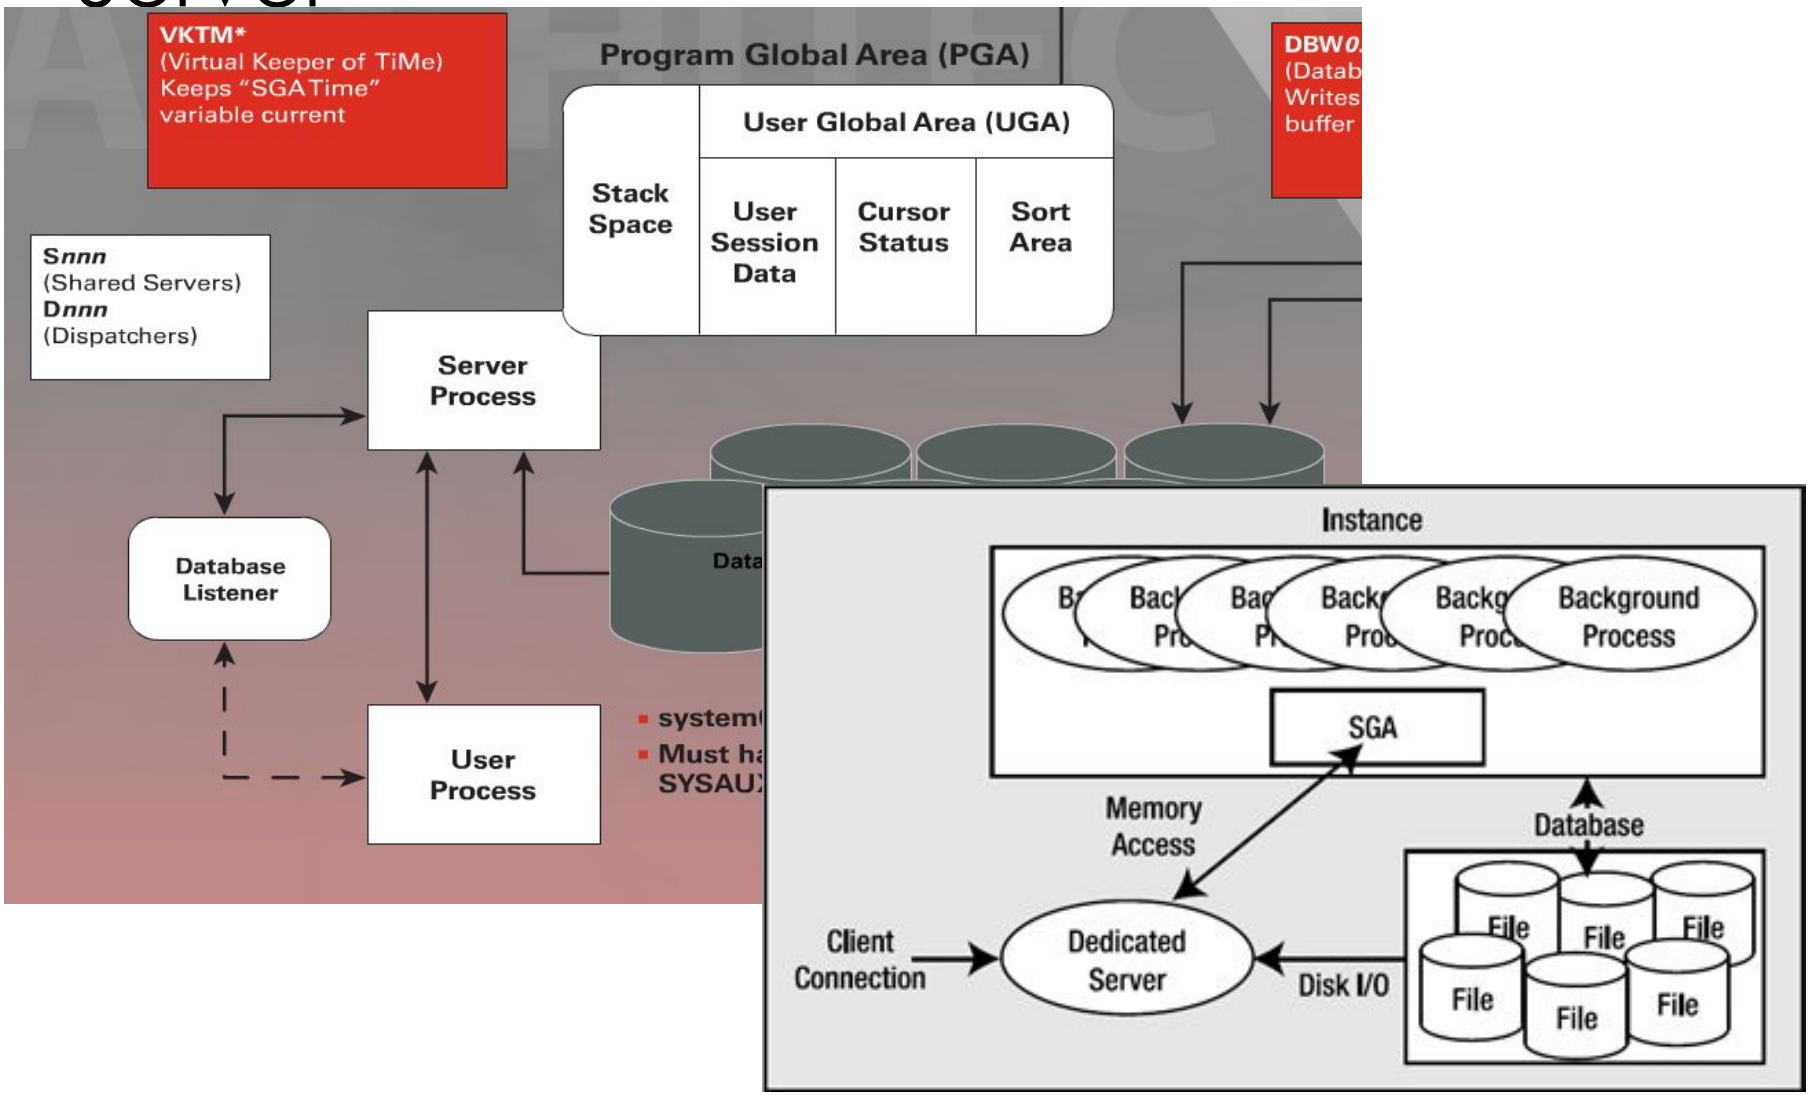

Connection to Oracle – Shared Server

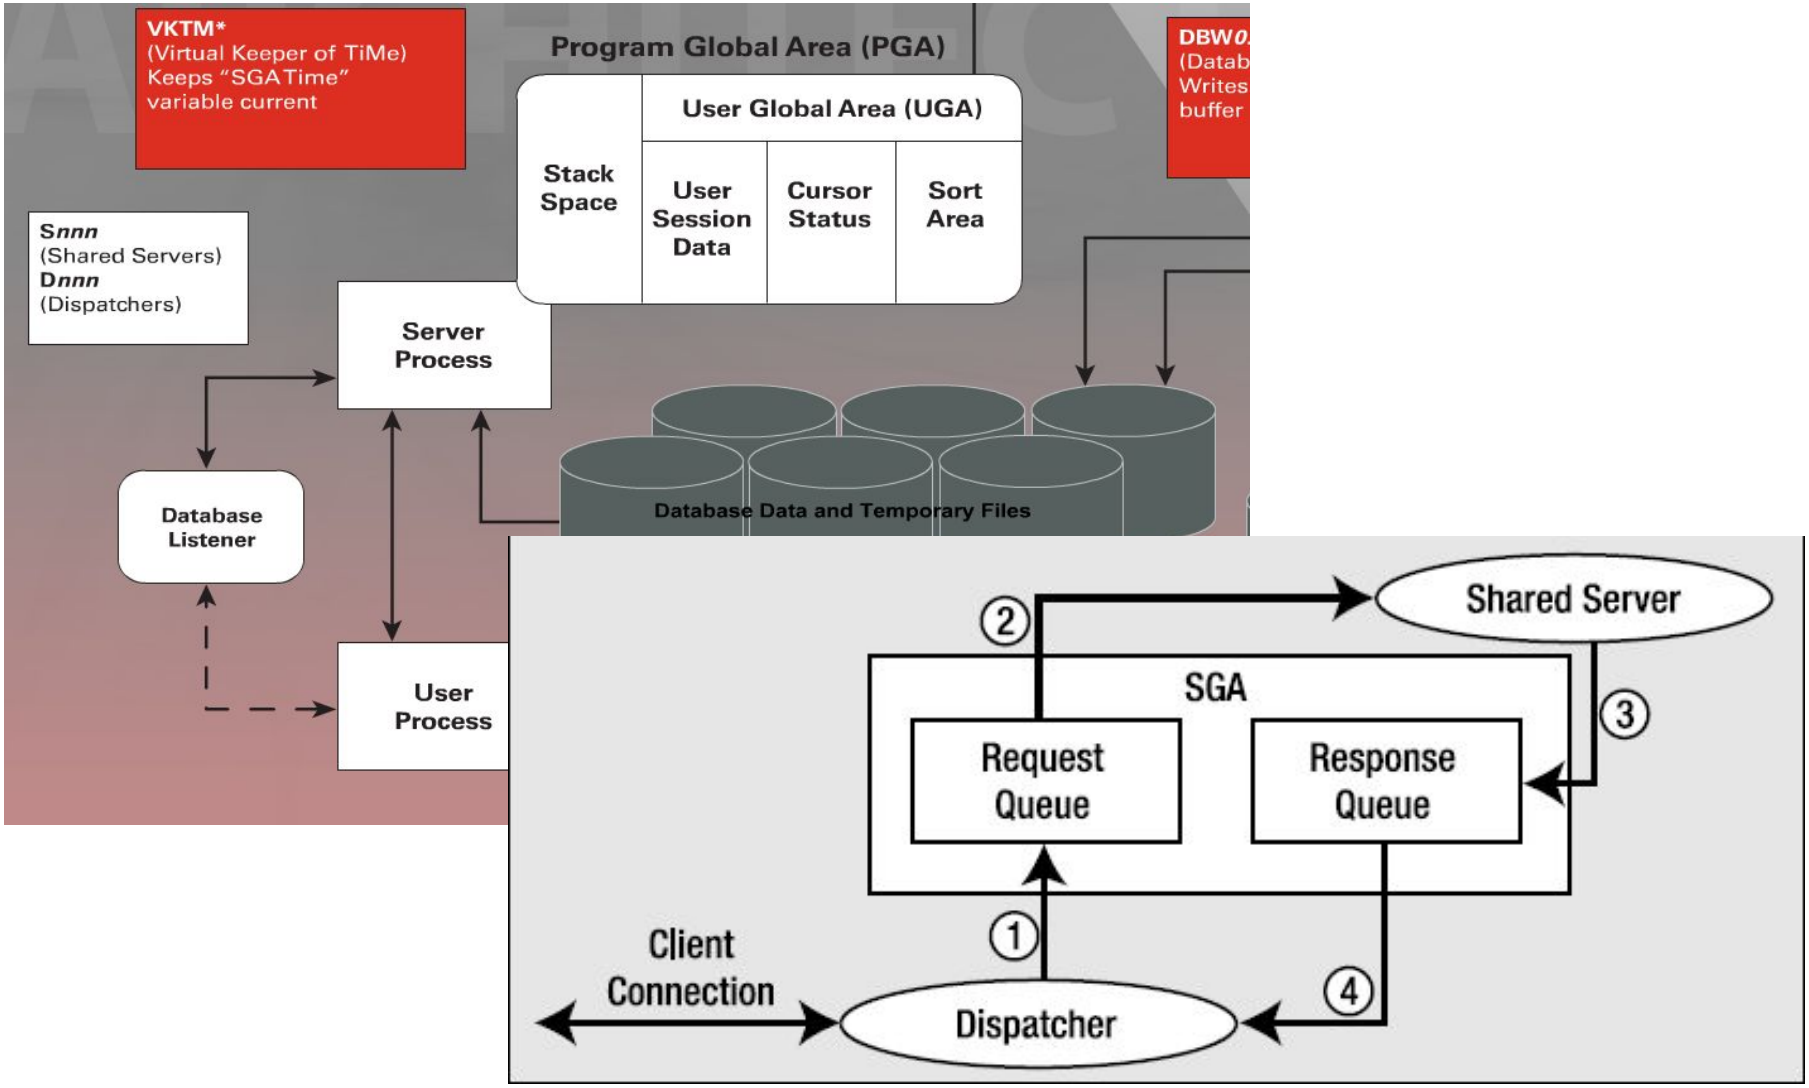

Oracle Database Data Files - Storage

A Brief Review of File System Mechanisms

- **"Cooked" operating system (OS) file systems**: You can use simple OS utilities such as xcopy on Windows or cp on UNIX to move them around. Cooked OS files are historically the most popular method for storing data in Oracle, but I see that changing with the introduction of ASM (more on that in a moment).
- **Raw partitions:** These are not files—these are raw disks. You don't ls them; you don't review their contents in Windows Explorer. They are just big sections of disk without any sort of file system on them. The entire raw partition appears to Oracle as a single large file. This is in contrast to a cooked file system, where you might have many dozens or even hundreds of database data files
- **Automatic Storage Management (ASM):** This is a new feature of Oracle 10g Release 1. In releases prior to 11g Release 2, ASM is a file system designed exclusively for use by the database. ASM is designed to work in either a single machine or clustered environment. Since Oracle 11g Release 2, ASM provides not only this database file system but optionally a clustered file system as well, which is described next.
- **Clustered file system:** This is specifically for a RAC (clustered) environment and provides what looks like a cooked file system that is shared by many nodes (computers) in a clustered environment. A traditional cooked file system is usable by only one computer in a clustered environment.

The Storage Hierarchy in an Oracle Database

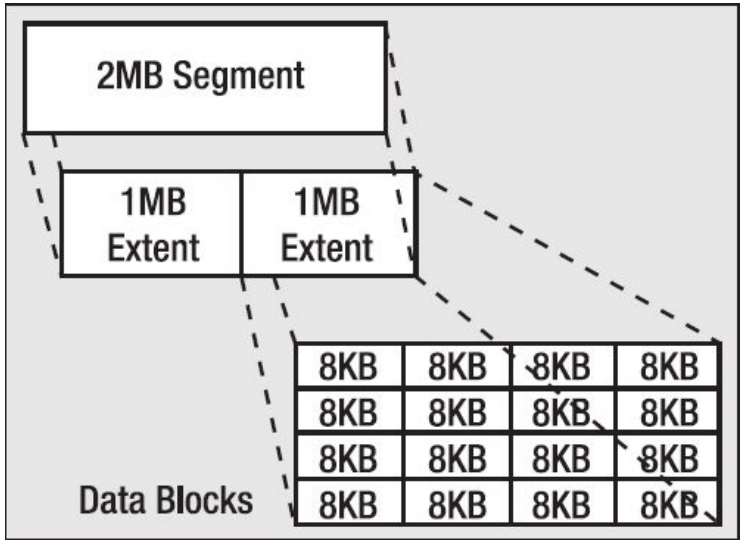

- ***Segments*** are the major organizational structure within a tablespace. Segments are simply your database objects that consume storage—typically objects such as tables, indexes, undo segments, and so on. Most times, when you create a table, you create a table segment.
- **Extents:** Segments consist of one or more *extent*. An extent is a logically contiguous allocation of space in a file. Traditionally, every segment starts with at least one extent.
- **Blocks:** Extents, in turn, consist of *blocks*. A block is the smallest unit of space allocation in Oracle. Blocks are where your rows of data, or index entries, or temporary sort results are stored.

Data Block – PCTFREE , PCTUSED

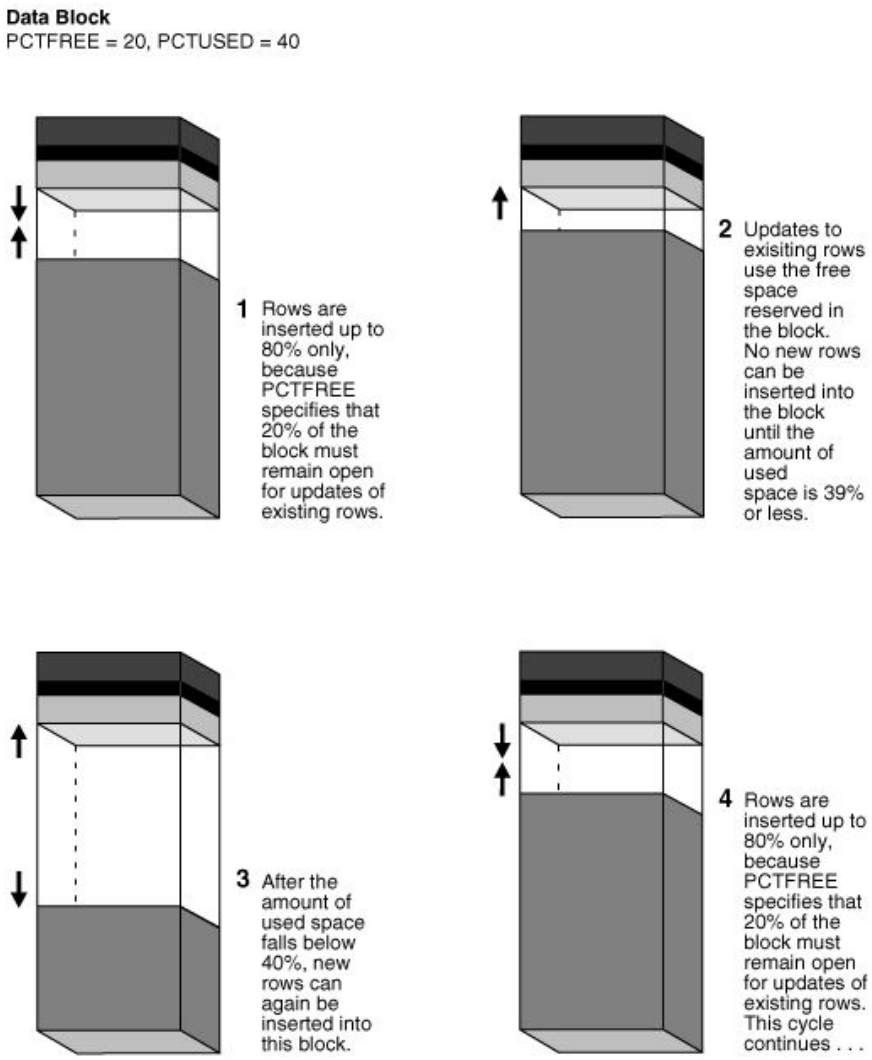

- In many cases, the OLTP database might be using a small block size, such as 2KB or 4KB, whereas the DW would be using a much larger one (8KB or 16KB).
- Often for OLAP tables very important follow best practices with technical attributes PCTFREE and PCTUSED.

The Memory Structures Oracle Database

The Memory Structures Oracle Database

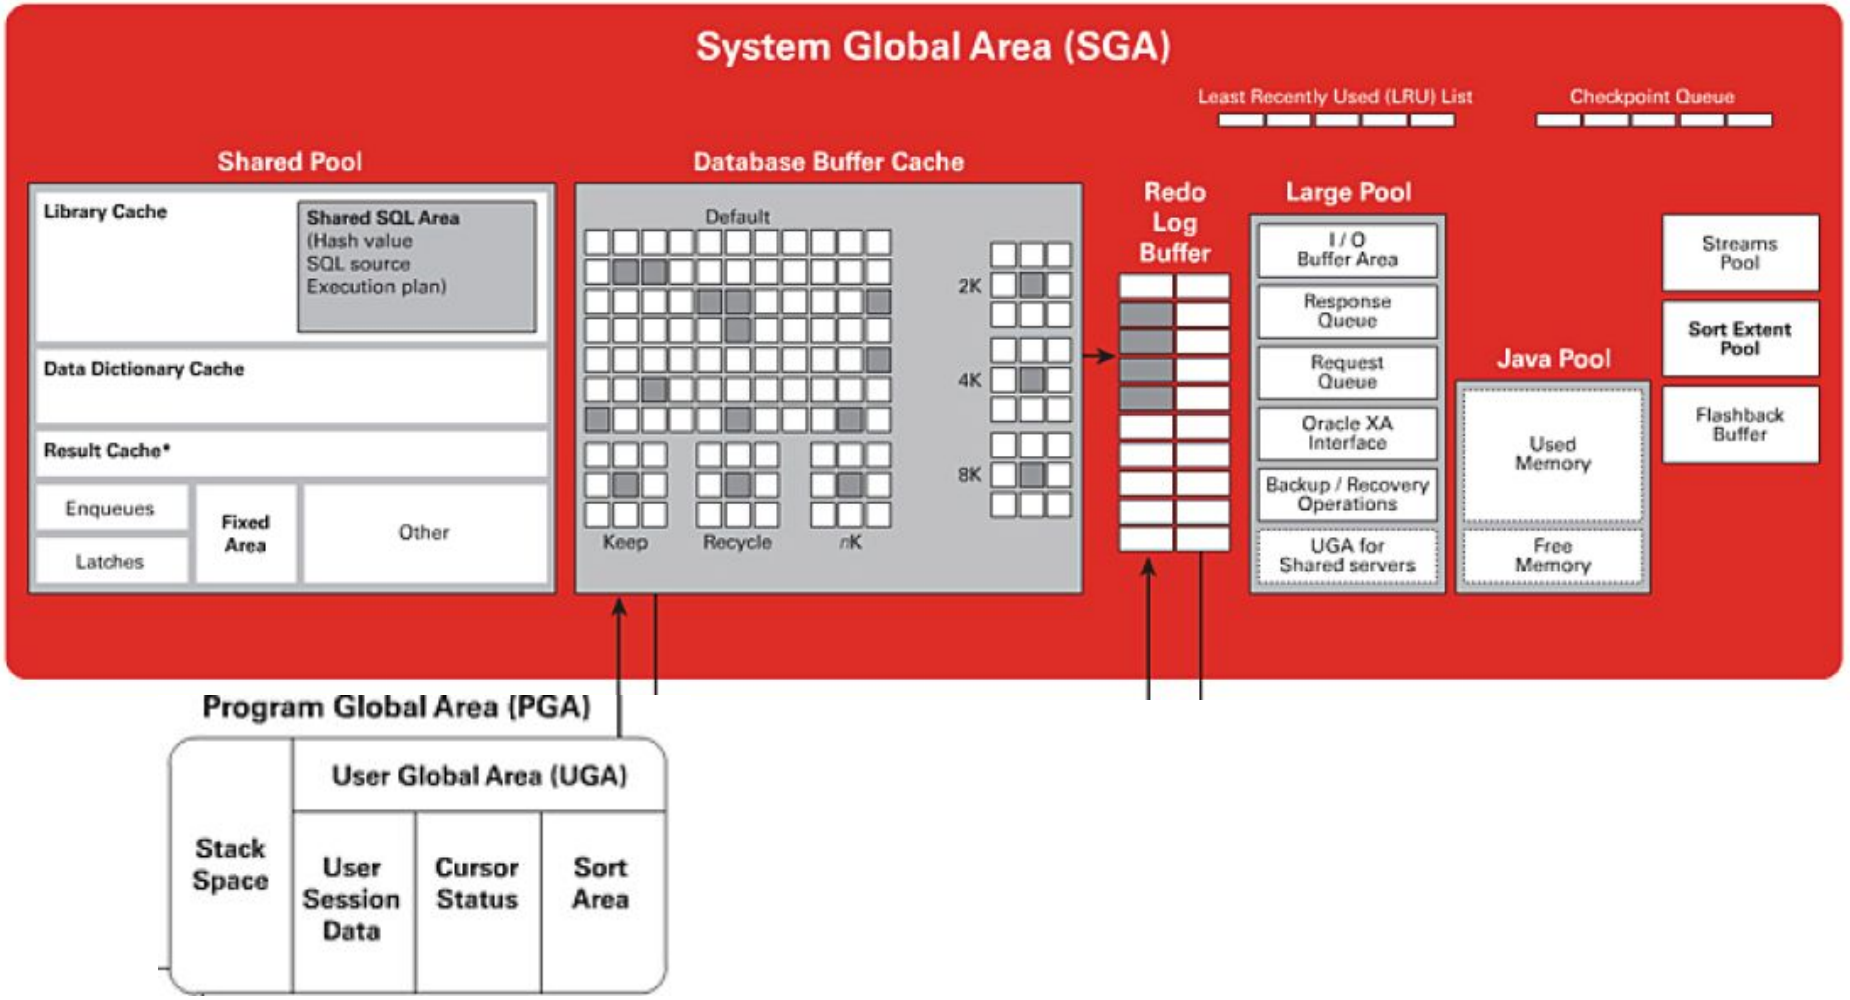

Select, Update, Insert Operations

Select, Update, Insert Operations

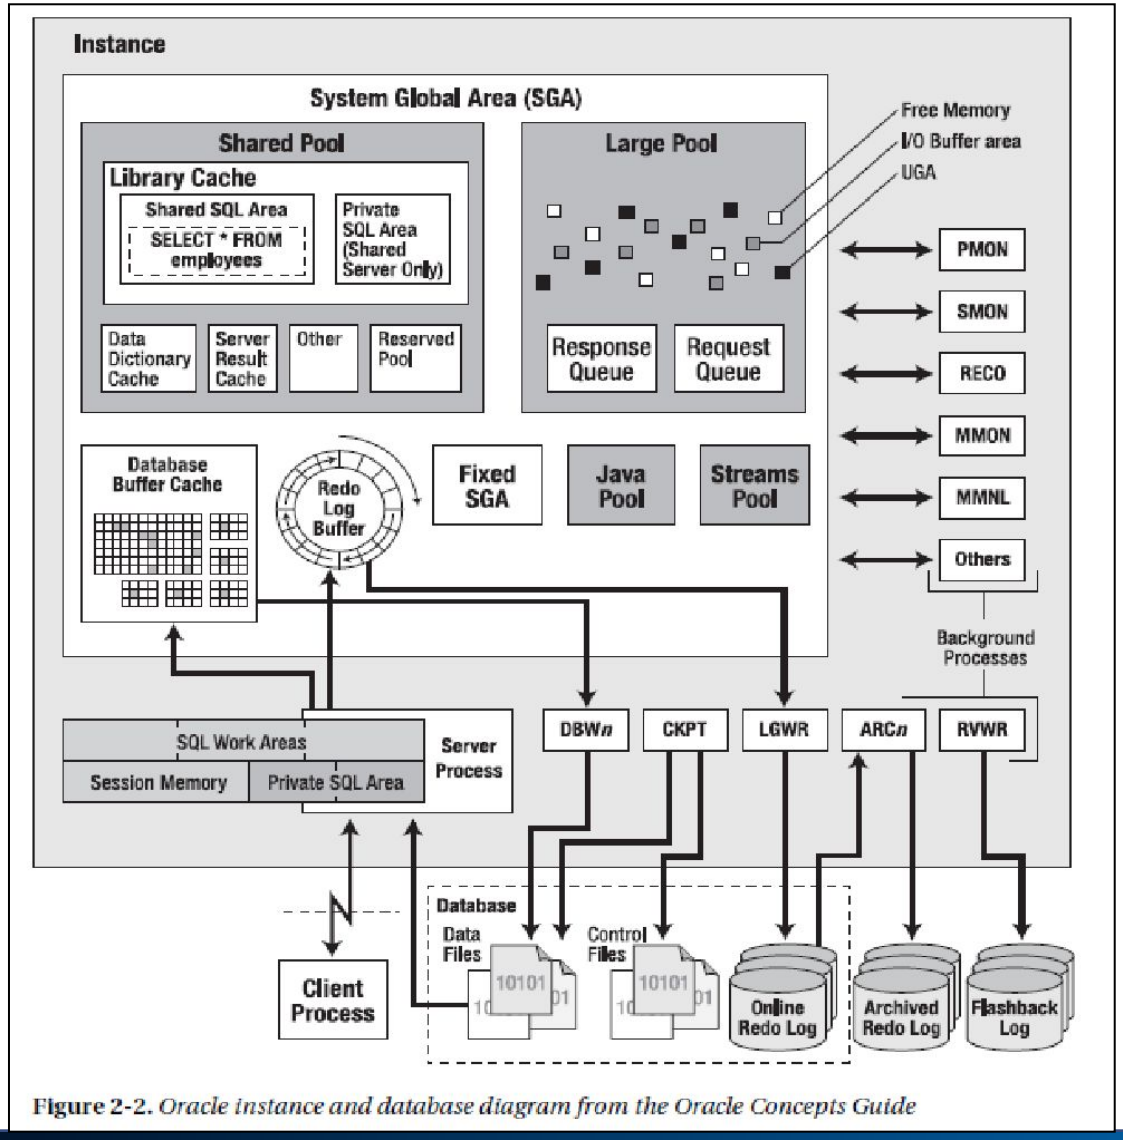

Figure 2-2. *Oracle instance and database diagram from the Oracle Concepts Guide*

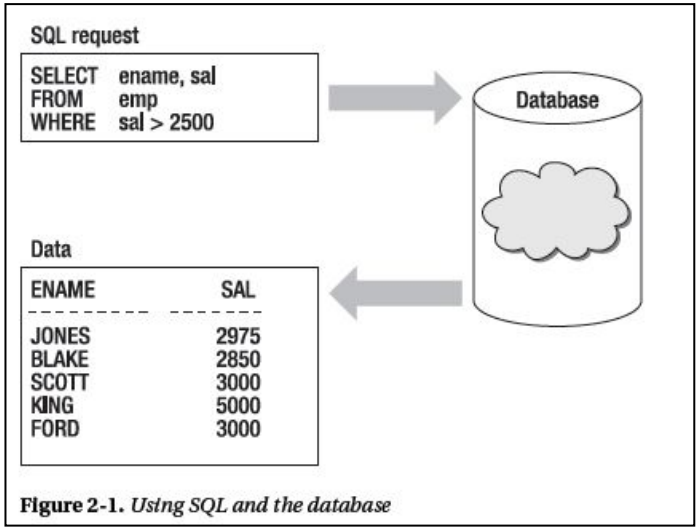

Figure 2-1. *Using SQL and the database*

Thank you
Time for questions

Kiryl Bucha