



DB

Database Services

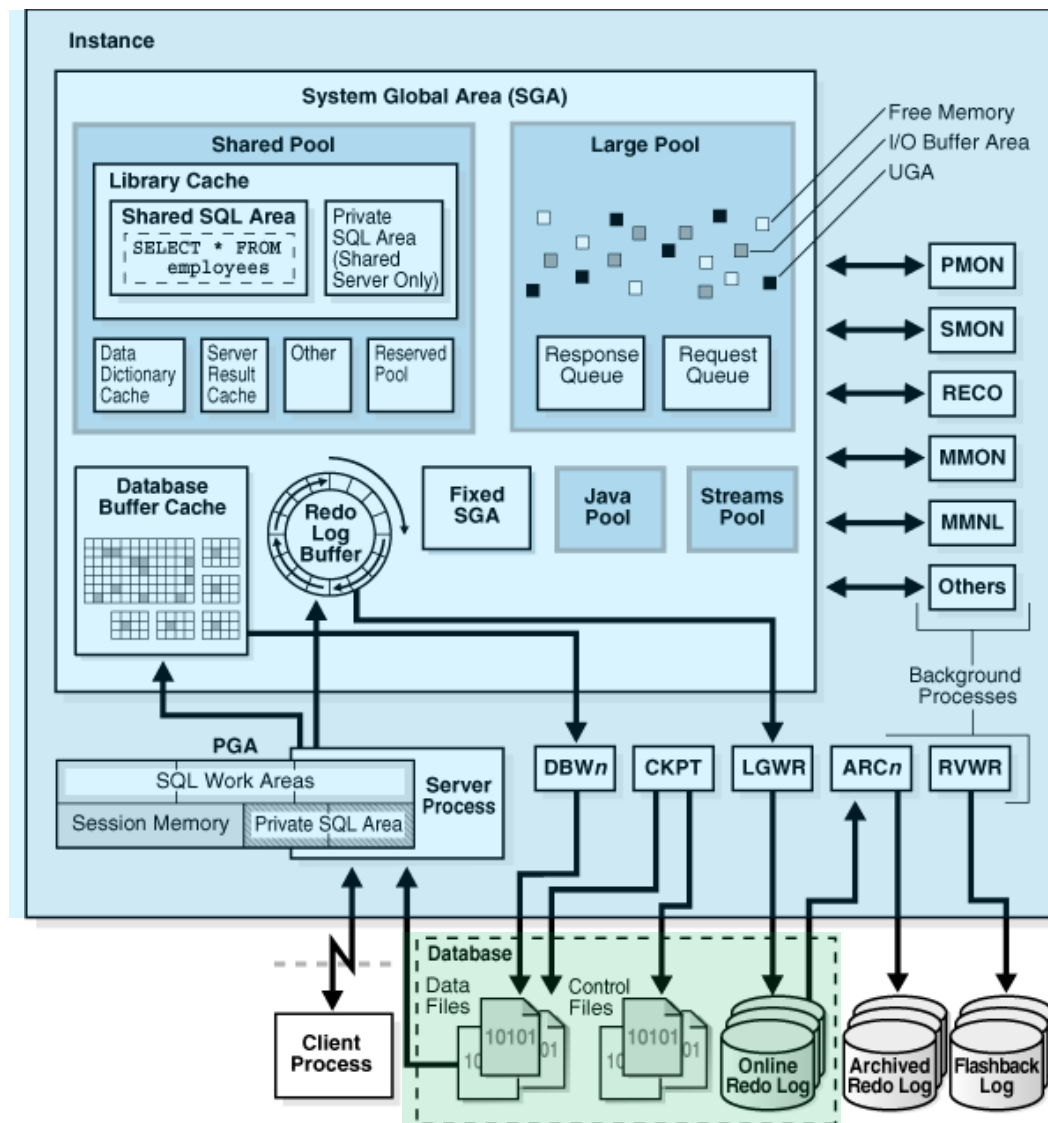
CERN IT
Department

Oracle Architecture

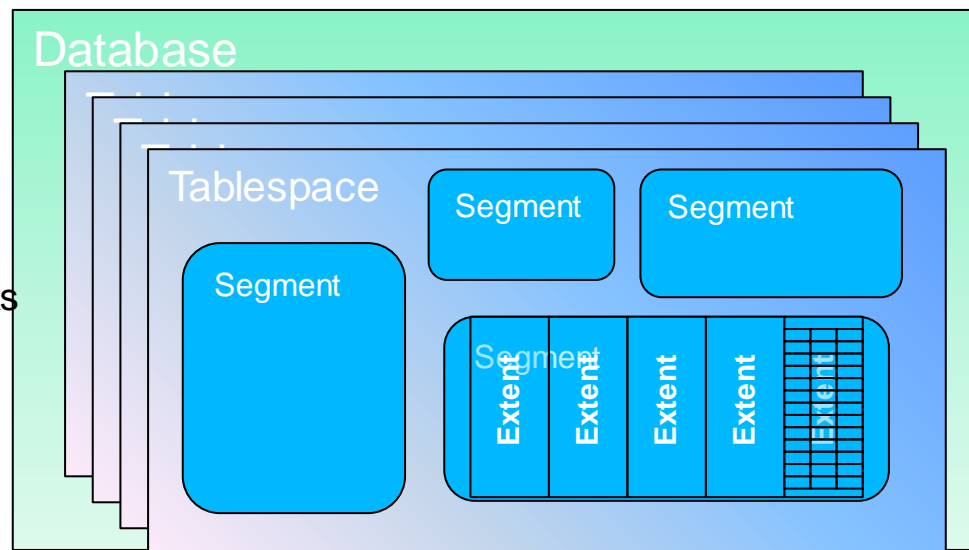
Przemyslaw Radowiecki
IT-DB-DBF

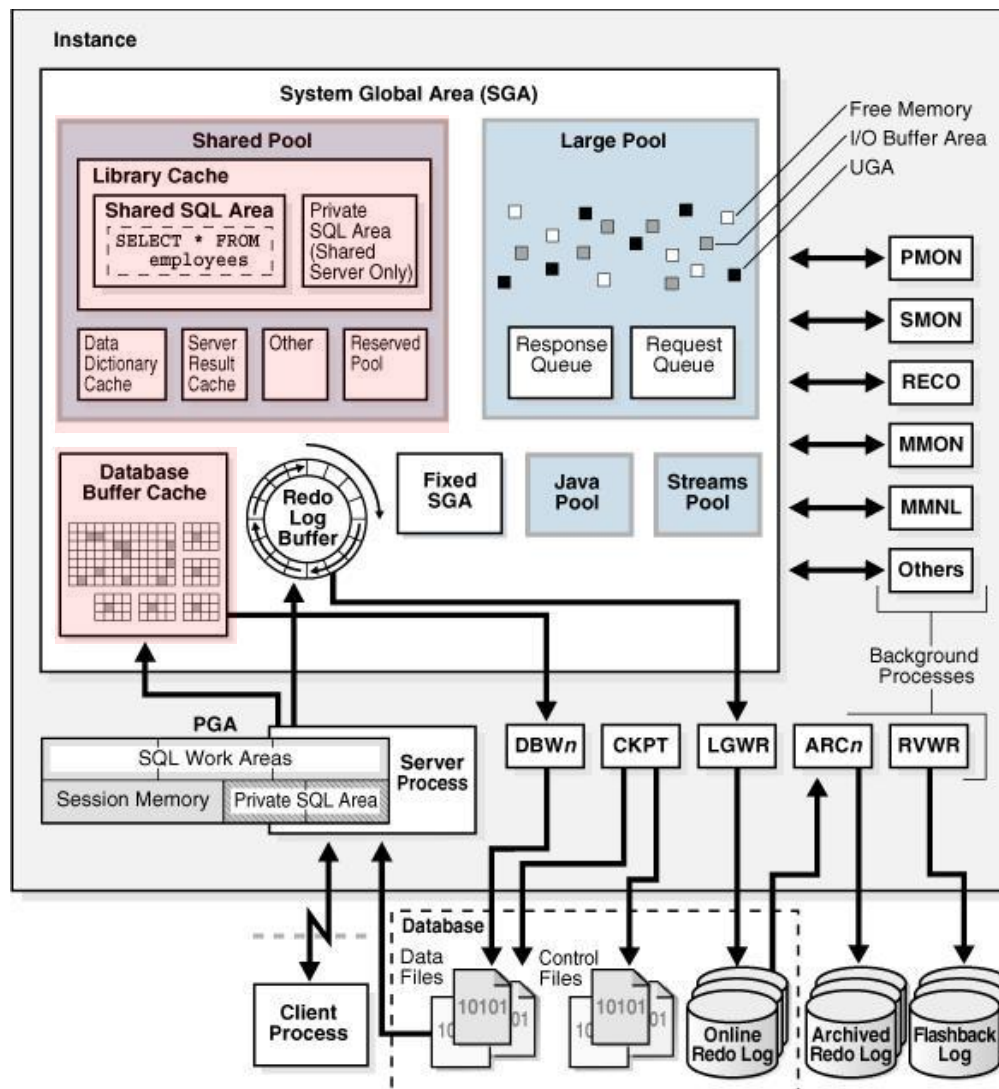
- Oracle Database
- Oracle Instance
- Real Application Clusters (RAC)
- Network connectivity
- SQL statements processing

- Organized collection of data treated as a unit. The purpose of a database is to store and retrieve related information.
- Oracle's implementation of a database consists of
 - data **files**
 - control **files**
 - online redo log **files**
- A database has its name (among other properties)
 - ATLR, CMSONR, LCGR, INT11R, DEVDB11, TEST2, ...



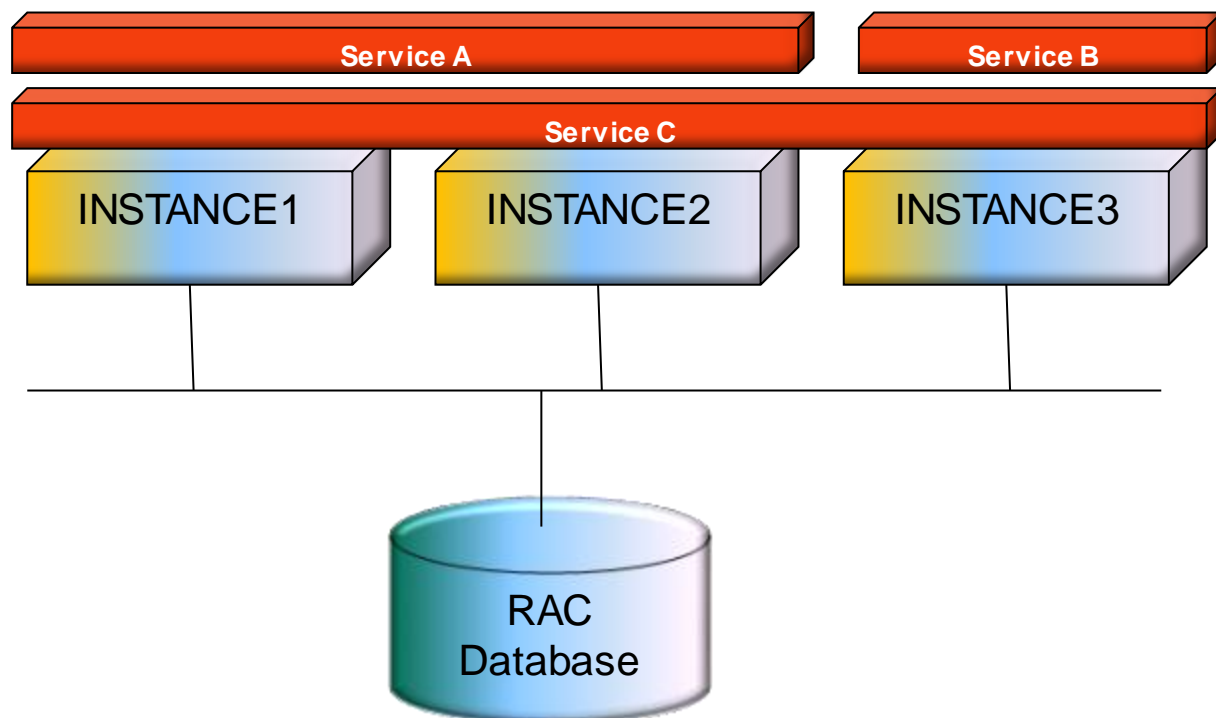
- **Block**
 - the smallest logical unit of data storage
- **Extent**
 - continuous set of blocks
- **Segment**
 - type (table, index, ...)
 - tablespace name
- **Tablespace**
 - name
 - maximal size
 - maps to one or more data files
- **Practical aspect:**
 - A user can have a limited quota on a tablespace
 - in most CERN databases (all physics) 1 block = 8KB
 - 1 database block - the smallest piece of data to be exchanged with physical storage





- Buffer cache
 - holds copies of data blocks
 - server processes manipulate data from buffer cache (not directly from data files)
 - managed using LRU algorithm (Least Recently Used)
- Shared pool
 - portion of the SGA that contains shared memory constructs such as shared SQL areas
 - most recently executed SQL statements
 - most recently used data definitions
 - Practical aspect:
 - no bind variables causes many shared SQL areas creation, which leads to high memory usage and fragmentation
 - ORA-04031:unable to allocate X bytes of shared memory

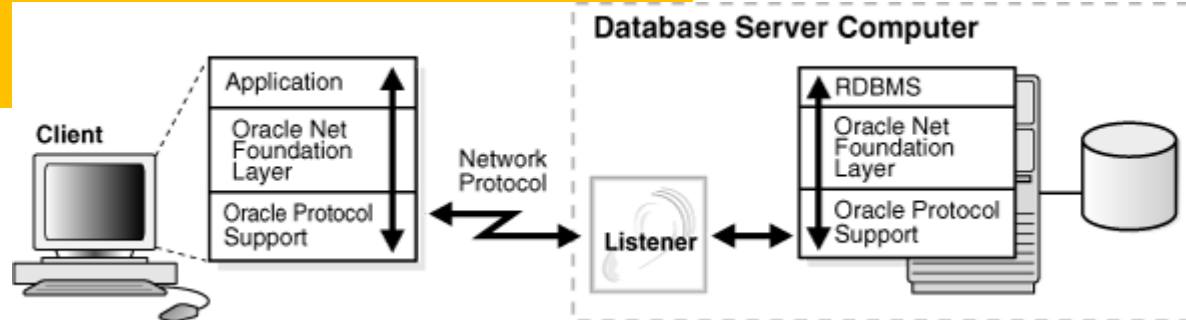
- High Availability (HA) + Scalability solution
- One database, multiple instances
 - managed by Oracle Clusterware
 - database services vs. database instance



- Practical aspects:
 - rolling interventions
 - sessions are transparently reconnected to working instance
 - transactions are broken and need to be rolled back
 - internal information sometimes bound to instance not visible immediately in GUI tools (EM, SM)
 - connection specification can cause losing of HA features, e.g. if instance specified instead of service
 - RAC aware applications can gain a lot, not RAC aware applications can lose a lot running on RAC

- Oracle Net manages connections between user process and server process
 - On top of popular network protocols (e.g. TCP, Inter-Process Communication (IPC))
- Oracle Net address

```
sqlplus user/password@'(
  DESCRIPTION=
  (ADDRESS= (PROTOCOL=TCP) (HOST=dbsrv3305.cern.ch) (PORT=10121))
  (CONNECT_DATA=
    (SID=DEVDB11)
  )
)'
```



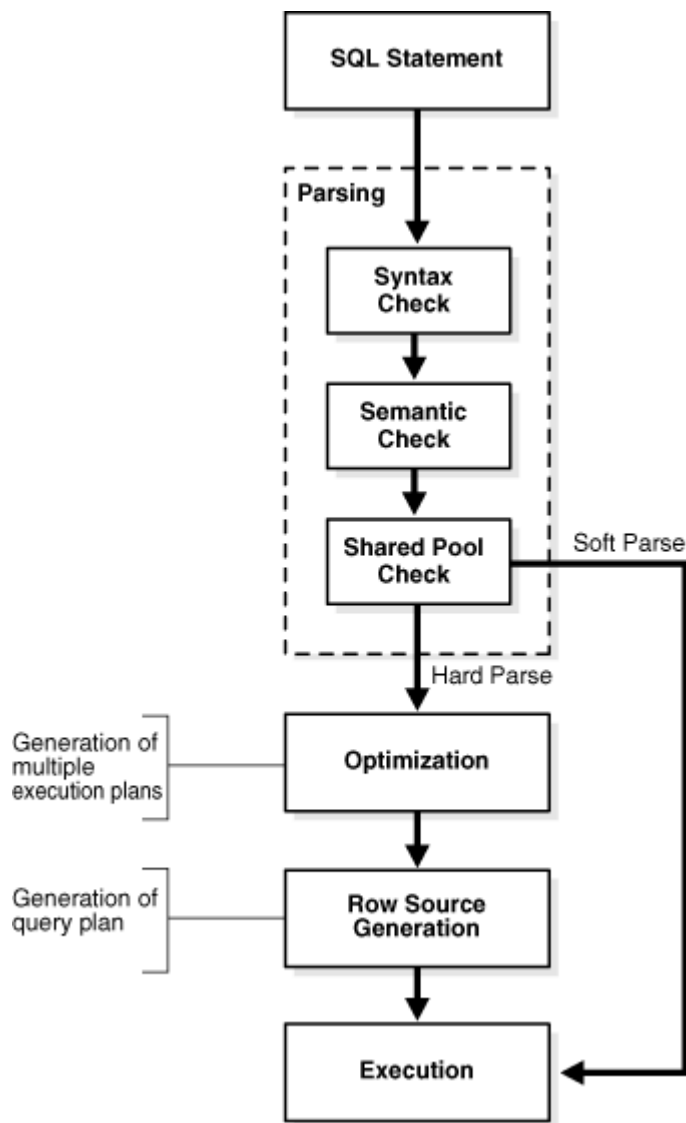
- Local naming method
 - Based on \$TNS_ADMIN/tnsnames.ora

```
lcgr_backup =
```

```
(DESCRIPTION =  
  (ADDRESS = (PROTOCOL = TCP)(HOST = lcgr1-v.cern.ch)(PORT = 10121))  
  (ADDRESS = (PROTOCOL = TCP)(HOST = lcgr2-v.cern.ch)(PORT = 10121))  
  (ADDRESS = (PROTOCOL = TCP)(HOST = lcgr3-v.cern.ch)(PORT = 10121))  
  (ADDRESS = (PROTOCOL = TCP)(HOST = lcgr4-v.cern.ch)(PORT = 10121))  
  (ADDRESS = (PROTOCOL = TCP)(HOST = lcgr5-v.cern.ch)(PORT = 10121))  
  (CONNECT_DATA =  
    (SERVER = DEDICATED)(SERVICE_NAME = lcgr_backup.cern.ch)  
  )  
)
```

```
sqlplus username/password@lcg_backup
```

- At CERN, centrally managed tnsnames.ora stored in:
 - \\cern.ch\dfs\Applications\Oracle\ADMIN
 - /afs/cern.ch/project/oracle/admin



- SQL statement processing phases:
 - Parse
 - Execute
 - Fetch
- Parse paths:
 - Soft parse
 - Hard parse
- Practical aspect:
 - Bind variables make SQL statements similar enough to use soft parse path

```
select name from emp where id=:id_val;
```

- Oracle Database Concepts

http://docs.oracle.com/cd/E11882_01/server.112/e25789/toc.htm

- Oracle 11g documentation

http://www.oracle.com/pls/db112/portal.all_books