INFM309

Oracle & Java

02. Oracle DB internals

Александър Станев
alex at stanev.org
http://stanev.org

Contents

- Oracle database internals
- Developing PL/SQL program units (based on Oracle University course programme)
- Web applications with Oracle
- Vivisection of eService – reallife case

Oracle Architecture

- Oracle Database Structures
- Physical Database Structures
- Oracle Processes
- Oracle Memory Structures

Oracle Database Structures

- ***Instance*** - collection of processes and utilities used to update, retrieve, and manage the data includes processes and system global area (SGA) memory a memory resident file
- ***Database*** - collection of files that store and organize data
- ***Tablespace*** - logical collection of data files; comprised of segments
- ***Segments*** - hold various types of data (e.g., tables, indexes, rollback); comprised of extents
- ***Extents*** - logical data entity; comprised of data blocks
- ***Blocks*** - smallest unit of data used by Oracle. Block size is determined when the database is first created

Key Point

Block size affects the efficiency of the database operation. The database administrator should determine the optimum block size considering operating system file system buffer size and database usage.

Tablespaces

- ***System tablespace*** - includes system data needed by the database to manage data includes data dictionary
- ***Temp tablespace*** - temporary clipboard used to manage transactions such as sorting data from an ORDER BY clause in a query
- ***Tools tablespace*** - stores objects used by tools that interact with the database
- ***User's tablespace*** - stores a personal user's objects
- ***Rollback tablespace*** - stores rollback segments; rollback segments are used by the database to roll back data (i.e., a transaction failure or user command to undo any uncommitted changes)
- ***Data and index tablespaces*** - store application data

Physical Database Structures

- ***Oracle data files*** - all files that store the database data includes rollback segments, redo logs, audit files and files that form basis for tablespaces and indexes. Typically have dbf extension
- ***Oracle software files*** - typically hold the database server's binary code and all other program code together constitute the database functionality includes library, Java and listener files
- ***Parameter files*** - store configuration information about the database server and the database instance includes location of control files, log files and init.ora file
- ***Control files*** - small binary file that contains information about the database instance and is needed to start the database. All major changes to the structure of the database are recorded in the control file
- ***Log/trace files***- log files contain sequential list of all changes to database; files trace files are generated by the database to facilitate troubleshooting. The file alert.out stores any server messages that are generated by the database server

Oracle Database Files

Init.ora - generic name of start-up parameter file for Oracle. File is read every time an Oracle database is opened to set a number of parameters including control backup and recovery, transaction journaling, security and audit logging, performance attributes and file locations. In addition, the server parameter file (SPFILE) may be used to set initialization parameters.

Oracle Database Files

The control files (named CONTROL_nn.CTL) are the key files that the Oracle DBMS maintains about the state of the database and includes:

- Checkpoint progress records
- Redo thread records
- Log file records
- Data file records
- Tablespace records
- Log file history records
- Archived log records
- Backup set and data file copy records

Warning

If the control file is corrupt or missing, it is very difficult to start the database instance. It is good practice to have at least two copies of the control file.

Oracle Processes

Each instance has a set of processes that interacts only with the data files associated with that particular instance. On UNIX systems, Oracle processes are referred to as background processes. On Microsoft NT/2000 systems, Oracle processes run as threads within a single system process.

Processes

- ***DBWn*** - writes data to the data files; up to ten of these can exist
- ***LGWR*** - reads and writes to the redo logs
- ***PMON*** - monitors for user sessions that are prematurely disconnected and handles cleanup
- ***SMON*** - manages database recovery that may be required at start-up
- ***ARCO*** - copies redo log files to the archive file destination
- ***CKPT*** - keeps track of the latest redo log files used, which are used during a database recovery
- ***RECO*** - handles distributed transactions that have failed
- ***LCK0*** - manages instances in a parallel server configuration
- ***SNPn*** - runs jobs from the database job queue, where n=1-10
- ***QMNn*** - manages message queues, where n=1-10

View processes

```
Connected to Oracle Database 10g Enterprise Edition Release 10.2.0.2.0
Connected as alex
SQL> select * from v$bgprocess where paddr <> '00';

PADDR    PSERIAL# NAME     DESCRIPTION
BCBF2058 1        PMON     process cleanup  0
BCBF2648 1        PSP0     process spawner 00
BCBF2C38 1        MMAN     Memory Manager   0
BCBF3228 1        DBW0     db writer process 0       0
BCBF8548 1        ARC0     Archival Process 0        0
BCBF8B38 1        ARC1     Archival Process 1        0
BCBF9128 1        ARC2     Archival Process 2        448
BCBF3818 1        LGWR     Redo etc.         0
BCBF3E08 1        CKPT     checkpoint        0
BCBF43F8 1        SMON     System Monitor Process    0
BCBF49E8 1        RECO     distributed recovery      0
BCBF4FD8 1        CJQ0     Job Queue Coordinator     0
BCBF9718 3        QMNC     AQ Coordinator   0
BCBF6D88 1        DMON     DG Broker Monitor Process 0
BCBF55C8 1        MMON     Manageability Monitor Process     0
```

Key Point

It is important to ensure that the key database processes are running on the database server. The termination of these processes can lead to availability, recoverability and connectivity issues. In many client environments, the monitoring of operating system processes is handled by system administrators and does not fall under the responsibility of database administrators. In these environments, it is imperative that database administrators ensure that there is an automated system in place to monitor key database system processes and notify appropriate personnel in the event that they are not running.

Oracle Memory Structures

- ***System Global Area (SGA)*** – stores pertinent information about instance
- ***Program Global Area (PGA)*** – contains data and control information for a single process; used by a single Oracle process and is not shared among processes

System Global Area

- ***Database buffer cache*** – memory buffers that the database can use to hold data that the server process reads from the data files on disk in response to user's requests
- ***Shared pool*** – holds executable PL/SQL code and SQL statements, as well as information regarding the data dictionary tables
- ***Library cache*** – component of shared pool holding already parsed and ready-to-execute form of SQL statements; shared by all database users
- ***Data Dictionary cache*** – component of shared pool primarily contains object definitions, usernames, roles, and privileges
- ***Redo log buffer*** – holds redo data prior to its writing to the redo logs

System Global Area

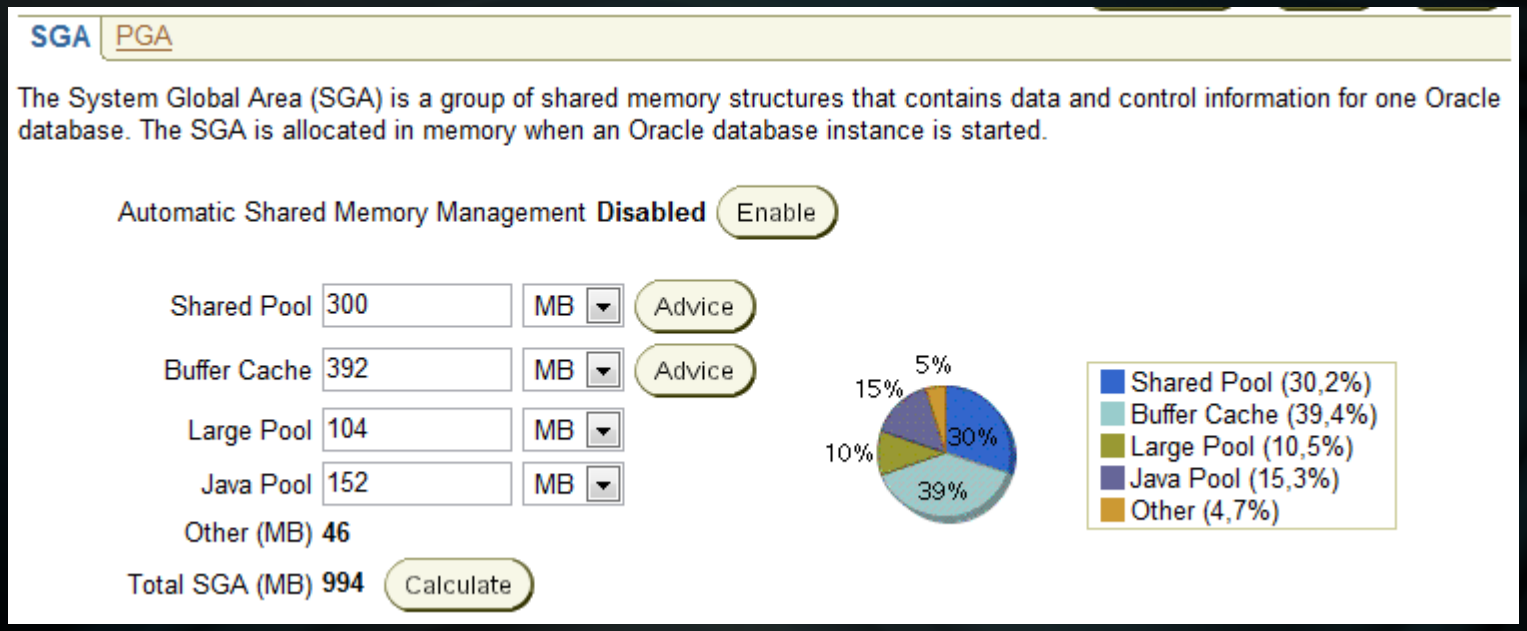

Program Global Area

- ***Private SQL area*** – Data for binding variables and runtime memory allocations
- ***Session memory*** – Memory that holds session variables and other session information.
- ***SQL work area*** –Memory allocated for sort, hash-join, bitmap merge, and bitmap create types of operations.
- ***Data Dictionary cache*** – component of shared pool primarily contains object definitions, usernames, roles, and privileges
- ***Redo log buffer*** – holds redo data prior to its writing to the redo logs

Program Global Area

Program Global Area

- Memory reserved for each user process connecting to an Oracle database
- Allocated when a process is created
- Deallocated when the process is terminated
- Used by only one process

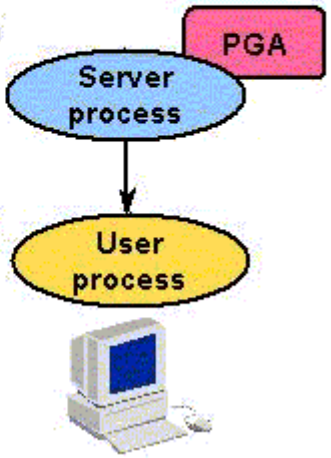

Connecting to Oracle instance

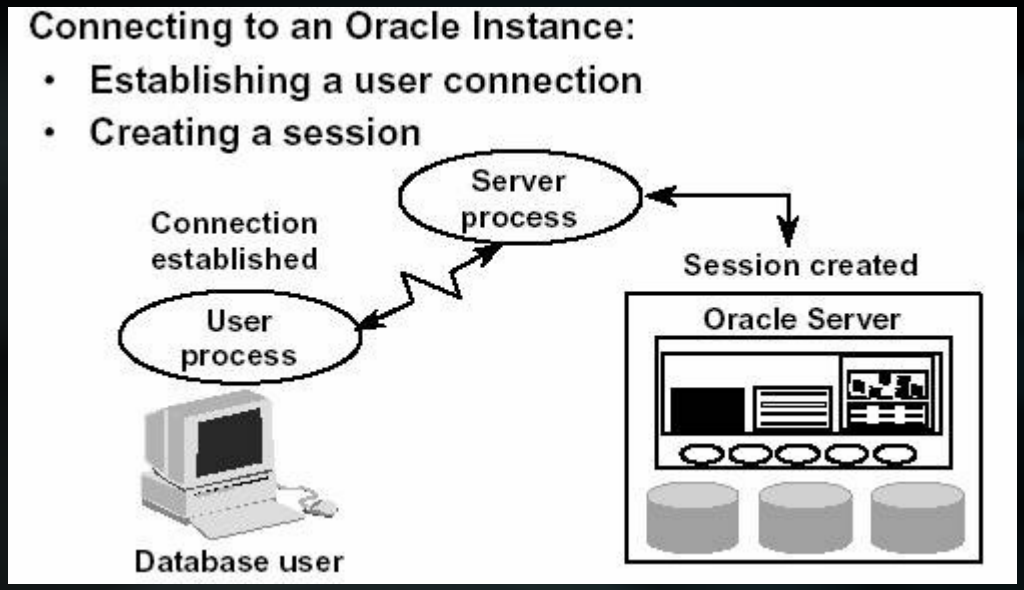

This is **Dedicated server** connection – one user process to one server
Shared server – many user processes to one server

Oracle Database Architecture under Windows

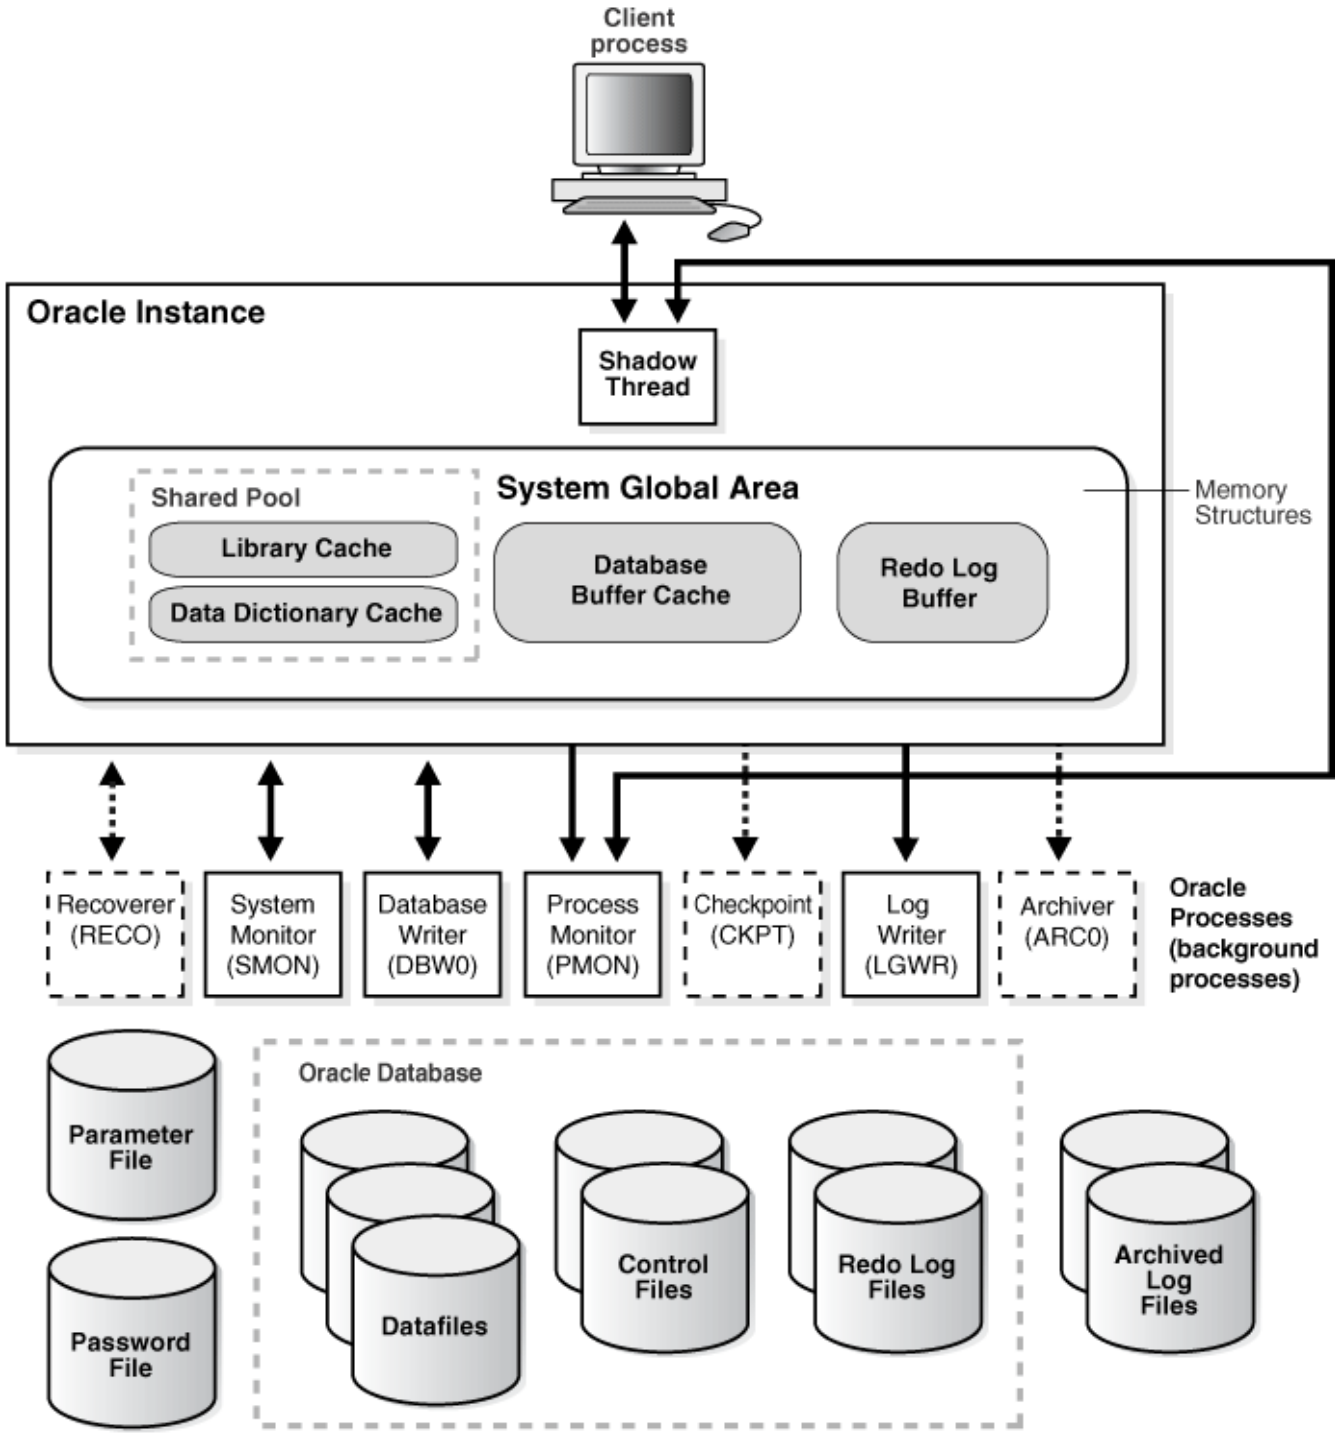

Oracle Database Objects/Terminology

- ***Schema*** - A collection of logical structures of data, or schema objects. A schema is owned by a database user and has the same name as that user.
- ***Database links*** - An object that resides in the local database and uniquely defines the path to the remote database. A DBLink is required for a local database to connect to a remote database.
- ***Indexes*** - A database object that provides fast access to individual rows in a table. Once created are automatically maintained and used for data access by the database engine whenever possible.
- ***Sequences*** - A schema object that generates sequential numbers. After creating a sequence, you can use it to generate unique sequence numbers for transaction processing.
- ***Synonyms*** - An alias for a table, view, sequence or program unit. A synonym is not actually an object itself; rather, it is a direct reference to its base object.

Oracle Database Objects/Terminology (cont.)

- ***Stored Procedures*** - A set of SQL and PL/SQL statements grouped together as an executable unit to perform a specific task. Unlike functions, procedures do not return a value to the caller.
- ***Triggers*** - similar to stored procedures (collection of SQL statements) but are embedded in the database and activated when certain conditions are met
- ***Functions*** - A PL/SQL subprogram that executes an operation and returns a single value at the completion of the operation. A function can be either built-in or user-named.
- ***Packages*** - A method of encapsulating and storing related procedures, functions and other package constructs together as a unit in the database. While packages provide the database administrator or application developer organizational benefits, they also offer increased functionality and database performance.

Oracle Database Objects/Terminology (cont.)

- ***Tables*** - The basic unit of storage in a relational database management system. A table represents entities and relationships and consists of one or more units of information (rows) each of which contains the same kinds of values (columns).
- ***Clusters*** - Two or more tables that are physically stored together to take advantage of similar columns between the tables
- ***Snapshots*** - Copies of tables or views.
- ***Views*** - Tailored presentation of the data contained in one or more tables. A view is defined using a query and can be thought of as a stored query or virtual table.
- ***Columns*** - Fields of the table.
- ***Rows*** - Records of the table.

Страница на курса

Слайдовете, материалите и допълнителна информация за курса могат да бъдат открити на адрес:

http://alex.stanev.org/training/nbu-infm309

Въпроси

TPILB.

This page intentionally left blank.