

Oracle 8i User Guide for CS2312

Carole Goble

1. Introduction

Oracle is a relational database management system that uses SQL as its data manipulation language. Information on SQL is given separately to these notes. These notes are a fast-start to the way we will use Oracle in the labs and supplement the *SQL*Plus User's Guide and Reference manual* available on the web.

The Oracle database resides on a server. This is where the binaries are and your databases. Your database tables DO NOT reside in your own file store and can only be accessed by using an Oracle tool. When you log onto a machine and start an Oracle tool the server must be notified to serve the tool to your machine. This can be slow. However, once the tool is loaded onto the client it runs the program locally so this is faster. Once you have started a tool don't close it down unless you are finished.

You will be developing your database implementations in pairs. To avoid confusion and wasted disk space do not create multiple tables with the same things in them. There should only be *one* team database, looked after by *one* person and the other should be granted permission to use that database (see later on how to do that).

2. Oracle

2.1 Oracle Architecture

Oracle, very simplistically, can be viewed as having a nested architecture (figure 1)

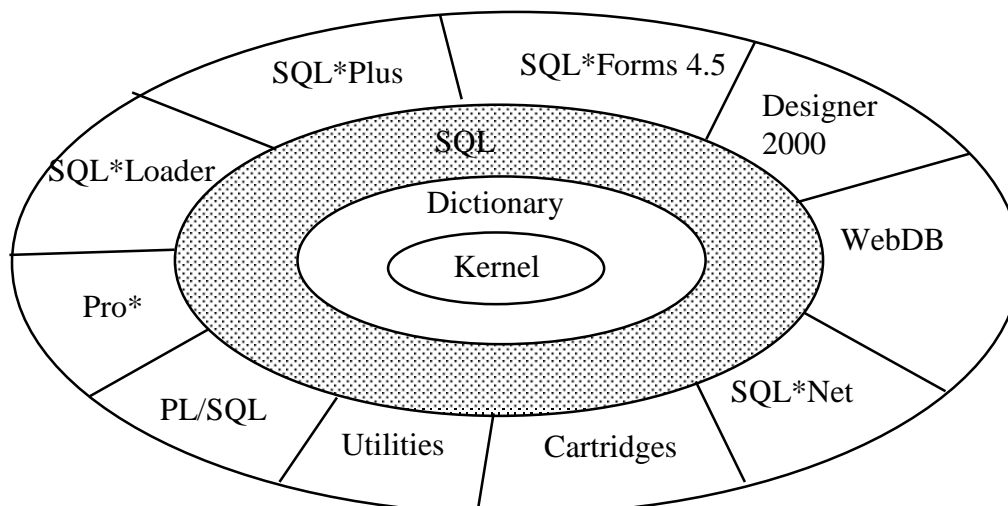


Figure 1: Basic architecture of Oracle

2.1.1 The Kernel

The internal code that communicates with the database, and handles all the facilities of DBMS: retrieval, recovery etc. It also optimises request execution to be more efficient, and deals with the accelerators and physical data storage exploiting the facilities of the host operating system, which in our case is Unix.

2.1.2 The Data Dictionary

Describes in detail the Oracle objects. It is used for documentation and management of the objects. Holds the schema of the tables and views. The description of database objects is often called metadata (data about data). The dictionary is a set of tables and views manipulated through SQL. Users can look at the content of the tables through a series of views. Only the DBA can actually manipulate the dictionary directly (see figure 2).

2.1.3 SQL layer

The interface between the kernel and the tools of the external layer. Every access operation to the data is expressed in SQL. SQL commands include:

- Data Definition commands to define the schema;
- Data Manipulation commands to manipulated the *contents* of the schema;
- Data Control commands to grant users privileges and access rights.

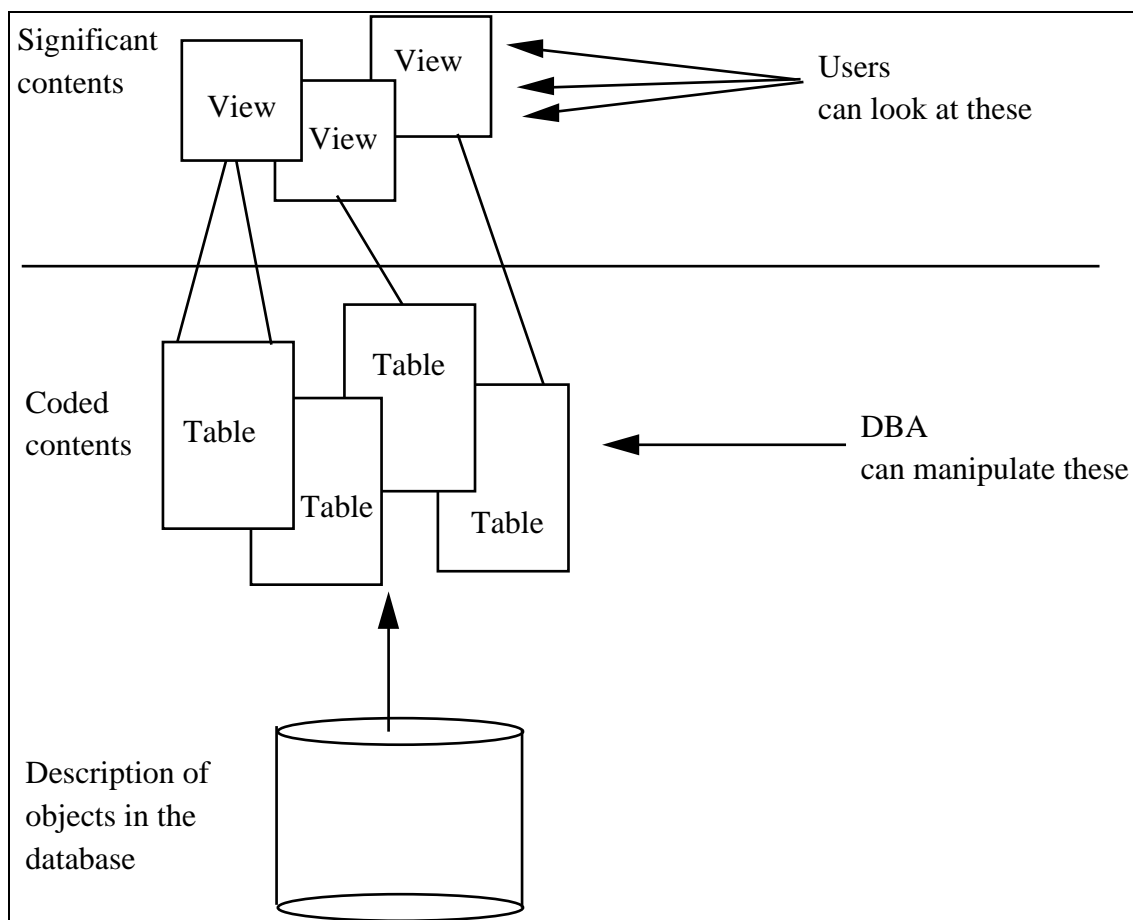


Figure 2: The Data Dictionary

2.1.4 The External Layer

A set of tools that enable all the facilities provided by the three basic underlying layers to be used easily. It is an applications development environment. We will only use SQL*Plus. This is the interactive interface to SQL, so you can put SQL commands directly to the database. Has a work environment that can be parameterise; has commands for the formatting of results and commands to store SQL and SQL*Plus commands in files.

2.2 Oracle objects

Oracle manipulates many objects, of which we are interested in five:

- Tables
- Views
- Integrity constraints
- Users and their privileges
- Accelerators on tables (indexes, clusters) -- these are data structures which give a fast access path to tables, a bit like an index in a book. See Chapters 4 & 5 of Elmasri and Navathe. We will cover these later in the course.

In Oracle, a user cannot explicitly define an object called a database; a user's database is the set of tables and views that a user has access privileges to. A user may access table and views other than those created by themselves by prefixing them with the name of that user, provided they have been granted the privilege of doing so.

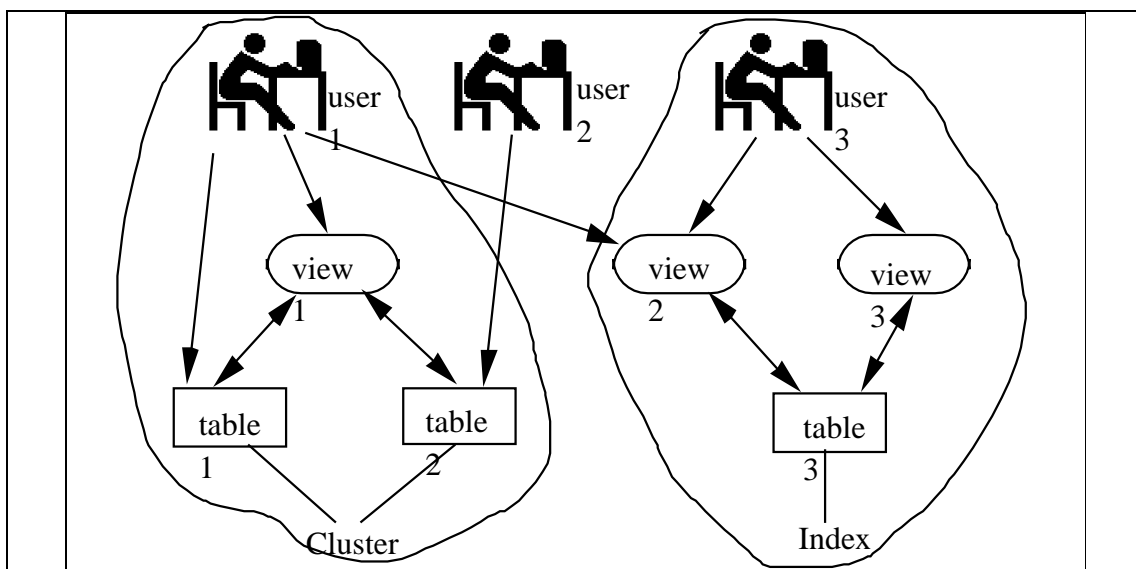


Figure 3: Objects manipulated by Oracle

In Figure 3, user1 has created two tables, *table1* and *table2*, and a view *view1*. The two tables are grouped in the same cluster, which means that they are stored internally close together on the disk because they are often accessed together. This same user can access table *table3* via view *view2*. User3 has created a table *table3*, which uses an index (a data structure which makes accessing data in it faster), and two views *view2* and *view3* on this table. Finally user2 has created no table or view, but has the privilege of accessing table *table2* by prefixing it with the name of user1.

2.2.1 User Privileges

Users are defined by a name (which is unique), a password and a set of privileges. You have all been given the privilege to create, modify and delete database objects (called RESOURCE privilege). The way you will build tables for your databases is as follows:

1. Design your tables

2. One of your team, user A, creates the tables
3. User A grants permission to the other member of the pair, user B, the privilege to access the tables using the SQL GRANT statement
4. User B will have to use user A's tables by prefixing the table name with user A's oracle user name. To get round this, users B should create SYNONYMS for user A's tables and views.
5. All users can create views, and organise which users have permission to use the views.

2.2.2 Databases and Tablespaces

The way Oracle has been set up on the SUN server is on a disk local to it. There is only one Database, run by the Data Base Administrator. The database is divided into six tablespaces. Tablespace project users are given permission to create objects in tablespace cags_course. Tablespace SYSTEM is where the Data Dictionary is held (figure 4).

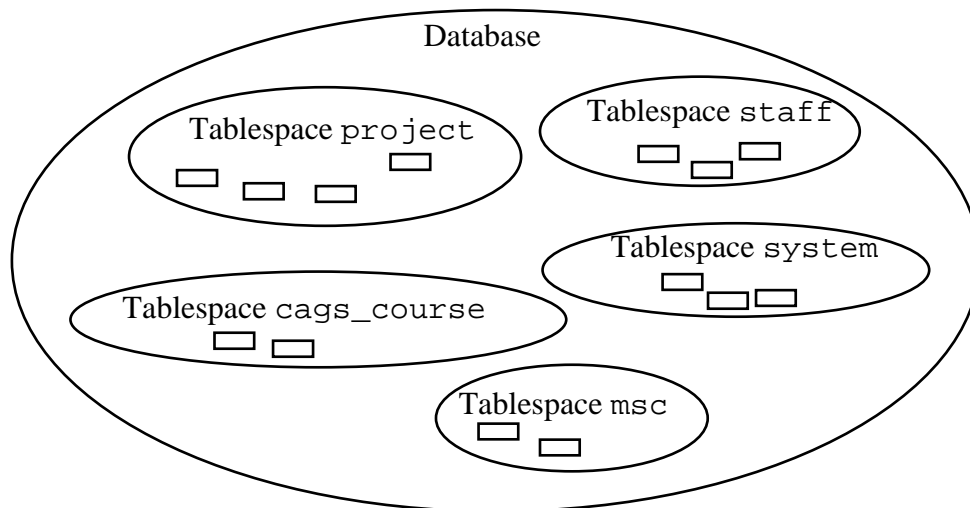


Figure 4: Tablespaces

3. Running SQL commands

Oracle requires each user to be registered as a user of Oracle with a user name and a password. This is separate to the Unix system. However, we have set up your Oracle username and password to be the same as your Unix ones automatically. All of the Oracle paths, environment variables and aliases should have been set up automatically with your account. If you have any problems using a command (for example the command cannot 'be found') then contact the lab supervisor.

Sometimes SQL commands refer to the user name of an Oracle user. Usually your user name will be the same as the Unix one.

There are two types of commands:

SQL commands:

the commands for controlling, defining, manipulating and querying the database tables. The *SQL Language Reference Manual* is the complete manual of the whole of the Oracle SQL language. The SQL language is used for creating, manipulating, updating, querying and deleting your tables and the data dictionary tables. It is very useful and is available on the web. For CS2311 you can ignore any statement that has the following keywords when referring to the manual: Rollback, Databases, Tablespace, Storage and blocks, Clusters and indexes, Distributed.

SQL*Plus commands:

additional commands governing the SQL*Plus environment, formatting or results and measuring performance. The complete manual is on the web. I have included in these notes the basics with a page number of the Guide for further details. When reading this Guide ignore any reference to PL/SQL and SQL*Forms.

Oracle is not case-sensitive, which means you can use upper or lower case for commands, names of tables and so on, but Oracle won't tell any difference between STUDENT, Student or student.

6. Backing Up

We do keep a backup. However, for an immediate fix if you muck up a relation is to use a SQL batch file to create tables, and a SQL batch file for loading the tuples. Then keep these files in a safe place in the teams collective file store. Then, if all goes wrong you can rebuild from scratch. Otherwise you will have to ask the DBA to restore your table which is rather difficult.

To compress these files so they use up less of your filestore quota, look at the Unix commands 'compress' and 'uncompress'.

7. The Oracle Catalog and Data Dictionary

Oracle has a data dictionary containing all the management information on the objects in the database. You cannot alter this dictionary directly, but you can look at it: the data

is held as tables so the usual SQL commands work. Appendix B of the *Oracle8 Server Administrator's Guide* has a list of all the Data Dictionary Views and the Oracle reserved words. The most interesting Data Dictionary Views for you are:

ALL_CATALOG	Tables, views, synonyms accessible to the user
ALL_CONSTRAINTS	Constraints on tables accessible to the user
ALL_CONS_COLUMNS	Information about accessible columns in constraint definitions
ALL_SYNONYMS	Synonyms accessible to the user
ALL_TABLES	Tables accessible to the user. This gives more information than ALL_CATALOG
ALL_VIEWS	Views accessible to the user
DICTIONARY (or DICT)	Data Dictionary tables and views
TABLE_PRIVILEGES	Grants on objects for which the user is the grantor, grantee or owner, or PUBLIC is the grantee
USER_CATALOG (or CAT)	Tables, views, synonyms and sequences owned
USER_COL_PRIVS	Grants on columns for which the user is the owner, grantor or grantee
USER_COL_PRIVS_MADE	Grants on columns of objects owned by the user
USER_COL_PRIVS_RECD	Grants on columns for which the user is grantee
USER_CONSTRAINTS	Constraints owned by the user
USER_CONS_COLUMNS	Columns in constraints owned by the user
USER_SYNONYMS	The user's private synonyms
USER_TABLES	Tables owned by the user
USER_TAB_COLUMNS	Columns of the user's tables, views and clusters
USER_TAB_PRIVS	Grants on objects for which the user is the owner, grantor or grantee
USER_TAB_PRIVS_MADE	Grants on objects owned by the user
USER_TAB_PRIVS_RECD	Grants on objects for which the user is the grantee
USER_USERS	Privileges of the current user
USER_VIEWS	Views owned by the user