



Introduction

The Oracle Relational Database Management System (RDBMS) has become one of the most powerful and flexible databases available. With ever-expanding functionality provided by options such as Data Guard replication and Real Application Clusters (RAC), it has increasingly blurred the divide between database administrators and system administrators. Nowadays it's becoming essential that sysadmins have a solid understanding of Oracle in order to interface effectively with database admins and architects.

This booklet won't teach you everything about Oracle, and it's not intended to. Instead, you're going to get an overview of the major components of Oracle, specifically Oracle10g. Even though you might never export an Oracle database, it's essential that you understand what it is, so that when your database admins do it, you'll know what they are doing and how to provide for their needs. Throughout the booklet you'll find URLs to Oracle documentation which will provide you with more information on a given topic.

By the end of this booklet you'll have the confidence to log in to Oracle, configure basic networking, manipulate user access, understand the programming interfaces, utilize major tools for importing and exporting data, and interface with the backup system—and, most important, you'll have enough knowledge to learn more about Oracle on your own.

The Relational Model

In this day and age it's easy to take the "simplicity" of the modern relational database for granted. If when you first learned about relational databases you said to yourself, "This is just a bunch of spreadsheets!" you'd have been more or less right. What makes a relational database special is how relationships between parts of the data are managed. In 1985 Dr. E.F. Codd, an IBM researcher, published a list of rules known as "Codd's 12 Rules" that defined how a true RDBMS should be evaluated. Understanding these rules will greatly improve your ability to understand RDBMSes in general and Oracle in particular:

1. Information is represented logically in tables.
2. Data must be logically accessible by table, primary key, and column.
3. Null values must be uniformly treated as "missing information," not as empty strings, blanks, or zeros.

2 / Introduction

4. Metadata (data about the database) must be stored in the database just as regular data is.
5. A single language (usually SQL) must be able to define data, views, integrity constraints, authorization, transactions, and data manipulation.
6. Views must show the updates of their base tables and vice versa.
7. A single operation must be able to retrieve, insert, update, or delete data.
8. Batch and end-user operations are logically separate from physical storage and access methods.
9. Batch and end-user operations can change the database schema without having to recreate it or the applications built upon it.
10. Integrity constraints must be available and stored in the relational database metadata, not in an application program.
11. The data manipulation language of the relational system should not care where or how the physical data is centralized or distributed.
12. Any row processing done in the system must obey the same integrity rules and constraints that set-processing operations do.¹

Of these rules the most difficult one to come to terms with is rule 4: this is where *system tables* or *data dictionaries* come in and why it's so difficult to start exploring the database without a basic understanding of the relational database design rules. It's not terribly different from building a Web server that must be configured using a CGI; while it isn't intuitive, it is an admirable design.

You by no means need to bother memorizing this list of rules, but if they're stashed in your long-term memory, retrieval may help you to understand some design decisions used by modern relational databases that might otherwise seem odd.

Play Along at Home

All of the examples and code used throughout this booklet were done using Oracle10g Enterprise Edition on Solaris 10 (Sun Ultra2 Dual UltraSPARC II workstation, 512MB memory) and Gentoo Linux (AMD AthlonMP Dual 1.2Ghz, 1GB memory, Kernel 2.6.8.1). The Oracle Database is available as a "free" download from Oracle.com. Oracle provides you with a 30-day trial period by which to test and explore the product. Please note that although the software will not stop you from running Oracle longer than that, it is illegal to do so.

1. I was unable to acquire a copy of Codd's *The Relational Model for Database Management, 2nd ed.*; these rules are quoted from Kevin Kline and Daniel Kline, *SQL in a Nutshell* (O'Reilly, 2004).