

2

Oracle

October 1997

In this issue

- 3 Joining character and number columns
- 4 Oracle database object dependency
- 16 Self-documenting 'exists' clauses
- 17 Running batch jobs in Unix
- 18 The Oracle environment
- 26 Performance tuning and capacity planning – part 2
- 31 Space usage monitoring
- 33 Very Rapid Application Development – part 2
- 44 Oracle news

© Xephon plc 1997

update

Oracle Update

Published by

Xephon
27-35 London Road
Newbury
Berkshire RG14 1JL
England
Telephone: 01635 38342
From USA: 01144 1635 38342
E-mail: xephon@compuserve.com

North American office

Xephon/QNA
1301 West Highway 407, Suite 201-405
Lewisville, TX 75067
USA
Telephone: 940 455 7050

Australian office

Xephon/RSM
GPO Box 6258
Halifax Street
Adelaide, SA 5000
Australia
Telephone: 088 223 1391

Subscriptions and back-issues

A year's subscription to *Oracle Update*, comprising twelve monthly issues, costs £165.00 in the UK; \$245.00 in the USA and Canada; £171.00 in Europe; £177.00 in Australasia and Japan; and £175.50 elsewhere. In all cases the price includes postage. Individual issues, starting with the September 1997 issue, are available separately to subscribers for £14.00 (\$20.50) each including postage.

Editor

Trevor Eddolls

Disclaimer

Readers are cautioned that, although the information in this journal is presented in good faith, neither Xephon nor the organizations or individuals that supplied information in this journal give any warranty or make any representations as to the accuracy of the material it contains. Neither Xephon nor the contributing organizations or individuals accept any liability of any kind howsoever arising out of the use of such material. Readers should satisfy themselves as to the correctness and relevance to their circumstances of all advice, information, code, PL/SQL, and other contents of this journal before making any use of it.

Contributions

Articles published in *Oracle Update* are paid for at the rate of £170 (\$250) per 1000 words and £90 (\$140) per 100 lines of code for original material. To find out more about contributing an article, without any obligation, please contact us at any of the addresses above and we will send you a copy of our *Notes for Contributors*.

Oracle Update on-line

Code from *Oracle Update* can be downloaded from our Web site at <http://www.xephon.com>; you will need the user-id shown on your address label.

© Xephon plc 1997. All rights reserved. None of the text in this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, without the prior permission of the copyright owner. Subscribers are free to copy any code reproduced in this publication for use in their own installations, but may not sell such code or incorporate it in any commercial product. No part of this publication may be used for any form of advertising, sales promotion, or publicity without the written permission of the publisher. Copying permits are available from Xephon in the form of pressure-sensitive labels, for application to individual copies. A pack of 240 labels costs \$36 (£24), giving a cost per copy of 15 cents (10 pence). To order, contact Xephon at any of the addresses above.

Printed in England.

The Oracle environment

INTRODUCTION

Before attempting to effectively administer the Oracle server environment, it is worthwhile examining the architecture of Oracle in more detail. In other words, let's see 'what makes Oracle tick'. This can best be achieved by investigating the basic 'architectural' blueprint of Oracle and examining its core components.

The Oracle architecture has five basic components:

- File structures
- Memory structures
- Processes
- Rollback segments
- Redo logs.

These core components operate in an integrated manner to provide support for the client tasks.

INSTANCE VERSUS DATABASE

An Oracle instance is the combination of all of the memory structures and background processes that are allocated when an Oracle database is started. Often, Oracle users tend to confuse an Oracle instance with an Oracle database.

An Oracle database has both physical structures (data files) and logical structures (table, index). Oracle separates the physical from the logical structure so that the physical storage of data may be manipulated without affecting user access to the logical structures. The physical structure of an Oracle database is determined by the files created at the operating system level by the Oracle instance during database creation (controlfile, logfile) or by the DBA during normal operation (create tablespace, create control file).

The three physical file structures that comprise an Oracle database are:

- Control file – a small administrative file that is used by the Oracle database.
- Redo log files – record changes made to data. The redo log files are used in recovery to ensure that no data is lost should a failure prevent changed data from being written to disk.
- Database files – contain the database information, including both system and application data.

ORACLE PARAMETER FILE

An Oracle parameter file contains all the configuration parameters for an Oracle instance. These parameters can be set to different values to tailor the Oracle instance to the system's operating configuration as well as the needs of the individual application or applications contained in the Oracle database. The parameter file specifications affect both memory and process settings of the associated Oracle instance. All parameters contained in the parameter are optional and will default if omitted (refer to the *Oracle DBA Guide* for default values).

ORACLE DATABASE FILES

Oracle database files contain the data associated with a particular database. All of the files discussed are not absolutely necessary for normal database operations, but this configuration is highly recommended for a well-designed efficient environment. Oracle database files can be grouped according to specific categories:

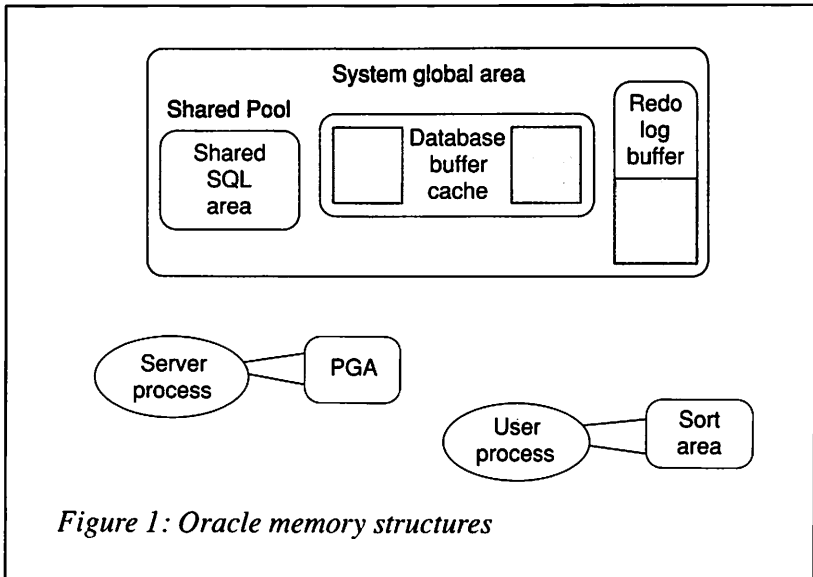
- Control files that record the physical structure of the Oracle database. The control file is read during every database start-up. It is stored in binary format and cannot be edited by normal means; however, it will be automatically modified when database changes are made.
- Data files that are associated with Oracle tablespaces including:

- System data files, which contain the contents of the system tablespace. The system tablespace consists of the data dictionary, and the names and locations of all the tablespaces, tables, indexes, triggers, procedures, and clusters for the database.
- Application data files and default data files, which are physical database structures. These files hold user data that is accessed by clients in the course of normal database operation. The difference between default and application data files is that the default tablespace/datafile combination is used when an explicit specification is not made.
- Temporary files, which contain temporary data that is used by queries and sorts. Regardless of their name, temporary data files are actually permanent files created to provide a transient work area for Oracle to sort and join tables and to execute complex SQL.
- Oracle rollback files, which contain information related to the rolling back of uncommitted data during error processing.
- Redo log files that record changes made to data. Every Oracle database has at least one set of two or more redo log files. These files record changes made to data and are used in recovery circumstances to prevent lost or damaged data.
- The CONFIG.ORA file is associated with the Oracle client, not the server. It specifies certain defaults, eg file and directory names for the client portion of the Oracle environment.

ORACLE MEMORY STRUCTURES

Oracle utilizes specific memory structures to perform DBMS-related tasks. These memory structures are contained in the main memory (RAM) of the computer running the Oracle instance.

The basic memory structures for an Oracle instance are the SGA (System Global Area), PGA (Program Global Area), and Oracle Sort Areas. This is illustrated in Figure 1.



The SGA is a group of shared memory structures allocated by the Oracle instance to contain data and control information. The SGA is automatically allocated when the instance is started and deallocated when the instance is shut down. If more than one user is connected to a multiple-process database, the information in the SGA is shared among the different users.

The SGA contains the following components:

- Database buffer cache – a staging area for database reads and writes, holding the most recently used data blocks.
- Redo log buffer – similar to the database buffer cache, the redo log buffer is a staging area for redo log entries.
- Shared pool – an area of memory set aside for the parsing and processing of SQL statements (shared SQL area), staging of Data Dictionary reads and writes (Data Dictionary cache), and, if the Oracle configuration is multithreaded, binding information and run-time buffers (private SQL area).

The PGA is a memory structure containing information for both user and background processes. Each process has its own PGA. Various processes can use the PGA as an internal 'scratch pad' or work area. The contents of the PGA vary depending on the type of process (user or background) and Oracle configuration (multithreaded server or conventional server).

Every time a user process requests a sort to be performed, a special memory structure is allocated called an Oracle sort area. The sort area exists in the memory of the user process that is requesting the sort. User processes typically perform sorting when query result sets are ordered or grouped.

ORACLE PROCESSES

Oracle processes are the true workhorses of the Oracle instance. Each process has a series of tasks assigned to it and its own internal memory allocated to it (the PGA). This memory allocation allows the process to perform the internal processing required to perform its designated function. Oracle has two general types of process – user processes and Oracle processes.

- A user process is created to execute the program code of an application program.
- An Oracle process is called by another process to perform certain specific functions on behalf of the invoking process. Oracle processes can be further broken down into server processes and background processes:
 - Server processes are created to interact between the user processes and the Oracle instance. The server process communicates with the user process and acts as a 'relay' between the user process and SGA information.
 - Background processes perform designated functions that would otherwise have to be performed by each user process that interacts with the Oracle instance. This specialization of functions is much more efficient than if each user process was required to perform a variety of tasks.

Oracle uses the following background processes:

- The PMON (process monitor) background process performs 'clean-up duties' when a user process fails with an error condition. PMON cleans up the cache, releases locks, and carries out other miscellaneous tasks.
- The SMON (system monitor) background process provides instance recovery during start-up. SMON also cleans up temporary segments that are no longer in use, compresses free space into contiguous extents, and, in a parallel server environment, provides instance recovery for a failed CPU.
- The DBWR (database writer) writes database buffers contained in cache to disk.
- The LGWR (log writer) is a background process that manages the redo log buffer.
- If an Oracle instance takes more checkpoints than the LGWR process can efficiently handle, checkpoint duties can be turned over to the CKPT (checkpoint) background process.
- The ARCH (archiver) process performs log archival by copying on-line redo log files to auxiliary storage devices as logs become full.
- The RECO (recover) automatically resolves failures involving distributed transactions.
- The final types of process are user and server processes. The user process provides the communication link between the client application and the Oracle instance. The user process does not execute SQL statements. Those duties fall to the server process. The server process parses and executes SQL and performs all the tasks required to send results back to the requesting application.

ORACLE CONFIGURATIONS

There are two levels of software that must be executed by users accessing the Oracle server. The application or software tool on the

client must be executed to initiate the connection to the Oracle server and transport the desired SQL statement to be executed. The server process on the Oracle server receives the connection request and its associated SQL statement and begins the process of executing the desired statements.

The Oracle server can be configured as either a single or multiple process instance. The different configurations are based on the operating system on which the instance will run and the processing requirements of the applications accessing Oracle data.

Some operating systems (such as MS-DOS) cannot run multiple background processes concurrently. All of the code for the various background required by the instance must be run in a single background process. These types of database can be accessed by only one user at a time. This requires a single process Oracle instance.

The most popular class of configuration, however, is the multiple process instance. There are three possible multiple process configurations:

- Combined user/server process – some operating systems (such as VAX VMS) can simultaneously execute the server process and the application's code in a single executable user process. The communication between the two processes running as one is done through the program interface. Although configured as a multiple process instance, only one Oracle connection at a time is allowed. Multiple user, concurrent access is not available.
- Dedicated server process – in this set-up, the background run separately. This is the most popular of all Oracle configurations. It is called the dedicated server process because each server process acts on behalf of only one user process.
- Multi-threaded server process – this process differs from the dedicated server process because several user processes have the capability of sharing one server process. A dedicated and shared server process perform the same functions.

