


Oracle Cluster POWERsolution Guide

December 1997



IBM

**International Technical Support Organization
Austin Center**

International Technical Support Organization

Oracle Cluster POWERSolution Guide

December 1997



Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix D, "Special Notices" on page 105

First Edition (December 1997)

This edition applies to IBM Interactive Network Dispatcher Version 1.2 for AIX.

Comments may be addressed to:

IBM Corporation, International Technical Support Organization
Dept. JN9B Building 045 Internal Zip 2834
11400 Burnet Road
Austin, Texas 78758-3493

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© **Copyright International Business Machines Corporation 1997. All rights reserved**

Note to U.S Government Users – Documentation related to restricted rights – Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

Contents	v
Figures	ix
Tables	xi
Preface	xiii
The Team That Wrote This Redbook	xiii
Comments Welcome	xiii
Chapter 1. Relational Database Concepts	1
1.1 Relational Databases	1
1.2 Database Structure	2
1.3 Structured Query Language (SQL)	2
1.3.1 SQL Statements	2
1.3.2 Context Areas and Cursors	3
1.3.3 Statement Parsing	3
1.4 Client/Server Architecture	4
1.4.1 Client/Server Connection	4
1.4.2 Network Connections	4
1.4.3 Shared-Memory Connections	4
1.5 Filesystems and Raw Logical Volumes	5
1.5.1 Filesystems	5
1.5.2 Raw Logical Volumes	5
Chapter 2. Oracle Concepts	7
2.1 Oracle Files	7
2.1.1 Data Files	7
2.1.2 Redo Log Files	9
2.1.3 Control File	9
2.1.4 The init.ora File	10
2.2 Oracle Memory Structure	10
2.2.1 System Global Area or Shared Global Area (SGA)	10
2.2.2 Program Global Area (PGA)	10
2.3 Oracle Processes	11
2.3.1 Database Writer (DBWR)	13
2.3.2 Log Writer (LGWR)	13
2.3.3 System Monitor (SMON)	13
2.3.4 Process Monitor (PMON)	13
2.3.5 Checkpoint (CKPT)	13
2.3.6 Recover (RECO) Process	14
2.3.7 Archiver (ARCH) Process	14
2.3.8 Parallel Server Lock Processes (LCKn)	14
Chapter 3. Database Planning, System Requirements, and Sizing	15
3.1 Planning Your Database	15
3.1.1 First Considerations	15
3.1.2 Database Application Type	15
3.1.3 Database Size	16
3.1.4 File Locations	17
3.2 System Requirements	17
3.2.1 Requirements for Oracle Product Installation on a Single Server	17

3.2.2	Requirements for Oracle Optional Products	17
3.3	Disk Space and Memory Sizing for Oracle	19
3.3.1	Required Calculations	19
3.3.2	Disk Space Requirements	19
3.3.3	Final Calculation of Disk Requirement	27
Chapter 4.	HACMP Cluster Configurations	29
4.1	Rotating Standby Configuration	30
4.2	Hot Standby Configuration	31
4.2.1	Hot Standby Cluster Description	31
4.3	Mutual Takeover Configuration	32
Chapter 5.	Quick Install Program Preparation	33
5.1	Choosing Oracle Products to Install	33
5.2	Disk Space Sizing for Selected Oracle Products	33
5.3	Copying Files From Floppy Disk	37
5.4	Installation Worksheet	37
5.5	Assign Network Adapters	37
5.6	Update /.rhosts and /etc/hosts	37
5.6.1	Update .rhosts on Both Nodes	38
5.6.2	Update /etc/hosts on Both Nodes	38
5.7	Find Available Group & User ID Numbers	39
5.8	Find the Correct Terminal Emulation	39
5.9	Creation of Logical Volumes	40
Chapter 6.	Running the Quick Install Program	41
6.1	Declaring Node Number	41
6.2	Declaring Failover Type	41
6.3	Choose TTY Port	42
6.4	Select Network Type	42
6.5	Adapter Configuration Screen	42
6.6	Information Menus	47
6.7	Change Default Names	48
6.8	Create DBA Group	48
6.9	Create Oracle User Account	49
6.10	Creating Filesystems, Directories, and Setting Permissions	50
6.11	Setting Oracle User Environment	51
6.12	Creating Volume Groups	51
6.13	Creation of Logical Volumes	53
6.14	Mount CD-ROM	54
6.15	Linking Oracle 7 Files to Staging Area	55
6.16	Running The Preinstallation Script	55
6.17	Oracle 7 Welcome	56
6.18	Enter Installer Activity	56
6.19	Installation Options	57
6.20	Entering the Mount Point	57
6.21	Create Database Objects	58
6.22	Information Prompts	58
6.23	OPS Install	58
6.24	Preinstallation Preparation	59
6.25	Install Products on All Nodes?	59
6.26	Installation Source	60
6.27	Source Staging Area	61
6.28	ORACLE_SID Variable	61

6.29	Specify National Language Support (NLS)	61
6.30	Relink Executables?	62
6.31	Location of Postinstallation File	62
6.32	Online Help and Documentation	63
6.33	Products Menu	63
6.34	Official Hostname	64
6.35	TCP Service Port	65
6.36	Passwords	65
6.37	Group Names	65
6.38	More Than One Instance?	65
6.39	Choose Storage Type	66
6.40	Select the Character Set	66
6.41	Enter the Password for the SYSTEM User and the SYS User	66
6.42	Do You Want To Set the Passwords for DBA and Operator?	66
6.43	MTS Configured and the SQL*Net Listener Automatically Started?	66
6.44	Database Domain Name	67
6.45	Name Control Files and Tablespaces	67
6.45.1	Name and Size Tablespaces and Logs:	69
6.46	Optional Product Documentation	78
6.47	Post-Installation Program	78
6.48	Copying and Linking Necessary Files	78
Chapter 7. Post-Installation of the Oracle Server Database		81
7.1	Verify Database Security and Integrity	81
7.2	Create Oracle Server User Logins	81
7.3	Check your Filesystems	81
7.4	Testing Cluster Synchronization	83
7.5	HACMP Cluster Services	83
Chapter 8. Common Errors, Problems and Recommendations		85
8.1	Common Errors and Problems	85
8.1.1	Errors and Problems During the Installation Process	85
8.1.2	SQL*Net Errors and Problems	86
8.2	Recommendations	87
Appendix A. Quick Installation Worksheets		89
A.1	Mutual Takeover Worksheet	89
A.2	Rotating Standby Worksheet	92
A.3	Hot Standby Worksheet	94
Appendix B. Configuration Files for Rotating & Hot Standby Resources .97		97
B.1	Customized Files	97
B.1.1	File: /etc/services	97
B.1.2	File: /home/oracle/.profile	97
B.1.3	File: /etc/oratab	98
B.1.4	File: //node1_scripts/listener.ora	98
B.1.5	File: /node1_scripts/tnsnames.ora	99
Appendix C. Oracle Configuration Files for Mutual Takeover Resources 101		101
C.1	Customized Files	101
C.1.1	File: /etc/services	101
C.1.2	File: /home/oracle/.profile	101
C.1.3	File: /etc/oratab	102
C.1.4	File: /node1_scripts/listener.ora	103

C.1.5 File: //node1_scripts/tnsnames.ora	104
Appendix D. Special Notices	105
Appendix E. Related Publications	107
E.1 International Technical Support Organization Publications	107
E.2 Redbooks on CD-ROMs	107
E.3 Other Publications	107
How To Get ITSO Redbooks	109
How IBM Employees Can Get ITSO Redbooks	109
How Customers Can Get ITSO Redbooks	110
IBM Redbook Order Form	111
List of Abbreviations	113
Index	115
ITSO Redbook Evaluation	117

Figures

1. Relationship between Two Tables	1
2. Logical Structure of a Typical Oracle Database	8
3. Oracle Architecture	12
4. Cluster Configuration for Rotating Resources	30
5. Cluster Configuration for Hot Standby Resources	31
6. Cluster Configuration for mutual takeover Resources	32
7. .rhosts File from node1	38
8. /etc/hosts file from node1 Preinstallation	38
9. Adapter Test on node1_standby	39
10. Logical Volumes for Oracle 7 Storage	40
11. Oracle Processes Running on Node with Database	84
12. Entry Needed in /etc/services for the Listener	97
13. /home/oracle/.profile File	97
14. /etc/oratab File for Rotating or Hot Standby Resources.	98
15. listener.ora File for Rotating or Hot Standby Resources	99
16. tnsnames.ora File for Rotating or Hot Standby Resources	99
17. Entry Needed in /etc/services File for the Listener.	101
18. /home/oracle/.profile File on First Node	101
19. /home/oracle/.profile File on Second Node	102
20. Node 1 oratab Entry for the Database	102
21. Node 2 oratab Entry for the Database	102
22. oratab Entries for Both Databases Running on One Node	103
23. listener.ora File for Mutual Takeover Resources, Node 1	103
24. listener.ora File for Mutual Takeover Resources, Node 2	104
25. tnsnames.ora File for Mutual Takeover Resources	104

Tables

1.	Hardware Requirements for Oracle Product Installation on a Single Server . .	17
2.	Software Requirements for Oracle Product Installation on a Single Server. . .	17
3.	Prerequisites for PL/SQL Release 2.3	18
4.	Prerequisites for SQL*Plus Release 3.3	18
5.	Prerequisites for SQL*Net Release 2.3	18
6.	Prerequisites for Oracle TCP/IP Adapter Release 2.3	18
7.	Disk Sizing Sheet for Oracle Server Products	20
8.	Disk Sizing Sheet for Oracle Networking Products	20
9.	Disk Sizing Sheet for Oracle Tools Products	22
10.	Disk Sizing Sheet for Oracle CASE Products	24
11.	First-Phase Disk Sizing Sheet for Oracle Precompiler Products	25
12.	Second-Phase Disk Sizing Sheet for Oracle Precompiler Products	26
13.	Total Disk Sizing Summary Sheet	27
14.	Disk Sizing Result Sheet for Oracle Server Products	34
15.	Disk Sizing Result Sheet or Oracle Tools Products	35
16.	Disk Sizing Result Sheet for Oracle Networking Products	36
17.	Total Disk Sizing Result Sheet	37
18.	Common Errors During the Installation Stage	85
19.	Node1 Adapters for Mutual Takeover Configuration	89
20.	Node2 Adapters for Mutual Takeover Configuration	89
21.	Database Logical Volumes (Node 1)	90
22.	Database Logical Volumes (Node 2)	90
23.	Node 1 Adapters for Rotating Configuration	92
24.	Node 2 Adapters for Rotating Configuration	92
25.	Database Logical Volumes for Both Nodes in Rotating Configuration	93
26.	Node1 Adapters for Hot Standby Configuration	94
27.	Node2 Adapters for Hot Standby Configuration	94
28.	Database Logical Volumes for Both Nodes in Hot Standby Configuration. . . .	95

Preface

This redbook describes the implementation of the Oracle database products in a highly available AIX cluster using HACMP. Its purpose is to provide documentation for a Quick Install Program that automates much of the procedures for configuring an HACMP cluster by creating necessary shared disk areas and installing the Oracle database products.

The Quick Install Program supports three types of HACMP cluster configurations with Oracle: Mutual Takeover, Rotating Standby and Hot Standby. It assumes the user has first set up the cluster hardware correctly, including connecting shared disks, connecting network adapters and RS232 cables, according to instructions in the HACMP for AIX Installation Guide. It also assumes the user has installed AIX 4.2 or higher and HACMP 4.2 or higher. From there, the Quick Install Program configures network adapters, creates shared volume groups and logical volumes, configures the HACMP cluster, and installs the Oracle products.

The first three chapters of the redbook provide general information about relational databases and Oracle, and the chapters beyond that directly deal with planning for, preparing for, and running the Quick Install Program. Some knowledge of AIX, the RISC System/6000, and HACMP for AIX is assumed.

The Team That Wrote This Redbook

This redbook was produced by a team of consultants working at the International Technical Support Organization, Austin Center.

Scott Brudner

John Neidhart

Neidhart Consulting Services of Austin, Texas.

Neidhart Consulting is a full-service provider for computer information systems. Consulting services for programming, networking, systems/process analysis and design, as well as hardware installation, and repairs are available. (512) 837-1491

Thanks to the following people for their invaluable contributions to this project:

Dave Thiessen

International Technical Support Organization, Austin Center

Marcus Brewer, Editor

International Technical Support Organization, Austin Center

Roger Feigelson, Product Line Marketing Manager

IBM Products Division, Oracle Corporation

Redwood Shores, CA

Comments Welcome

Your comments are important to us!

We want our redbooks to be as helpful as possible. Please send us your comments about this or other redbooks in one of the following ways:

- Fax the evaluation form found in “ITSO Redbook Evaluation” on page 117 to the fax number shown on the form.
- Use the electronic evaluation form found on the Redbooks Home Pages at the following URLs:

For Internet users: <http://www.redbooks.ibm.com>

For IBM Intranet users: <http://w3.itso.ibm.com/redbooks>

- Send us a note at the following address:

redbook@vnet.ibm.com

Chapter 1. Relational Database Concepts

A database is a means of organizing a collection of data. The basic tasks required of a database are:

- Entering data
- Storing data
- Manipulating data

This chapter provides an overview of the features that are common to all of the databases referred to in this redbook. The following topics are discussed:

- Relational databases
- Database structure
- Structured Query Language
- Client/Server architecture
- File systems and raw logical volume

1.1 Relational Databases

All of the database management systems discussed in this redbook are relational. In this type of database, the information is stored in tables. Each table has a name and is logically represented as being made up of rows and columns. Tables can be related to each other if they contain columns with common types of information.

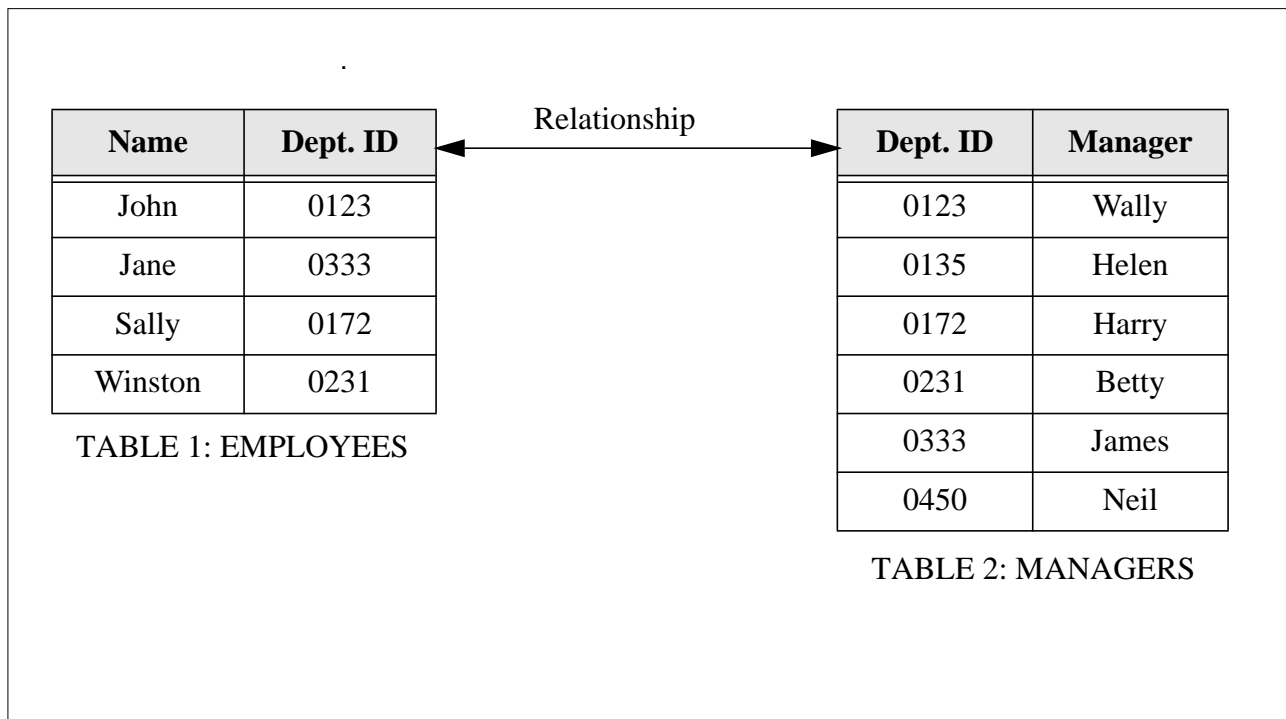


Figure 1. Relationship between Two Tables

Although the information stored in each table may be independent, there is obviously a link between the two tables shown. The department ID in the Employees table is “related” to the department ID in the Managers table. In order to find the name of the manager of a particular employee, you look up his/her department ID (in the Employees table) and then find the manager of this department (in the Managers table).

Relational databases are data driven; that is, the tables and relationships between them are defined once, and then only the data changes over time. Data changes have no impact on the applications using them.

1.2 Database Structure

The structure of a relational database can be divided into three parts:

1. Physical
2. Logical
3. External

The physical part is made up of the files, the directories, and the physical storage elements. The logical part is made up of the objects that are referenced by the database. These include tables, tablespaces, and other elements that make up the relational database model. The boundaries of logical structures are independent of physical space allocation. The external part represents the data as seen by the users. This includes views of the data, the clustering of the information, and accessibility to the various tables.

This three-tiered approach allows for an independence between the data and the means of storage of the data, and it allows for different types of access to the data depending upon the needs of the user.

1.3 Structured Query Language (SQL)

SQL (pronounced sequel or S.Q.L.) is an “English-like,” non-procedural language that is used for most database actions. It was developed by IBM research and has since been refined by ANSI.

1.3.1 SQL Statements

The SQL language consists of about 30 statements. Each SQL statement begins with a *verb* (which describes what the statement does) followed by one or more *clauses*.

The main types of SQL statements are:

- Queries
- Data Manipulation
- Data Definition
- Access Control
- Transaction Control

Queries are statements that retrieve (but do not change) data and begin with the reserved word `SELECT`.

Data Manipulation statements are used to change data by using instructions such as:

- `INSERT`—to add new rows of data
- `DELETE`—to delete rows of data
- `UPDATE`—to change column values in existing rows

Data Definition statements are used to create, alter and drop database objects. They differ from the previous two types of statements because they require a write access to the data dictionary. Typical reserved words are `CREATE TABLE`, `DROP TABLE`, and `ALTER TABLE`.

Access Control statements are used for two types of access:

- Access to the database system
- Access to the database data

Access control statements like `GRANT` and `REVOKE` can control user access privileges.

Transaction Control statements are of two types:

- `COMMIT`—to end the current transaction
- `ROLLBACK`—to abort the current transaction

1.3.2 Context Areas and Cursors

Each SQL statement is associated with a *cursor* or *context area*. This is a memory buffer created on the server that contains the current status of one SQL statement.

A cursor can be thought of as a file handle (or name of the context area) that is opened to gain access to the SQL statement results. The statement is processed or parsed within the cursor, and as long as the cursor remains open, the statement can be re-executed without being reparsed.

Just as a file handle keeps track of the program's current position within a file, the cursor keeps track of the state or phase of the statement. Cursors are only capable of forward sequential processing. When the statement has been fully parsed and no longer needs to be executed, the cursor is closed either explicitly or by the client terminating the connection to the server.

1.3.3 Statement Parsing

An SQL statement is parsed (processed), and a representation of the statements is loaded into the cursor. Parsing consists of:

- Translating the SQL statement
- Loading it into a cursor
- Verifying user access privileges (access control)
- Determining access paths to be used when executing the statement
- Determining resource requirements
- Reserving the resources required

1.4 Client/Server Architecture

In a database using client/server architecture, the database application is divided into two parts:

1. A front-end or *client* portion
2. A back-end or *server* portion

The client part executes the database front-ends, such as interactive query tools or report writers or the application that accesses the database. It often interacts with the user through a graphical user interface. The clients make SQL requests to the server.

The server part executes the database-management-system software and handles all of the functions required for data access and storage. The server processes the request sent to it by the client, calculates the result, and sends the data back to the client.

Although both the client and server parts can execute on the same machine, it is more efficient when client and server are separated in a networked environment.

1.4.1 Client/Server Connection

The client/server connection can be made in various ways. A connection here means the logical association between two applications, such as a client application and the database server. The connection must be established before data transfer and must be maintained for the duration of the transfer.

Most databases use two types of connections:

1. Network connections
2. Shared memory connections

1.4.2 Network Connections

Network connections themselves may be of two types.

1.4.2.1 Remote Network Connections

Remote network connections are used generally when the client and server are on different computers. In this case, the connection is made across a physical network linking the two machines.

1.4.2.2 Local Loopback Network Connections

When the client and server are on the same machine, the setup is called a local loopback connection. The connection is established using the networking facilities of the machine and behaves as though the client and server are on two separate machines.

1.4.3 Shared-Memory Connections

Shared-memory connections can only be established when the client and the server are on the same machine. The client and server share the system memory for very fast data transfer. This type of connection can be more prone to data integrity problems if the shared-memory communication between client and

server is badly managed. There may also be problems arising from memory being overwritten by nondatabase applications.

1.5 Filesystems and Raw Logical Volumes

The files of a database provide the actual physical storage for database information. All of the databases used allow data to be stored either as files within a filesystem or directly on the disk through the raw device interface in logical volumes.

1.5.1 Filesystems

Using filesystem files is easier than using raw devices. In this case, the database simply creates a file in the filesystem. However, when using the filesystem, the database must contend with the disk caching that is inherent in block devices. Block devices do buffered I/O, where data is collected in a buffer until an entire block can be transferred at one time. Generally, writes need to be done synchronously in order to ensure coherency.

On reads, there are overheads due to the data being read into the disk cache before becoming exploitable by the database. There is also the problem that the files in the filesystem may not be placed contiguously on the disk, and additional overheads will be incurred due to disk seek times.

1.5.2 Raw Logical Volumes

A raw logical volume is an area of contiguous physical and logical disk space that is under the direct control of an application rather than under control of the operating system and filesystem.

The applications use character (raw) input and output, carrying out a data transfer with every read or write rather than the block input and output of filesystems. As data is written directly to the disk, bypassing all disk caching and filesystem overheads, performance is generally improved.

Backup and recovery of raw devices is more complicated because file copies do not work. Backups are carried out by using either the UNIX `dd` command or a third-party backup utility.

Chapter 2. Oracle Concepts

The Oracle database follows the three-tiered database structure outlined in Chapter 1, “Relational Database Concepts” on page 1. An Oracle database is generally divided into tablespaces and schema objects such as tables, views and indexes. A description of the three levels as well as the relationship between them is stored in the *data dictionary*.

This chapter describes the more important Oracle files that will be installed on your system and also gives an outline of the architecture of an Oracle instance.

2.1 Oracle Files

An installation of Oracle on your system consists of the following sets of files:

- Data Files
- Redo Log Files (online and archive)
- Control Files
- Initialization Parameter File (INIT.ORA)

2.1.1 Data Files

An Oracle database contains one or more data files. The data files make up the physical repository for all information stored in the database. Two main groups of information are stored in the data files:

- User-created objects (tables, indexes, clusters)
- The data dictionary that is found in the SYSTEM data file created during the Oracle installation process

Figure 2 shows the logical structure of a sample Oracle database.

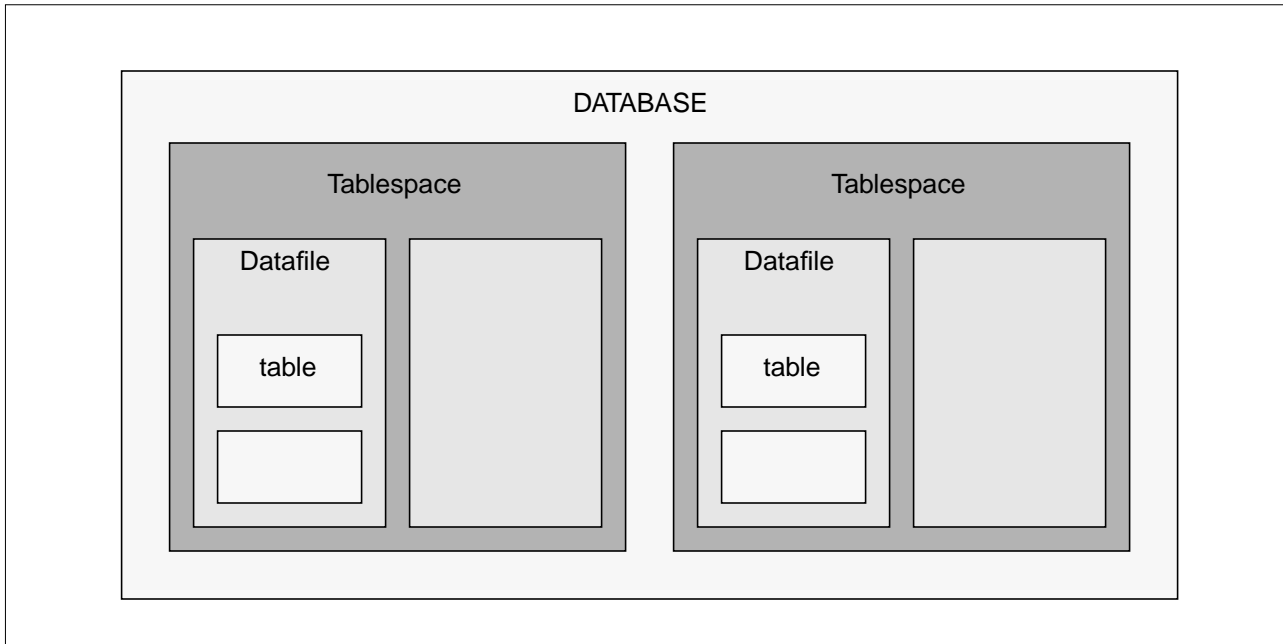


Figure 2. Logical Structure of a Typical Oracle Database

Database files have the following characteristics:

- A data file is only associated with one database.
- The size of a data file cannot be changed after its creation.
- One or more data files can be grouped together to form a tablespace.

2.1.1.1 Tablespaces

A database is made up of one or more tablespaces. Tablespaces are useful in helping you to organize your data for ease of management, security and performance.

System Tablespace is used to store the data dictionary. This holds information such as names of tablespaces and what datafiles are contained in each one.

User Tablespace(s) are used to hold personal data.

2.1.1.2 Segments

A segment is a generic name given to any object that occupies storage in the database. A segment is made up of a group of extents (contiguous blocks) that are in the same tablespace (but not necessarily in the same data file).

Extents are used to minimize the amount of wasted storage and can grow or shrink as required. Extents themselves are made up of data blocks, which are the smallest pieces that make up an Oracle database and are physically related to the disk partition size. An Oracle database has the following types of segments:

Rollback Segment: Whenever data is altered, this change must either be committed or rolled back. The rollback segment holds the previous version of the data being modified; this allows for recovery from aborted or incomplete transactions.

Temporary Segment: This is used by Oracle when an SQL statement needs a temporary work area. It is destroyed upon completion of the statement.

Data and Index Segments: Data segments store user data within the database. Index segments store indexes used by Oracle to look up data quickly. An index scan of the database is much quicker than a full table scan, where Oracle would look at every row in the database.

2.1.2 Redo Log Files

The redo log files are journal files that record all changes made to the database. These files are in memory (for performance reasons, disk I/O is roughly a thousand times slower than actions in memory), and Oracle can then write the changes to the data files (on disk) at its leisure. Every Oracle database has a set of two or more log files.

These files are used in case a failure (machine, disk, and so forth) prevents modified data from being permanently written to the data files. During recovery, Oracle will apply the data in these files to bring the database to a consistent state without losing any committed transactions.

2.1.2.1 Online Redo Logs

Since these files are open or online during normal operation, they are referred to as the online redo log files. The online redo logs work in a circular fashion. As transactions take place, they are recorded in the first redo log. When this is full, a log switch occurs. Now, transactions are recorded in the second redo log. When this is full, another log switch occurs. All new transactions are once again recorded in the first log, overwriting the previous contents.

2.1.2.2 Archived Redo Logs

Oracle offers the possibility of running in either ARCHIVELOG mode or NOARCHIVELOG mode.

In ARCHIVELOG mode, the contents of the redo logs are copied to an archive area (on disk) before they are overwritten. These archive files are known as the archived redo log files. Because they are not open during normal operation of the database and are only required during data recovery, they are also known as the offline redo logs.

In NOARCHIVELOG mode, old redo logs are not kept, and the redo logs are simply overwritten.

2.1.3 Control File

Every Oracle database has a control file. It is highly recommended to have more than one copy to guard against data loss. This file records the physical structure of the database and contains the following types of information:

- Name of the database
- Creation date and time
- Location of the database
- Status and state of all the data files
- Location of the redo logs

Whenever there is a change to the structure of the database, the control file is updated.

2.1.4 The init.ora File

The init.ora file is a parameter file (text) containing Oracle system parameters. It is provided by Oracle and should be customized for your site. This file is read during database startup to determine the size of the system global area (see Section 2.2.1, “System Global Area or Shared Global Area (SGA)” on page 10) and to locate the control files. The actual name of the init.ora file has the Oracle instance identifier appended to it.

2.2 Oracle Memory Structure

Oracle uses system memory to run user processes, cache data, and indexes, and to store shared program code. Oracle has two types of memory structures:

3. The system global area
4. The program global area

2.2.1 System Global Area or Shared Global Area (SGA)

The SGA is an area of shared memory used by Oracle to store data and control information for one Oracle instance. All of the information contained within the SGA is shared by all of the users connected to the instance. The SGA is allocated when the instance starts up and is automatically deallocated when the instance shuts down. The SGA is made up of the following key components:

2.2.1.1 Data Buffer Cache

The data buffer cache stores the most recently used data blocks— that is, those blocks which have been modified but not yet written to disk (dirty blocks) as well as those which have been written to disk (clean blocks).

Before a user process can access a piece of data, it has to be in the data buffer cache. A least recently used algorithm is used to free up space when new data is requested by a user.

2.2.1.2 Redo Log Buffer

Before any transactions can be recorded into the redo logs, they must first reside in the redo log buffer. They are then written to the redo logs by a database background process (log writer).

2.2.1.3 Shared Pool

The shared pool is a cache containing all the parsed SQL statements that are ready to run. This is useful for reducing overheads (memory, processing time, execution planning time) when multiple applications issue the same SQL statement.

2.2.2 Program Global Area (PGA)

A PGA is allocated when a user connects to the database. The PGA is not shared and contains data and control information for single Oracle server processes or for Oracle background processes. The PGA is made up of the following components:

- Stack space: memory holding session variables
- Session information: stored here for dedicated servers and in the SGA for multithreaded servers
- Private SQL area: information on binding variables and run-time buffers

2.3 Oracle Processes

An Oracle instance consists of:

- User processes
- Oracle server processes (shadow processes)
- Oracle background processes
- Shared memory used by these processes

The database can be organized further to run with a dedicated server or with a multithreaded server. In the dedicated server configuration, each user process has a dedicated server process. In the multithreaded server configuration, multiple user processes are serviced by a few shared server processes. Figure 3 on page 12 provides a graphical depiction of the Oracle architecture.

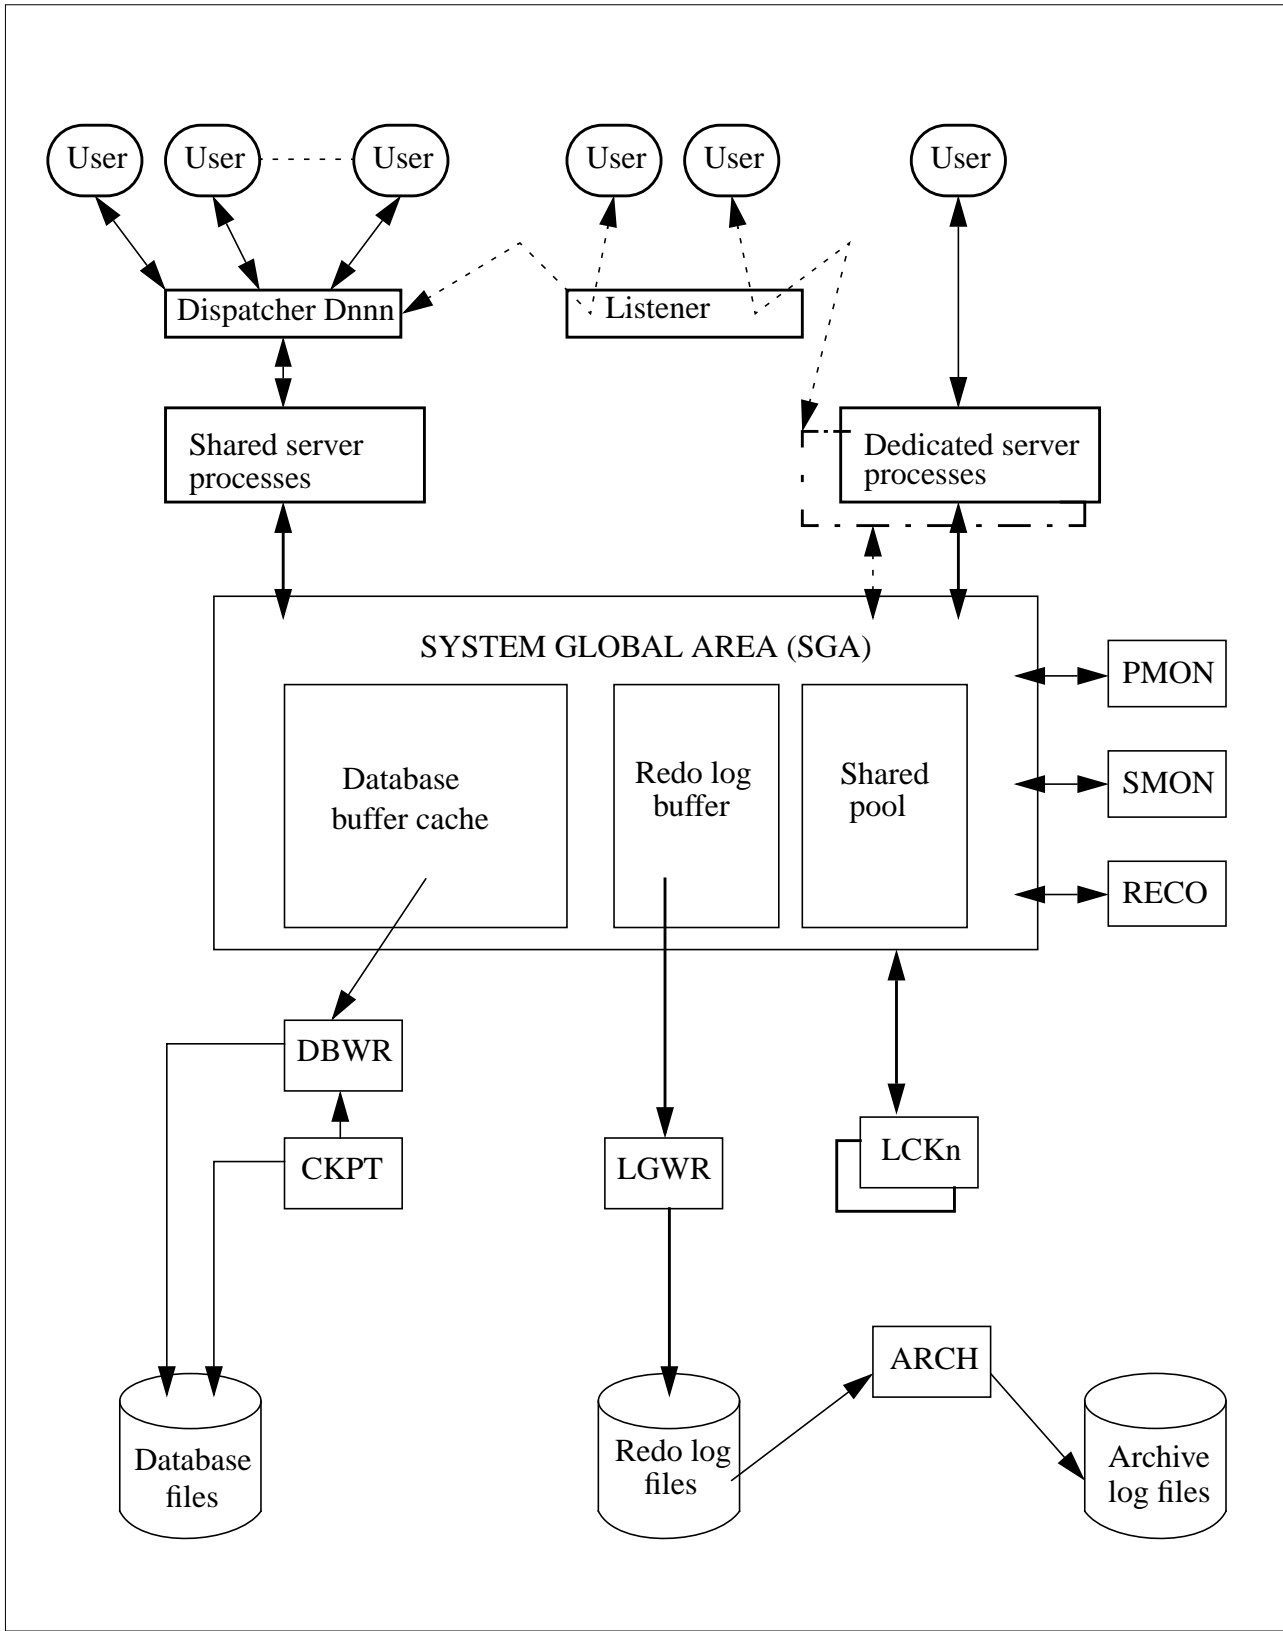


Figure 3. Oracle Architecture

There are four Oracle system background processes that must always be up and running for the database to be usable. They are the following:

1. Database Writer (DBWR)
2. Log Writer (LGWR)
3. System Monitor (SMON)
4. Process Monitor (PMON)

2.3.1 Database Writer (DBWR)

The database writer process is responsible for writing modified (dirty) data blocks from the database buffer cache to the database files. The DBWR only does a write when buffers are needed for data to be read in. There may be more than one DBWR process.

When using the Oracle Parallel Server, the DBWR might be forced to do a write, not because buffers are needed but because another user needs to modify the same buffers. This process is known as *pinging* and can be responsible for degrading database performance.

2.3.2 Log Writer (LGWR)

The log writer process writes redo log entries from the redo log buffer to the redo log files on disk.

2.3.3 System Monitor (SMON)

The system monitor process performs instance recovery at database startup time or, in the case of Oracle parallel server, when another instance belonging to the database has crashed or terminated abnormally.

SMON also releases temporary segments that are no longer needed, compacts the free space fragments in the database files, and detects deadlock situations.

2.3.4 Process Monitor (PMON)

The process monitor process keeps track of database processes. If a user process fails, PMON cleans up the cache and frees up any resources that the failed process was using.

PMON also monitors the dispatcher and shared server processes and restarts them if necessary.

Along with the four processes described above, there are a number of other optional system processes, such as:

- Checkpoint (CKPT)
- Recover (RECO)
- Archiver (ARCH)
- Parallel Server Lock Processes (LCKn)

2.3.5 Checkpoint (CKPT)

A checkpoint is an event in which all modified data blocks are written by the DBWR process to the data files. This usually occurs at a redo log switch. At a

checkpoint, all of the database file headers and the redo log file headers are updated to record the fact that a checkpoint has occurred.

The CKPT is optional. If it is not present, the LGWR process performs the tasks of the CKPT process. It is recommended that the CKPT process is enabled when there are a large number of data files.

2.3.6 Recover (RECO) Process

The recover process is used when there is a failure in a distributed transaction. A distributed transaction is one where two or more locations of the data have to be kept synchronized. In this environment, there may be multiple databases on multiple interconnected servers and either a node or a network fails.

Any transaction that may have completed in one site but not in another is referred to as *in doubt*. The RECO process attempts to establish communication with the remote servers. When the connection is reestablished, the RECO process automatically resolves all the in doubt transactions.

The RECO process is optional and is only needed in instances that carry out distributed transactions.

2.3.7 Archiver (ARCH) Process

The archiver optional process is used when the data is running in ARCHIVELOG mode and automatic archiving is enabled. ARCH copies the redo entries from the online redo log files to the archive area.

2.3.8 Parallel Server Lock Processes (LCKn)

The lock processes are used only if running Oracle Parallel Server. LCKn uses the distributed lock manager for inter-instance locking to prevent simultaneous changes to the same data from different instances.

Chapter 3. Database Planning, System Requirements, and Sizing

This chapter provides information for planning different types of databases. It also discusses AIX-specific issues, including system hardware and software requirements. Included is disk sizing information for your database.

3.1 Planning Your Database

The following section provides information for planning different types of databases.

3.1.1 First Considerations

When planning database installations, you need to determine your processing priorities and the amount of system resources you have.

3.1.1.1 Priorities

Your priorities will decide the characteristics of your database. Things to consider are:

- Is transaction speed more important than data safety?
- Will the database be used for a large number of short transactions or for a small number of large transactions?
- Is the database to be accessed by a single application or by multiple applications running on multiple nodes?

3.1.1.2 System Resources

The system resources available to you will also affect your database installation, usage and backup strategies. The following considerations are important:

- The number of users and applications running at one time will depend upon the physical memory available and the system running the database.
- You need to know how many disk drives are available, their device names, the speeds of the various disks (are some faster than others), and the number of disk controllers available.
- For backup purposes, you should know the number and type of tape drives available.
- What UNIX kernel parameters need to be changed or added?
- What network resources are available?

3.1.2 Database Application Type

The type of database application you are planning is a primary criterion in determining how you will plan and structure the database. Depending on the purpose for which they will be used, there are four basic types of database applications:

- On Line Transaction Processing (OLTP)
- Decision Support
- Development
- Mixed

3.1.2.1 OLTP Database

An OLTP database, such as a bank automated teller machine or ATM, has a very high volume of transactions (measured in transactions per second). When planning a database for such purposes, you usually need to:

- Split database files so that disk I/O is shared between many disk subsystems
- Split the logical database design
- Implement redo log archiving
- Be available 7 days a week, 24 hours a day
- Use raw disk partitions
- Use high availability hardware (mirrored disks, for example)
- Institute performance tuning
- Set *initsid.ora* parameters larger than the default values to accommodate more users and increase performance

3.1.2.2 Decision Support Database

A decision support database, such as an inventory system, has a relatively low number of database updates (measured in transactions per hour). The users tend to make few queries, and they look at the results of these queries for many minutes at a time. They may prepare reports from these queries. A decision support database:

- Has less need to distribute I/O across multiple disks and controllers
- Has less need to split the logical database design
- Has less need to implement redo log archiving
- Can usually be brought down for maintenance or backups

3.1.2.3 Development Database

A development database is used to develop new applications. Such a database:

- Does not usually need to split I/O
- Can usually be brought down for maintenance or backups
- Has little need for backup more than once a day
- Has little need for performance tuning

3.1.2.4 Mixed Database

A mixed database combines various functions, such as decision support and transaction processing. Most databases fall within this category. Such a database:

- Has an even mixture of queries and updates
- Has some performance tuning needs, but not as much as an OLTP database

3.1.3 Database Size

When considering the size of a database, you must always consider it relative to the system upon which it will run. The basic considerations are:

- A machine with more physical memory can support a larger database
- A machine with more processors can support a larger database

3.1.4 File Locations

Several files that will be prompted for when you install and create your database or object database must be put in the external shared disk if you need a highly available database within an HACMP environment. The files to be put into the external shared disk are as follows:

- Control file
- Redo log
- System tablespace
- Rollback segment
- Tools tablespace
- User tablespace
- Temporary tablespace

3.2 System Requirements

This section lists the hardware and software requirements for installing your Oracle database.

3.2.1 Requirements for Oracle Product Installation on a Single Server

This section outlines the requirements for installing various Oracle products.

- Hardware Requirements

Table 1. Hardware Requirements for Oracle Product Installation on a Single Server

Hardware Item	Requirement
CPU	IBM RISC System/6000.
Memory	Please refer to the disk space and memory requirement chart in the next section.
Tape Device	An 8 mm tape drive is required for the tape distribution of Oracle. The tape drive should be set to 512 block size.
CD-ROM	ISO 9660 format is supported.
Terminal	The following terminal types are supported: ANSI, HFT, LFT, 3151, VT100, WYSE50, and so on.
Controller	SQL*Net TCP/IP requires an adapter that will support TCP/IP.

- Software Requirements

Table 2. Software Requirements for Oracle Product Installation on a Single Server

Software Item	Requirement
Operating System	AIX Version 4.1.4 or later

3.2.2 Requirements for Oracle Optional Products

For the prerequisite requirements of Oracle optional products, please refer to the Oracle manual. The Oracle products are usually grouped into the following categories:

- Tools

- Precompilers
- CASE Products
- SQL*Net V2

Note that when you install an individual product, the dependent products are installed automatically by the Installer.

You can find more information about prerequisites or requirements for Oracle products that are included in one of the four categories mentioned above in *Chapter 2, "Requirements,"* in the *Oracle7 for AIX-Based Systems Installation and Configuration Guide*.

Below, we only show you some examples of the tables that we installed in this book, such as the Oracle Tools, PL/SQL Release 2.3 and SQL*Plus Release 3.3, and the SQL*Net with TCP/IP Protocol Adapter requirements or prerequisites for installation.

PL/SQL Release 2.3

Table 3. Prerequisites for PL/SQL Release 2.3

Software Requirement	Version	State During Installation
Oracle Server	7.3	Installed
Hardware Requirement		
None		

SQL*Plus Release 3.3

Table 4. Prerequisites for SQL*Plus Release 3.3

Software Requirement	Version	State During Installation
Oracle Server	7.3	Running
Hardware Requirement		
None		

SQL*Net Release 2.3

Table 5. Prerequisites for SQL*Net Release 2.3

Software Requirement	Version	State During Installation
Oracle Server	7.3	Installed
Oracle TCP/IP Adapter	2.3	Installed
TCP/IP for AIX 4.1.4	Included in O/S	Installed
Hardware Requirement		
Token-Ring Adapter or Ethernet Adapter supported for IBM RS/6000		

Oracle TCP/IP Adapter Release 2.3

Table 6. Prerequisites for Oracle TCP/IP Adapter Release 2.3

Software Requirement	Version	State During Installation
Oracle Server	7.3	Installed
SQL*Net	2.3	Installed
Hardware Requirement		
Token-Ring Adapter or Ethernet Adapter supported for IBM RS/600		

3.3 Disk Space and Memory Sizing for Oracle

This section is designed as a planning tool. It provides the information required to calculate the distribution space, database space, and memory requirements for each Oracle product available for installation on your system. Also we will discuss the memory requirement based on users. Follow the step-by-step procedure to calculate the disk storage and memory requirements for your specific installation. Note that these instructions provide only generalized estimates of your disk and memory needs, not a precise calculation.

3.3.1 Required Calculations

The calculations result in two values for total disk storage and two values for system memory. These values are:

Distribution Space	Space required in megabytes to store the products on disk.
Database Space	Megabytes of space the product requires in the database space. Assess the number of Oracle blocks required by dividing this amount by the default Oracle block size for your machine.
First User Memory	An estimate of virtual memory space, in kilobytes, used when the first user runs the product.
Memory/Additional User	An estimate of the virtual memory space, in kilobytes, used when each additional user runs the product.

3.3.2 Disk Space Requirements

In this section, we will explain the disk space sizing sheet for use as a tool for sizing and estimating the space required for an Oracle installation. There are six kinds of sheets, one for each category of Oracle product.

3.3.2.1 Oracle Server and Networking Products

To arrive at a basic estimate of your space needs for Oracle and the networking products, perform these calculations:

1. Select the product you want to install by placing a check mark in the first column of each product.
Note: For server installation, several products, such as Oracle Server and its components, Oracle Common Libraries and Utilities, Oracle Toolkit, and SQL*Net, are mandatory.
2. Calculate the Distribution (**Dist**) Space by subtotaling the values for the selected products in “A” at the bottom of the column.
3. Calculate the Database Space (**DB Sp**) by subtotaling the values for the selected products in “B” at the bottom of the column.
4. Transfer the First User Memory Space (**#1 User**) value for the Oracle Server to the Subtotal “C” at the bottom of the column.
5. Calculate the **Additional Users** memory space under the **Memory Space Requirements** heading.

Enter the number of additional users you estimate for the Oracle Server in the **Users** column.

Multiply **Users** by **KB per** and place the result in **Total**. For example:

$$9 \text{ users} \times 186 \text{ KB per} = 1674 \text{ KB}$$

Transfer this value into Subtotal **D** at the bottom of the column.

6. In the table's **Summary** section, enter **A**, **B**, **C**, and **D** subtotals.

Add the #1 User Memory (**C**) and Additional User Memory (**D**), and enter the result in the Virtual Memory Total (**E**).

Below are the tables showing the disk space sizing sheets for the Oracle server products and Oracle networking products.

Table 7. Disk Sizing Sheet for Oracle Server Products

Space Requirements for Oracle Server Products							
Disk Storage Requirements				Memory Space Requirements			
v		Dist.	DB Sp	#1 User	Additional Users		
Use	Product	(MB)	(MB)	(KB)	Users	KB per	Total
v	Oracle7 Server	32.37	12.5	6978		x 521	
	SQL*DBA			2589		x 283	
	SQL*Loader			1873		x 102	
	Export			1741		x 121	
	Import			1675		x 105	
	Server Manager (Line Mode only)	12.57	N/A	1908		x 120	
	Server Manager (Motif bitmapped & Line Mode)	12.57	N/A	5058		x 486	
	+ Distributed Opt	0.10	N/A				
	+ Parallel Server Opt.	0.08	N/A				
v	+ Parallel Query Opt	0.08	N/A				
v	Toolkit II	68.39	N/A				
v	Oracle Common Lib & Utilities	31.48	N/A				
	Migration Utility	3.00	N/A	1816		x 247	
	Oracle Parallel Backup/Restore Util.	30.00	10.0	1024			
		Σ A	Σ B	Σ C			Σ D
	Subtotals =						
Summary							
A	Tot. Distribution Space (A)			(MB)			
B	Tot. Database Space (B)			(MB)			
C	#1 User Memory (C)			(KB)			
D	Additional User Memory (D)			(KB)			
E	Virtual Memory Total (C+D)			(KB)			

Table 8. Disk Sizing Sheet for Oracle Networking Products

Space Requirements for Oracle Networking Products							
Disk Storage Requirements				Memory Space Requirements			
v		Dist.	DB Sp	#1 User	Additional Users		
Use	Product	(MB)	(MB)	(KB)	Users	KB per	Total
SQL*Net V2:							

Space Requirements for Oracle Networking Products						
v	SQL*Net V2	11.08	0.2			
	TNS Listener			1710		x 94
	DECnet	0.13	N/A			
	LU6.2	0.18	N/A	7		x 2
	OSI	0.16	N/A			
	SPX/IPX	1.77	N/A	695		x 152
	Names Server	6.59	N/A	3221		x 340
v	TCP/IP	0.15	N/A			
	Interchange	7.30	N/A			
	Interchange Control Utility			1302		x 70
	Interchange Listener			2127		x 172
	Navigator			926		x 56
	Data Pump	N/A	N/A	1001		x 56
	Configuration Tool	N/A	0.2			
		Σ A	Σ B	Σ C		Σ D
	Subtotals =					
Summary						
A	Tot. Distribution Space (A)			(MB)		
B	Tot. Database Space (B)			(MB)		
C	#1 User Memory (C)			(KB)		
D	Additional User Memory (D)			(KB)		
E	Virtual Memory Total (C+D)			(KB)		

3.3.2.2 Oracle Tools

The calculation procedures for sizing the space requirements of Oracle Tools products is the same as above. You must keep in mind that First User Memory (**#1 User**) is extremely difficult to predict because it depends heavily on the application. Furthermore, the number of cursors opened by applications can significantly affect the size of associated shadow processes. The numbers given in this sheet are to be understood as the bare minimum.

1. Select the product you want to install by placing a check mark in the first column of each product.
2. Calculate the Distribution (**Dist**) Space by subtotalling the values for the selected products in **A** at the bottom of the column.
3. Calculate the Database Space (**DB Sp**) by subtotalling the values for the selected products in **B** at the bottom of the column.
4. Transfer the First User Memory Space (**#1 User**) value for the Oracle Tools to the Subtotal **C** at the bottom of the column.
5. Calculate the Additional Users Memory Space.
6. Enter the number of additional users you estimate for the Oracle Tools in the **Users** column.
7. Multiply **Users** by **KB per** and place the result in **Total**. For example:
8. 9 additional users of SQL*Plus x 778 KB per = 7002 KB
9. Transfer this value into Subtotal **D** at the bottom of the column.
10. In the table's **Summary** section, enter **A**, **B**, **C**, and **D** subtotals.
11. Add the #1 User Memory (**C**) and Additional Users Memory (**D**), and enter the result in the Virtual Memory Total (**E**).

Below are the tables showing the disk space sizing sheets for the Oracle Tools products.

Table 9. Disk Sizing Sheet for Oracle Tools Products

Space Requirements for Oracle Tools							
Disk Storage Requirements				Memory Space Requirements			
v		Dist.	DB Sp	#1 User	Additional Users		
Use	Product	(MB)	(MB)	(KB)	Users	KB per	Total
	Easy*SQL						
	Oracle Data Query	17.81	0.4	3424		x 756	
	Oracle*Mail	48.50	N/A	2628		x 362	
	Oracle*Terminal	5.77	N/A	667		x 194	
	PL/SQL 1.0		N/A				
	PL/SQL 2.0	0.28	N/A				
	SQL*Calc	2.99	N/A				
	Oracle Toolkit I	1.29	N/A				
	Oracle XA Library	0.15	N/A				
	SQL*Forms V3.0	24.99	1.2				
	Design (Char Mode)			4614		x 1034	
	Runtime (Char Mode)			3511		x 710	
	SQL*Menu V5.0	21.84	0.8				
	Design (Char Mode)			4388		x 998	
	Runtime (Char Mode)			3565		x 760	
	SQL*Plus	5.65	0.6	3472		x 778	
	SQL*QMX	2.97					
	SQL*Report	2.37	N/A				
	rpt			1748		x 230	
	rpf			46		x 9	
	SQL*ReportWriter	2.37	N/A				
	Design			2642		x 230	
	Runtime			2424		x 302	
		Σ A	Σ B	Σ C			Σ D
	Subtotals =						
Summary							
A	Tot. Distribution Space (A)			(Mb)			
B	Tot. Database Space (B)			(Mb)			
C	#1 User Memory (C)			(Kb)			
D	Additional User Memory (D)			(Kb)			
E	Virtual Memory Total (C+D)			(Kb)			

3.3.2.3 Oracle CASE Products

Although in this redbook we did not install any CASE products, it is necessary for you to know about disk size spacing for these products.

The categories of space requirements given for CASE*Designer are as follows:

casefront The executable run by all users at all times. It is the startup window from which users invoke the different CASE*Designer diagrammers, depending on the stage of development.

caseerd The diagrammer invoked by all users at the entity relationship definition stage of development.

- casefhd** The diagrammer invoked by all users at the function hierarchy definition stage of development.
- casedfd** The diagrammer invoked by all users at the dataflow definition stage of development.
- casemd** The diagrammer invoked by all users who are checking completeness and validity links between pairs of element types.

The calculation procedures for sizing space requirements of Oracle CASE products are the same as above. The numbers given in this sheet are to be understood as the bare minimum.

1. Select the product you want to install by placing a check mark in the first column of each product.
2. Calculate the Distribution (**Dist**) Space by subtotaling the values for the selected products in **A** at the bottom of the column.
3. Calculate the Database space (**DB Sp**) by subtotaling the values for the selected products in **B** at the bottom of the column.
4. Transfer the First User Memory Space (**#1 User**) value for the Oracle Tools to the Subtotal **C** at the bottom of the column.
5. Calculate the Additional Users Memory Space.

Enter the number of additional users you estimate for the Oracle Tools in the **Users** column.

Multiply **Users** by **KB per** and place the result in **Total**. For example:

3 addtl. users of CASE*Designer casemd x 930 KB per = 2790 KB

Transfer this value into Subtotal **D** at the bottom of the column.

6. In the table's **Summary** section, enter **A**, **B**, **C**, and **D** subtotals.
Add the #1 User Memory (**C**) and Additional Users Memory (**D**), and enter the result in the Virtual Memory Total (**E**).

Below are the tables showing the disk space sizing sheets for the Oracle CASE products.

Table 10. Disk Sizing Sheet for Oracle CASE Products

Space Requirements for CASE Products							
Disk Storage Requirements				Memory Space Requirements			
v		Dist.	DB Sp	#1 User	Additional Users		
Use	Product	(MB)	(MB)	(KB)	Users	KB per	Total
	CASE*Dictionary	53.18	7.5				
	Runtime			1730		x 237	
	CASE*Generator	33.65	4.5				
	Runtime			1895		x 281	
	CASE*Designer	15.11	3.5				
	casefront Runtime			2107		x 330	
	caseerd Runtime			2575		x 426	
	casefhd Runtime			2504		x 430	
	casedfd Runtime			2689		x 496	
	casemd Runtime			2904		x 930	
		Σ A	Σ B	Σ C			Σ D
	Subtotals =						
Summary							
A	Tot. Distribution Space (A)			(MB)			
B	Tot. Database Space (B)			(MB)			
C	#1 User Memory (C)			(KB)			
D	Additional User Memory (D)			(KB)			
E	Virtual Memory Total (C+D)			(KB)			

3.3.2.4 Oracle Precompiler Products

Although in this redbook we did not install any precompiler products, it is necessary for you to know about disk size spacing for these products.

Calculating memory usage of these products also requires differentiation of application development and production cycles.

During application development, a few programmers will be using the prox (procob for COBOL, proc for C, proada for Ada, propas for Pascal, and profor for FORTRAN) executable extensively. During the production cycles, many more users may be executing the finished programs, which have been linked with the SQLLIB run-time library. Remember that each precompiler application may spawn an Oracle shadow process; so the same rule for including the per-additional-user memory consumption of the Oracle Kernel applies to precompilers.

Disk space for the precompilers is somewhat redundant. It is calculated assuming only one language is installed on your system. When a second language is installed, the additional free space required is not as great as listed because different languages share common libraries.

The calculation procedures for sizing space requirements of Oracle precompiler products have two phases. First, you have to calculate the development cycle calculations, which are the same as for other products above.

1. Select the product you want to install by placing a check mark in the first column of each product.

1. Calculate the Distribution (**Dist**) Space by subtotalling the values for the selected products in **A** at the bottom of the column.
2. Calculate the Database space (**DB Sp**) by subtotalling the values for the selected products in **B** at the bottom of the column.
3. Transfer the First User Memory Space (**#1 User**) value for the Oracle Tools to the Subtotal **C** at the bottom of the column.
4. Calculate the Additional Users Memory Space.
5. Enter the number of additional users you estimate for the Oracle Tools in the **Users** column.
6. Multiply **Users** by **KB per** and place the result in **Total**. For example:
7. 2 addtl. users of Pro*C V2 x 329 KB per = 658 KB
8. Transfer this value into Subtotal **D** at the bottom of the column.
9. In the table's **Summary** section, enter **A**, **B**, **C**, and **D** subtotals.
10. Add the #1 User Memory (**C**) and Additional Users Memory (**D**), and enter the result in the Virtual Memory Total (**E**).

Below are the tables showing the first-phase disk space sizing sheets for the Oracle precompiler products.

Table 11. First-Phase Disk Sizing Sheet for Oracle Precompiler Products

Space Requirements for CASE Products							
Disk Storage Requirements				Mem3 ory Space Requirements			
v		Dist.	DB Sp	#1 User	Additional Users		
Use	Product	(MB)	(MB)	(KB)	Users	KB per	Total
	Pro*Ada	1.10	N/A	3251		x 307	
	Pro*C v1.6	1.10	N/A	3261		x 396	
	Pro*C v2.2	1.10	N/A	3534		x 329	
	Pro*COBOL	1.10	N/A	3271		x 308	
	Pro*FORTRAN	1.10	N/A	3254		x 306	
	SQL*Module for Ada	8.83	N/A	3286		x 306	
	SQL*Module for C	8.83	N/A	3286		x 306	
		Σ A	Σ B	Σ C			Σ D
	Subtotals =						
Summary							
A	Tot. Distribution Space (A)			(MB)			
B	Tot. Database Space (B)			(MB)			
C	#1 User Memory (C)			(KB)			
D	Additional User Memory (D)			(KB)			
E	Virtual Memory Total (C+D)			(KB)			

After you have finished the first-phase calculation above, you must continue to calculate the second phase, which is the use cycle calculation. To calculate the minimum memory requirement for a finished user program, perform the following steps:

1. Determine the text and data size of each executable by using the `size` command.
2. Add the text and data total to determine the first user memory requirement for the program (F).

Space (A), in the final calculation sheet of disk space sizing before you prepare your installation tasks.

For more information about this documentation library and how to use it, please refer to the *Getting Started with Oracle Online Documentation for AIX* manual from the Oracle Corporation.

3.3.3 Final Calculation of Disk Requirement

Enter the **A**, **B**, and **E** totals from the individual tables into the table below. Total the three columns (including the additional Disk Space requirement) to determine the total distribution space, database space, and virtual memory space you will need.

Table 13. Total Disk Sizing Summary Sheet

Total Space Requirements (MB)			
Transfer Table	Disk Space (A)	Database Space (B)	Virtual Memory (E)
Oracle Server Products			
Precompiler Products: Development Cycle			
Precompiler Products: User Cycle	n/a	n/a	
Oracle Networking Products			
Oracle CASE Products			
Oracle Tools Products			
Oracle Online Documentation		n/a	n/a
Additional installation space for files generated during installation	+ 21 MB (Minimum)	n/a	n/a
Temporary space for relinking	+ 40 MB (minimum)		
TOTALS			

Chapter 4. HACMP Cluster Configurations

This chapter describes the HACMP with Oracle 7 configurations that are available using the Quick Install Program included with this redbook. The levels of software that have been tested with the program are:

- AIX 4.2 and 4.2.1
- HACMP 4.2 and 4.2.1
- Oracle Server 7.3.2

In each case, the cluster consists of two nodes. The cluster configurations available are the following:

- Rotating Standby
- Hot Standby
- Mutual Takeover

In each case, the configuration pictured uses the default settings in the Quick Install Program for network addresses and labels, shared volume groups and so on. These items are modifiable by the user.

4.1 Rotating Standby Configuration

This cluster configuration, consisting of two nodes, is set up in a rotating standby configuration. The cluster was configured as shown in Figure 4.

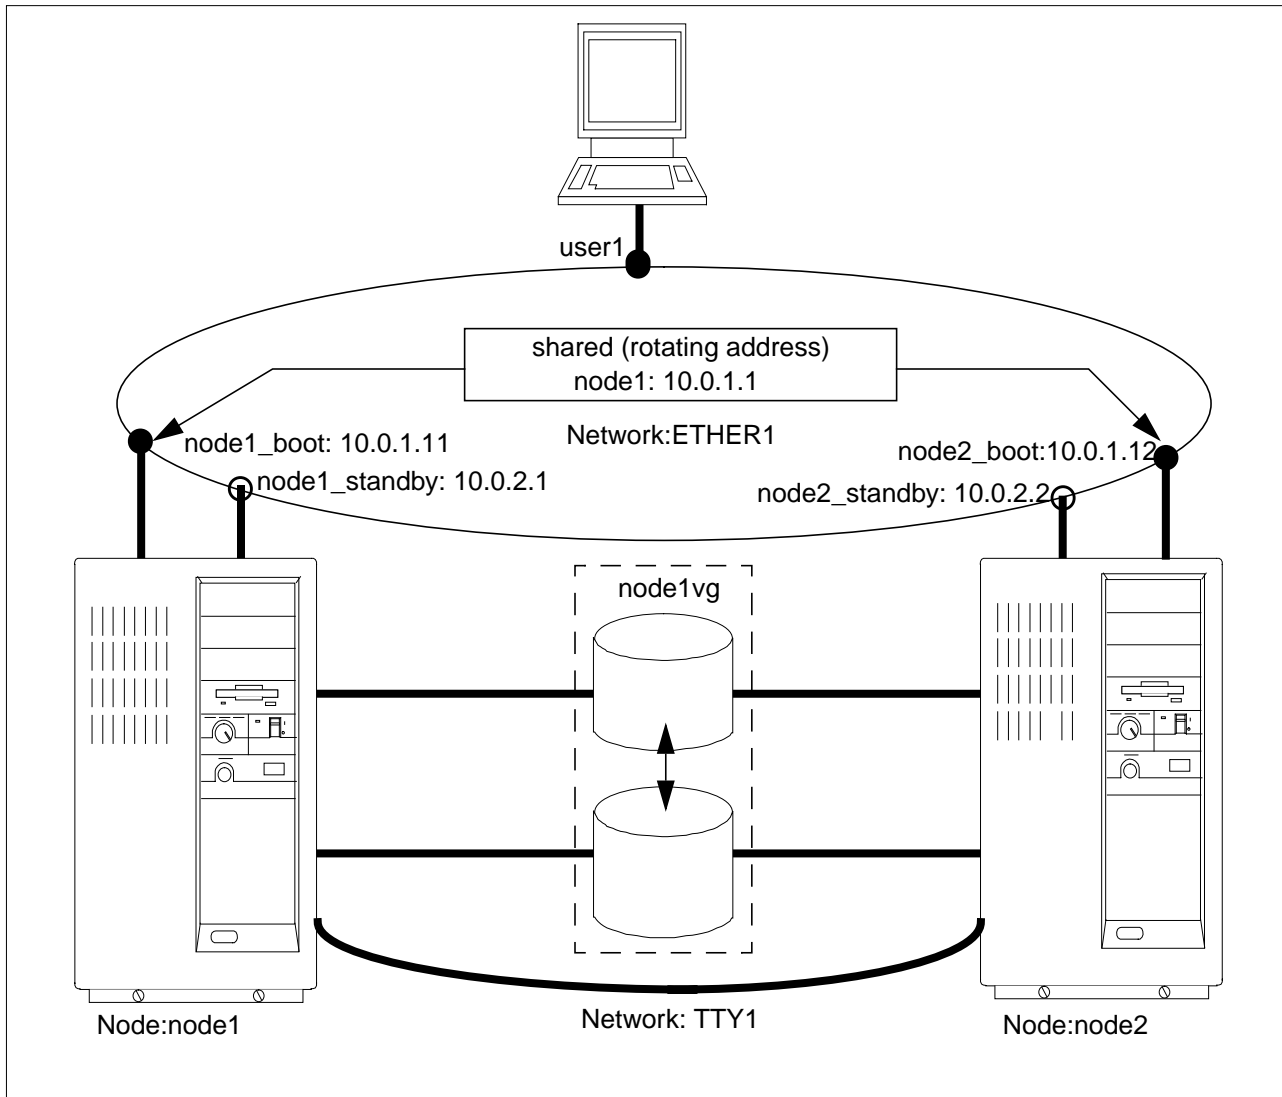


Figure 4. Cluster Configuration for Rotating Resources

The Rotating Standby cluster has a single resource group and application server. A "shared" adapter (we call it shared, although in the actual HACMP configuration panels it is called a service adapter) is configured for both nodes. The resource group containing the Oracle database is acquired and started by the first node to enter the cluster (start HACMP). On failure of the node that is serving the resources, the other node acquires the resources and will not release them unless it fails or leaves the cluster using the "graceful with takeover" option. If a node leaves the cluster (gracefully or because of failure) and then rejoins, it assumes the backup role until the other node leaves the cluster.

4.2 Hot Standby Configuration

This cluster, consisting of two nodes, is set up in a Hot Standby configuration. The cluster is configured as shown in Figure 5.

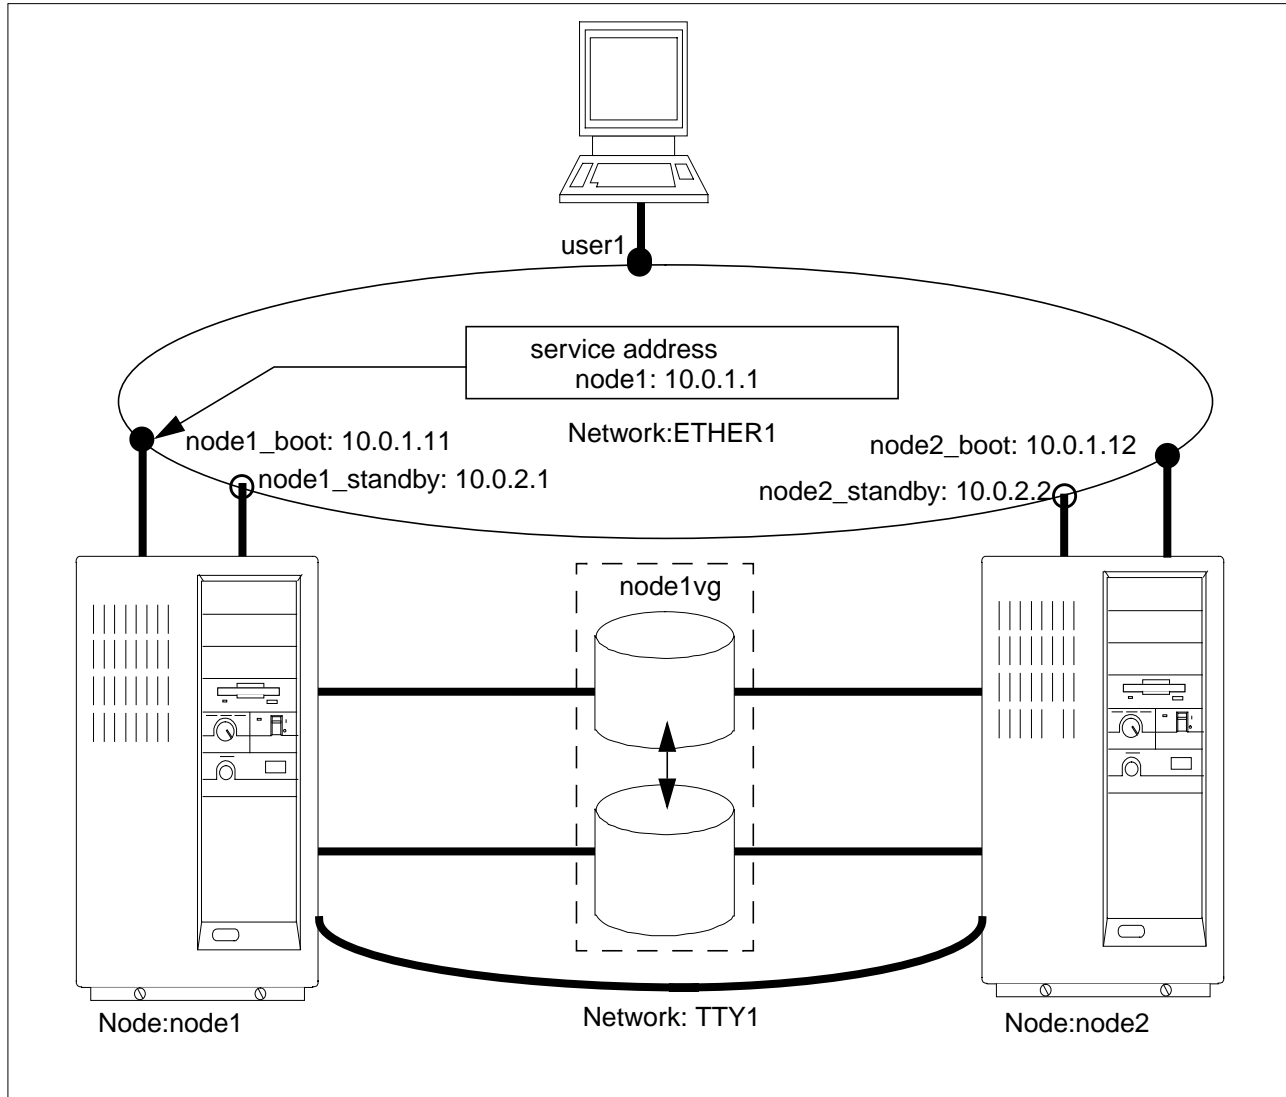


Figure 5. Cluster Configuration for Hot Standby Resources

4.2.1 Hot Standby Cluster Description

The Hot Standby cluster has a single resource group and application server. The node1vg volume group contains a single Oracle 7 database. By default, node 1 is assigned the high priority (server) role, and node2 is assigned the low priority (backup) role. Whenever node1 joins the cluster, it acquires the resources and starts the Oracle database. If node1 fails, the database is taken over by node2. When node1 rejoins the cluster, it reacquires the resources from node2.

4.3 Mutual Takeover Configuration

This cluster, consisting of two nodes, is set up in what is traditionally called a Mutual Takeover configuration. The cluster is configured as shown in Figure 6.

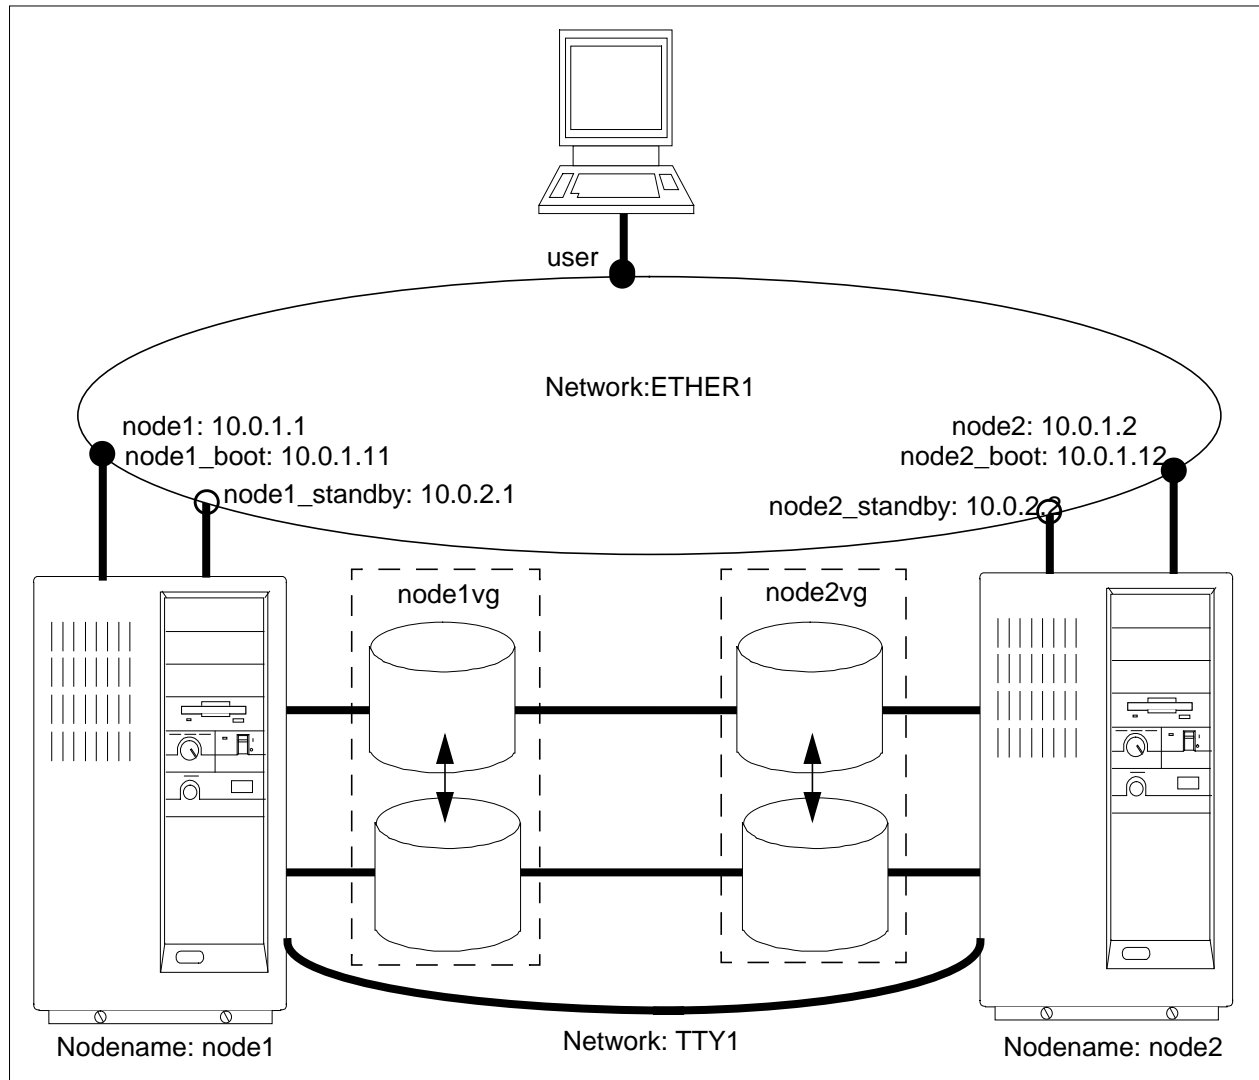


Figure 6. Cluster Configuration for mutual takeover Resources

Mutual Takeover is a form of cascading configuration where both nodes have their own resource groups and application servers. There are two Oracle 7 databases in the cluster, one contained in volume group node1vg and the other in volume group node2vg. Node1 is the high-priority node for the database in node1vg, and node2 is the backup. For the database in node2vg, node2 is the high-priority node, and node1 is the backup. In the event that either node fails, the opposite node acquires its resources. When the failed node reintegrates, its resources are returned to it.

Chapter 5. Quick Install Program Preparation

This chapter will get the cluster and user prepared to run the Quick Install Program.

5.1 Choosing Oracle Products to Install

The Oracle products selected for our test installation on the Rotating, Hot Standby, and Mutual Takeover configurations are listed below.

1. Oracle Server Products:

- Oracle7 Server with the Parallel Query option and its dependent products, such as the Oracle Common Libraries and Utilities and Toolkit II.

2. Oracle Tools Products:

- SQL*Plus Release 3.3

3. Oracle Networking Products:

- SQL*Net V2.3
- TCP/IP Protocol Adapter V2.3

We did not install any CASE products, precompiler products, or Oracle on-line documentation.

5.2 Disk Space Sizing for Selected Oracle Products

After selecting the products listed above (or any you chose), we now must calculate the disk space requirement based on disk space sizing worksheets that we described in Chapter 3, "Database Planning, System Requirements, and Sizing" on page 15.

The install program accounts for all space requirements. In most cases, the sizes used are double the minimum recommended. To make any adjustments, follow the disk space sizing worksheets below:

Table 14. Disk Sizing Result Sheet for Oracle Server Products

Space Requirements for Oracle Server Products							
Disk Storage Requirements				Memory Space Requirements			
v		Dist.	DB Sp	#1 User	Additional Users		
Use	Product	(MB)	(MB)	(KB)	Users	KB per	Total
v	Oracle Server	36.60	12.5	6458	1	x 489	489
	SQL*DBA			2589	1	x 283	283
	SQL*Loader			1778	1	x 204	204
	Export			1654	1	x 223	223
	Import			1590	1	x 207	207
	Server Manager (Line Mode only)	20.19	N/A	1752		x 221	
	Server Manager (Motif bitmapped & Line Mode)	42.88	N/A	3515		x 364	
	+ Distributed Opt	0.10	N/A				
	+ Parallel Server Opt.	0.08	N/A				
v	+ Parallel Query Opt	0.08	N/A				
v	Toolkit II	43.06	N/A				
v	Oracle Common Lib & Utilities	22.10	N/A				
	Migration Utility	1.85	N/A				
	Oracle Parallel Backup/Restore Util.	30.00	10.0	1024			
		Σ A	Σ B	Σ C			Σ D
	Subtotals =	101.84	12.5	14069			1612
Summary							
A	Tot. Distribution Space (A)		101.84				(MB)
B	Tot. Database Space (B)		12.50				(MB)
C	#1 User Memory (C)		14069				(KB)
D	Additional User Memory (D)		1612				(KB)
E	Virtual Memory Total (C+D)		15681				(KB)

Note: Each product selected for installation is indicated by **(v)**. The total disk space required for each product is written in **bold** type.

Table 15. Disk Sizing Result Sheet for Oracle Tools Products

Space Requirements for Oracle Tools							
Disk Storage Requirements				Memory Space Requirements			
v		Dist.	DB Sp	#1 User	Additional Users		
Use	Product	(MB)	(MB)	(KB)	Users	KB per	Total
	Easy*SQL						
	Oracle Data Query	17.81	0.4	3424		x 756	
	Oracle*Mail	48.50	N/A	2628		x 362	
	Oracle*Terminal	5.77	N/A	667		x 194	
	PL/SQL 1.0		N/A				
v	PL/SQL 2.0	0.28	N/A				
	SQL*Calc	2.99	N/A				
v	Oracle Toolkit I	1.29	N/A				
	Oracle XA Library	0.15	N/A				
	SQL*Forms V3.0	24.99	1.2				
	Design (Char Mode)			4614		x 1034	
	Runtime (Char Mode)			3511		x 710	
	SQL*Menu V5.0	21.84	0.8				
	Design (Char Mode)			4388		x 998	
	Runtime (Char Mode)			3565		x 760	
v	SQL*Plus	5.65	0.6	3472	1	x 778	778
	SQL*QMX	2.97					
	SQL*Report	2.37	N/A				
	rpt			1748		x 230	
	rpf			46		x 9	
	SQL*ReportWriter	2.37	N/A				
	Design			2642		x 230	
	Runtime			2424		x 302	
		Σ A	Σ B	Σ C			Σ D
	Subtotals =	7.22	0.6	3472			778
Summary							
A	Tot. Distribution Space (A)		7.22	(MB)			
B	Tot. Database Space (B)		0.60	(MB)			
C	#1 User Memory (C)		3472	(KB)			
D	Additional User Memory (D)		778	(KB)			
E	Virtual Memory Total (C+D)		4250	(KB)			

Note: Each product selected for installation is indicated by **(v)**. The total disk space required for each product is written in **bold** type.

Table 16. Disk Sizing Result Sheet for Oracle Networking Products

Space Requirements for Oracle Networking Products							
Disk Storage Requirements				Memory Space Requirements			
v		Dist.	DB Sp	#1 User	Additional Users		
Use	Product	(MB)	(MB)	(KB)	Users	KB per	Total
SQL*Net V1:							
	SQL*Net Async	0.27	N/A				
	SQL*Net DECnet	0.12	N/A				
	SQL*Net LU6.2	0.13	N/A				
	SQL*Net SPX/IPX	2.55	N/A				
	SQL*Net Named Pipes		N/A				
	SQL*Net TCP/IP	0.32	N/A				
SQL*Net V2:							
v	SQL*Net V2	35.59	0.2				
	TNS Listener			1183	1	x 176	176
	DECnet	0.12	N/A				
	LU6.2	0.14	N/A	7		x 2	
	OSI	0.16	N/A				
	SPX/IPX	4.51	N/A	695		x 152	
	Names Server	3.93	N/A	1806		x 215	
v	TCP/IP	0.12	N/A				
	Interchange	5.81	N/A				
	Network Manager	22.51	N/A	3266		x 323	
	Interchange Control Utility			1260		x 169	
	Interchange Listener			1224		x 168	
	Navigator			1187		x 168	
	Configuration Tool	N/A	0.2				
		Σ A	Σ B	Σ C			Σ D
	Subtotals =	35.71	0.20	1183			176
Summary							
A	Tot. Distribution Space (A)	35.71					(MB)
B	Tot. Database Space (B)	0.20					(MB)
C	#1 User Memory (C)			1183			(KB)
D	Additional User Memory (D)			176			(KB)
E	Virtual Memory Total (C+D)			1359			(KB)

Note: Each product selected for installation is indicated by (v). The total disk space required for each product is written in bold type.

After completing the calculations on these sizing sheets, we must fill in the **A**, **B**, and **E** totals from the individual sheets above into the final calculation sheet shown below.

Table 17. Total Disk Sizing Result Sheet

Total Space Requirements (MB)			
Transfer Table	Disk Space (A)	Database Space (B)	Virtual Memory (E)
Oracle Server Products	101.84	12.5	15.68
Precompiler Products: Development Cycle			
Precompiler Products: User Cycle	n/a	n/a	
Oracle Networking Products	35.71	0.2	1.36
Oracle Tools Products	7.22	0.6	4.25
Oracle Online Documentation		n/a	n/a
Additional installation space for files generated during installation	+ 21.00	n/a	n/a
Temporary space for relinking	40.00	n/a	n/a
TOTALS	205.77	13.3	21.29

Note: From the above Total Disk Space Sizing Summary sheet, we can estimate that the minimum space for those selected products for nonparallel database installation on the first sample cluster is about 250 MB.

5.3 Copying Files From Floppy Disk

You must (as root user) copy the install program scripts from the included floppy disk to the /tmp/hascripts directory.

```
# tar -xvf/dev/rfd0
```

5.4 Installation Worksheet

Make a copy of the worksheet from Appendix A that is appropriate for your failover type. Fill out the information on this worksheet as we go along, and keep it on hand. Much of the install is dependent on information from this worksheet.

5.5 Assign Network Adapters

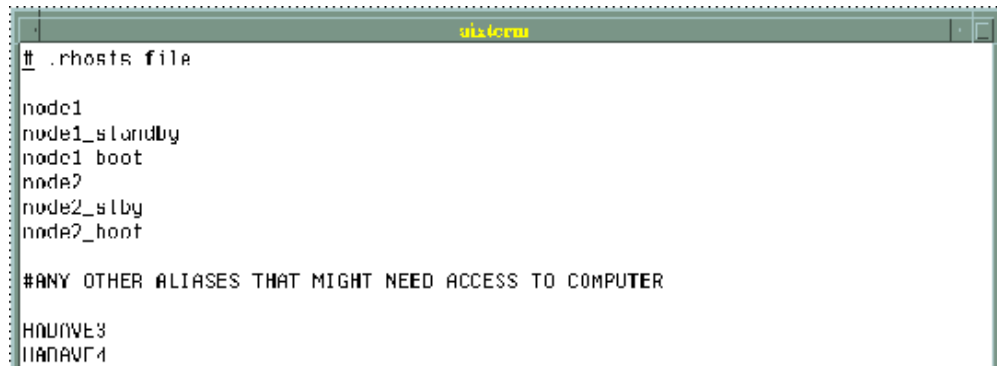
If your network is not already established, you need to assign and name all of the network adapters you will be using in the cluster. Each adapter has a label and an IP address (as seen on your worksheet). For more information on this step please contact your network administrator. Enter this information on your worksheet.

5.6 Update /rhosts and /etc/hosts

These files control communication between adapters on different nodes. These allow for name resolution as well as communication by IP address. The IP labels used in these examples are the defaults used by the Quick Install Program.

5.6.1 Update .rhosts on Both Nodes

As the root user, edit the .rhosts file adding the names of all adapters involved in cluster communication as well as the level of user allowed



```
aixterm
# .rhosts file

node1
node1_standby
node1_boot
node2
node2_slby
node2_hoot

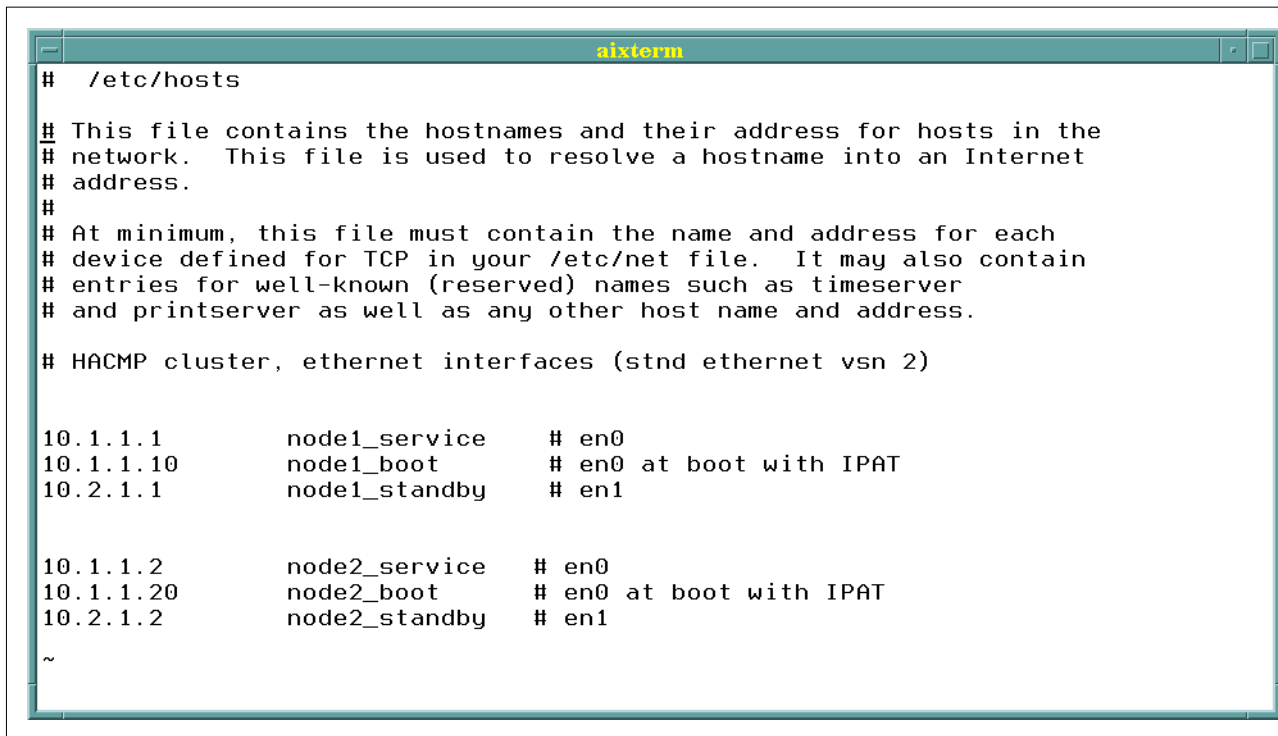
#ANY OTHER ALIASES THAT MIGHT NEED ACCESS TO COMPUTER

HACMPV3
HACMPV4
```

Figure 7. .rhosts File from node1

5.6.2 Update /etc/hosts on Both Nodes.

As the root user, edit the /etc/hosts file adding the names and IP addresses of all adapters involved in the cluster.



```
aixterm
# /etc/hosts

# This file contains the hostnames and their address for hosts in the
# network. This file is used to resolve a hostname into an Internet
# address.
#
# At minimum, this file must contain the name and address for each
# device defined for TCP in your /etc/net file. It may also contain
# entries for well-known (reserved) names such as timeserver
# and printserver as well as any other host name and address.

# HACMP cluster, ethernet interfaces (stand ethernet vsn 2)

10.1.1.1      node1_service    # en0
10.1.1.10    node1_boot       # en0 at boot with IPAT
10.2.1.1     node1_standby    # en1

10.1.1.2     node2_service    # en0
10.1.1.20   node2_boot       # en0 at boot with IPAT
10.2.1.2     node2_standby    # en1

~
```

Figure 8. /etc/hosts file from node1 Preinstallation

NOTE: Once you have entered these two files, you must check your adapters. After shutting down the system, and logging back in, type `host` followed by the adapter name. In Figure 9, we tested `node1_standby`:

A terminal window titled 'aixterm' showing the command 'host node1_standby' and its output 'node1 standby is 10.0.2.1'. The prompt is '# _'.

Figure 9. Adapter Test on `node1_standby`

You need to test every adapter in this way, and check that the IP address returned is the address expected. If not, you have a name resolution conflict on your network, and need to contact your network administrator.

5.7 Find Available Group & User ID Numbers

You need to find an available user ID number and an available group ID number. These numbers must be available on both nodes. Make sure you have these before continuing. The recommended user ID number is 201; the group ID number is 210. To check these, simply enter `smit user` or `smit group` at the command prompt, and enter `List All` at the menu. When you have done so, enter this information on your worksheet.

5.8 Find the Correct Terminal Emulation

`ORACLE_TERM` is the Oracle terminal emulation variable. When you have found one, enter it on your worksheet. The Oracle installer utility supports the following terminal types:

3151, 386*, Q303*, Q310*, ansi, dec, hft*, hp, iris, lft, sun, vt100, vt220, wy50, wy150, xlft, xsun, xsun50

* Multiple versions of these emulations

Once you have picked an emulation you think will work, you must test it. At the command prompt, enter `TERM=` to whatever emulation you picked. Then export `TERM`.

```
# TERM=lft
# export TERM
```

If everything seems to work fine, and you have the ability to backspace and destroy old text, you will be OK with that terminal emulation for the installation.

5.9 Creation of Logical Volumes

These raw logical volumes will store your database. They are created during the Quick Install program, but you are given the opportunity to adjust the default sizes. Check the Oracle 7 *Installation Guide* for a better explanation of the logical volumes. When you have decided upon the perfect sizes, enter them on your worksheet.

Remember, if you are using a Mutual Takeover configuration, you will have a set of logical volumes for each database. Therefore, you must pick the sizes for both sets of logical volumes.

Oracle Object	LV name	Oracle File Name	Suggested LV size (MB)	# of Logical Partitions
Control file 1	c1lv	/dev/rnode1_c1lv	N/A	N/A
Control file 2	c2lv	/dev/rnode1_c2lv	N/A	N/A
Control file 3	c3lv	/dev/rnode1_c3lv	N/A	N/A
System tablespace	dblv	/dev/rnode1_dblv	80 MB	20
Logfile 1 (instance 1)	log1lv	/dev/rnode1_log1lv	12 MB	3
Logfile 2 (instance 1)	log2lv	/dev/rnode1_log2lv	12 MB	3
Logfile 3 (instance 1)	log3lv	/dev/rnode1_log3lv	12 MB	3
Rollback tablespace	rollv	/dev/rnode1_rollv	20 MB	5
Temp tablespace	templv	/dev/rnode1_templv	12 MB	3
User tablespace	userlv	/dev/rnode1_userlv	12 MB	3
Tools tablespace	toollv	/dev/rnode1_toollv	40 MB	10

Figure 10. Logical Volumes for Oracle 7 Storage

The sizes given for the logical volumes are the recommended to load Oracle 7 and the sample database that comes with it. These are more than double the minimum sizes laid out by Oracle. For more detailed information, please see the Oracle7 documentation.

Note: A logical partition is equivalent to 4 MB.

Chapter 6. Running the Quick Install Program

This chapter outlines the steps involved in installing an Oracle database on a highly available cluster. The Quick Install Program used here does not apply to Oracle Parallel Server.

Now that you have the necessary information to begin installation, follow these steps, and let the Quick Install program lead the way. The steps shown are for node 1. Any changes for node 2 are specified along the way.

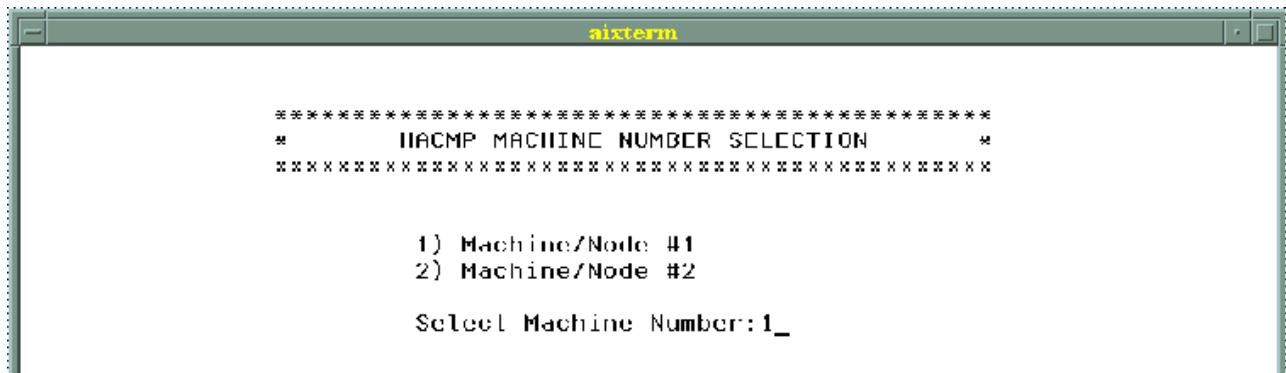
NOTE: If you choose to leave the Installer, or are forced out due to an error, running the script called `cleanup` and then shutting down the node will prepare the machine for the installation process again.

Start the Quick Install Program by running the following commands as root:

```
# cd /tmp/hascripts
# ./setup
```

6.1 Declaring Node Number

This first prompt asks you to enter the node which you are currently working on. If you are working on node 1, enter **1**. If you are working on node 2, enter **2**.



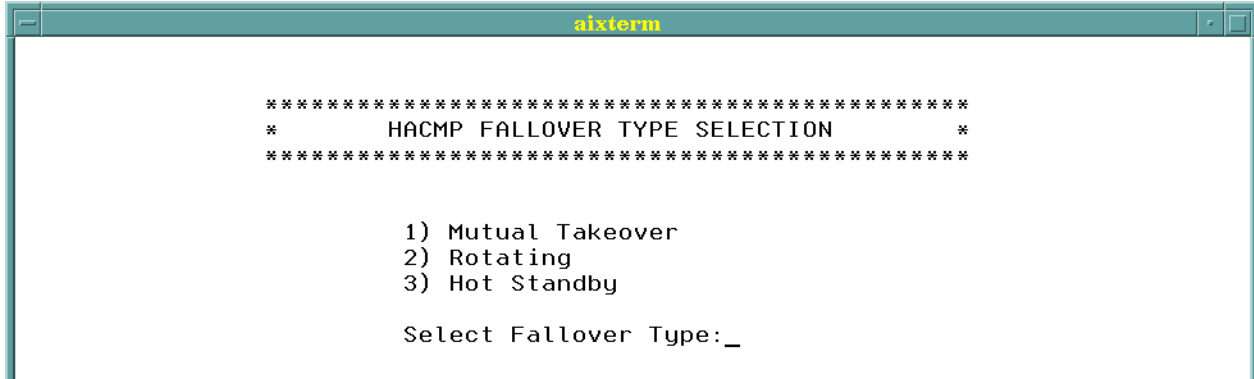
```
aixterm
*****
*           HACMP MACHINE NUMBER SELECTION           *
*****

1) Machine/Node #1
2) Machine/Node #2

Select Machine Number:1_
```

6.2 Declaring Failover Type

You should have already chosen which failover type you wish to use for this cluster configuration. If you want a Mutual Takeover configuration, enter **1**; for Rotating Standby, enter **2**, and for Hot Standby, enter **3**.



```
aixterm

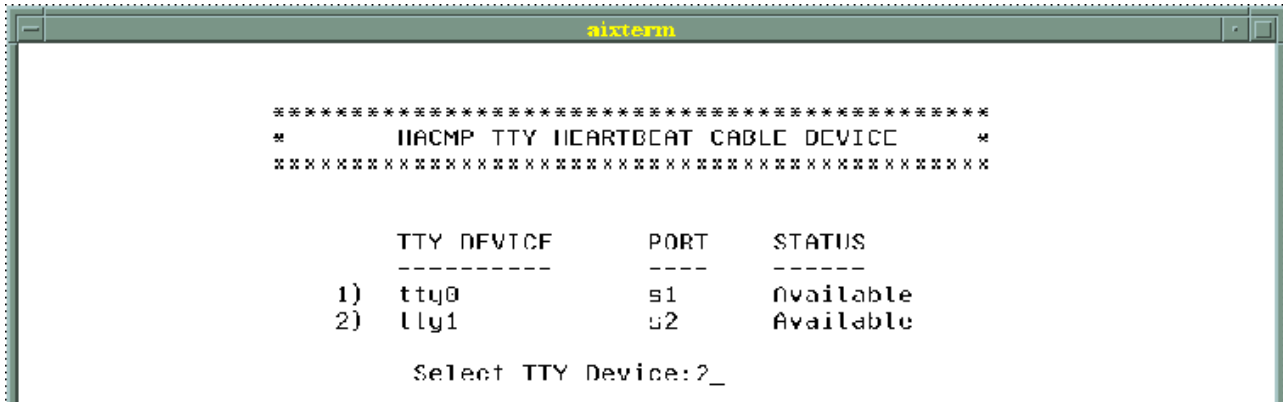
*****
*      HACMP FALLOVER TYPE SELECTION      *
*****

    1) Mutual Takeover
    2) Rotating
    3) Hot Standby

Select Fallover Type:_
```

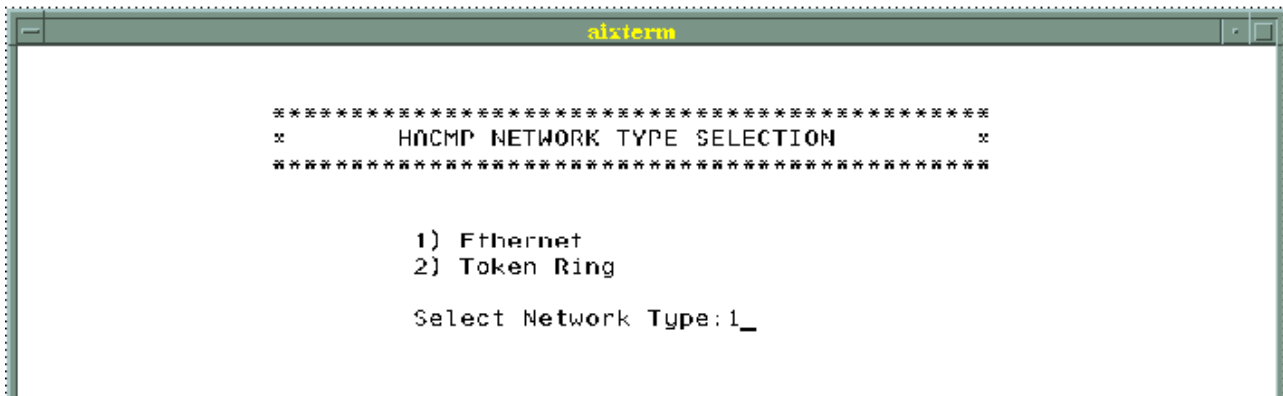
6.3 Choose TTY Port

The available TTY lines are presented to you. This line supports the “heartbeat” between the nodes in the cluster. Please select the appropriate choice for your machine:



6.4 Select Network Type

Select **1** if your network is using Ethernet connections, or select **2** if your network is using token-ring connections.



6.5 Adapter Configuration Screen

The next screen does all the adapter configuration for you. You must enter all the appropriate information, however. You should have gathered this information while completing the configuration worksheet. If you have an existing network, and are using IP addresses from that network, the Quick Install Program will locate this information, and present it on screen. If you do not have an existing network, this screen will come up with the default information. Either way, consult your worksheet for the proper adapter information and make any changes on this screen.

To make any changes to the label, address, mask, or slot position of the adapters, just follow these instructions:

- To enter/change an adapter label, enter **L** and a number (Ex. **L0**, **L1**).
- To enter/change an IP Address, just enter **A** and a number (Ex. **A0**, **A2**).
- To enter/change the subnet mask for an adapter just enter **M** and a number (Ex. **M0**, **M1**).
- To swap the position of two adapters, just enter **S** followed by the numbers you want switched (Ex. **S12**).

When you have filled in all the appropriate information enter **C** to move on.

Below, we will display the screens of node1 and node2 for all three configurations, as they are all unique. Make sure you are using the appropriate information for your configuration.

Mutual Takeover Node 1:

```
*****
*           HACMP ADAPTER CONFIGURATION           *
*                   Mutual Takeover                   *
*****

   HATYPE  IP_LABEL      IP_ADDR      IP_MASK      ADAPTER  SLOT
   -----  -
0)  service node1        10.0.1.1     255.255.255.0  N/A
1)  boot   node1_boot    10.0.1.11    255.255.255.0  ent0    4
2)  standby node1_standby 10.0.2.1     255.255.255.0  ent1    6

A)ddress L)abel M)ask S)wap adapter C)ontinue: _
```

Mutual Takeover Node 2:

```
aixterm

*****
*      HACMP ADAPTER CONFIGURATION      *
*      Mutual Takeover                    *
*****

  HATYPE  IP_LABEL      IP_ADDR      IP_MASK      ADAPTER  SLOT
  -----  -
0)  service node2        10.0.1.2      255.255.255.0  N/A
1)  boot   node2_boot    10.0.1.12    255.255.255.0  ent0      4
2)  standby node2_standby 10.0.2.2     255.255.255.0  ent1      6

A)ddress L)abel M)ask S)wap adapter C)ontinue: _
```

```
aixterm

*****
*      HACMP ADAPTER CONFIGURATION      *
*      Rotating                          *
*****

  HATYPE  IP_LABEL      IP_ADDR      IP_MASK      ADAPTER  SLOT
  -----  -
0)  shared node1        10.0.1.1     255.255.255.0  N/A
1)  boot   node1_boot    10.0.1.11    255.255.255.0  ent0      4
2)  standby node1_standby 10.0.2.1     255.255.255.0  ent1      6

A)ddress L)abel M)ask S)wap adapter C)ontinue: _
```

Rotating Standby Node 2:

```
aixterm
*****
*          HACMP ADAPTER CONFIGURATION          *
*                   Rotating                   *
*****

HATYPE  IP_LABEL      IP_ADDR      IP_MASK      ADAPTER  SLOT
-----  -
0)  shared  node1         10.0.1.1     255.255.255.0  N/A
1)  boot    node2_boot     10.0.1.12   255.255.255.0  ent0     4
2)  standby node2_standby  10.0.2.2    255.255.255.0  ent1     6

A)ddress L)abel M)ask S)wap adapter C)ontinue: _
```

Hot Standby Node 1:

```
aixterm
*****
*          HACMP ADAPTER CONFIGURATION          *
*                   Hot Standby                *
*****

HATYPE  IP_LABEL      IP_ADDR      IP_MASK      ADAPTER  SLOT
-----  -
0)  service node1         10.0.1.1     255.255.255.0  N/A
1)  boot    node1_boot     10.0.1.11   255.255.255.0  ent0     4
2)  standby node1_standby  10.0.2.1    255.255.255.0  ent1     6

A)ddress L)abel M)ask S)wap adapter C)ontinue: _
```

Hot Standby Node 2:

```
*****
*           HACMP ADAPTER CONFIGURATION           *
*                   Hot Standby                   *
*****

  HATYPE  IP_LABEL      IP_ADDR      IP_MASK      ADAPTER  SLOT
  -----  -
1)  service node2       10.0.1.2      255.255.255.0  ent0     4
2)  standby node2_standby 10.0.2.2      255.255.255.0  ent1     6

A)ddress L)abel M)ask S)wap adapter C)ontinue: _
```

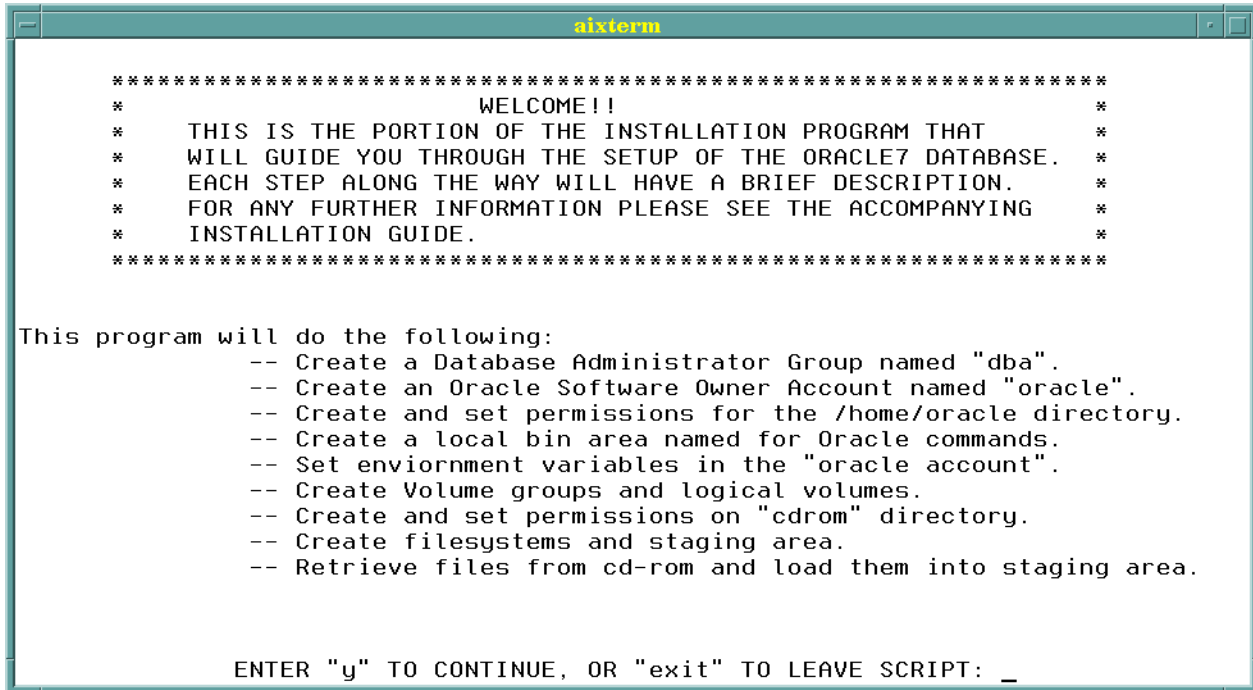
NOTE: On node 2, after the adapters are configured, you will be prompted for the label of node 1's boot adapter. This information is necessary for communication between the nodes during installation.

```
*****
*           HACMP NODE1 BOOT LABEL ENTRY           *
*****

Enter Node #1 Boot Adapter Label (ex. node1_boot): _
```

6.6 Information Menus

You are now entering the Oracle portion of the install. This Welcome screen lays out what will be done in this section of the Quick Install. Enter **y** to continue, or type `exit` to leave the Installer.



```
*****
*                                     *
*                WELCOME!!          *
*   THIS IS THE PORTION OF THE INSTALLATION PROGRAM THAT   *
*   WILL GUIDE YOU THROUGH THE SETUP OF THE ORACLE7 DATABASE. *
*   EACH STEP ALONG THE WAY WILL HAVE A BRIEF DESCRIPTION. *
*   FOR ANY FURTHER INFORMATION PLEASE SEE THE ACCOMPANYING *
*   INSTALLATION GUIDE.          *
*****

This program will do the following:
-- Create a Database Administrator Group named "dba".
-- Create an Oracle Software Owner Account named "oracle".
-- Create and set permissions for the /home/oracle directory.
-- Create a local bin area named for Oracle commands.
-- Set environment variables in the "oracle account".
-- Create Volume groups and logical volumes.
-- Create and set permissions on "cdrom" directory.
-- Create filesystems and staging area.
-- Retrieve files from cd-rom and load them into staging area.

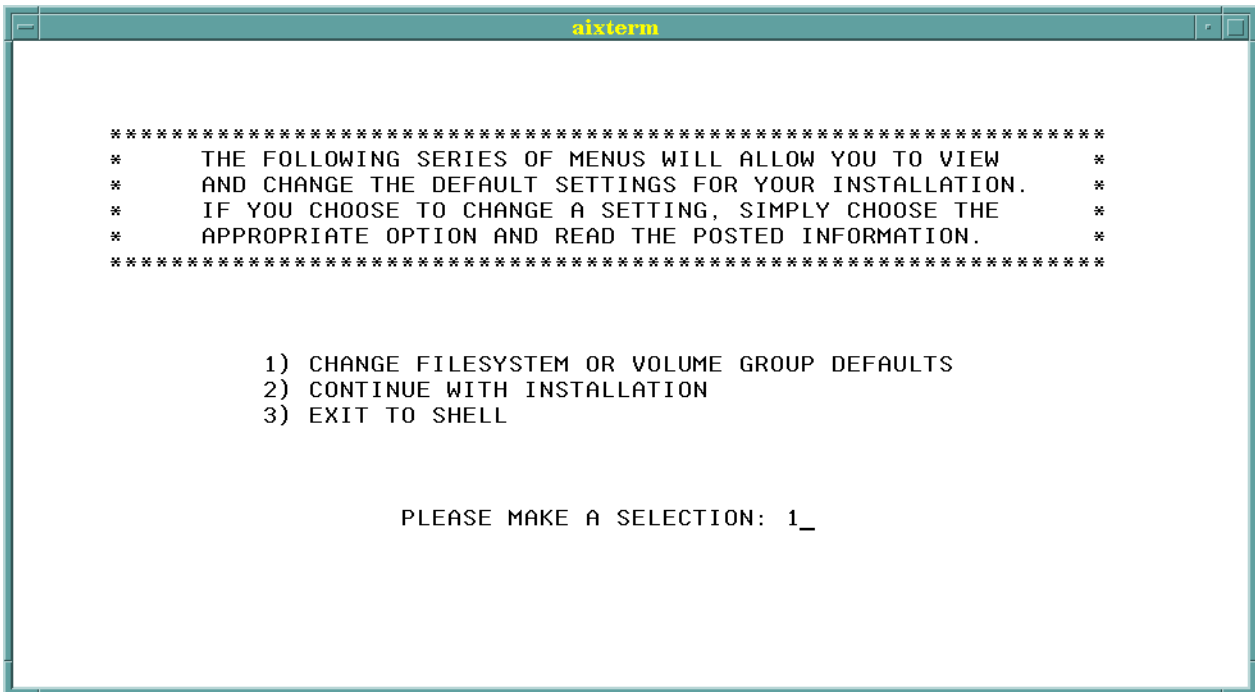
ENTER "y" TO CONTINUE, OR "exit" TO LEAVE SCRIPT: _
```

The next screen is a warning to make sure that you have planned the disk space sizing properly for installation. You should have planned this in the previous chapter; however, if you need to leave the program to make adjustments at this time, type `exit`. There are several explanations of disk sizing in this manual. If you have already done this, enter **y** to continue.

6.7 Change Default Names

After the warning on space requirements, you will come to a set of menus that allows you to change the default names for the volume groups and filesystems that will be created on your cluster. The volume groups will float between the nodes of the cluster during failover and will hold the filesystems where database scripts and data reside. Follow the menu options to make changes to the default names.

NOTE: These menus will not appear in node 2. The names have been read from node 1.



```
*****
*   THE FOLLOWING SERIES OF MENUS WILL ALLOW YOU TO VIEW   *
*   AND CHANGE THE DEFAULT SETTINGS FOR YOUR INSTALLATION. *
*   IF YOU CHOOSE TO CHANGE A SETTING, SIMPLY CHOOSE THE *
*   APPROPRIATE OPTION AND READ THE POSTED INFORMATION.    *
*****

1) CHANGE FILESYSTEM OR VOLUME GROUP DEFAULTS
2) CONTINUE WITH INSTALLATION
3) EXIT TO SHELL

PLEASE MAKE A SELECTION: 1_
```

The defaults for a Mutual Takeover configuration are:

```
node1vg node1_scripts
node2vg node2_scripts
```

The defaults for a Rotating or Hot Standby configuration are:

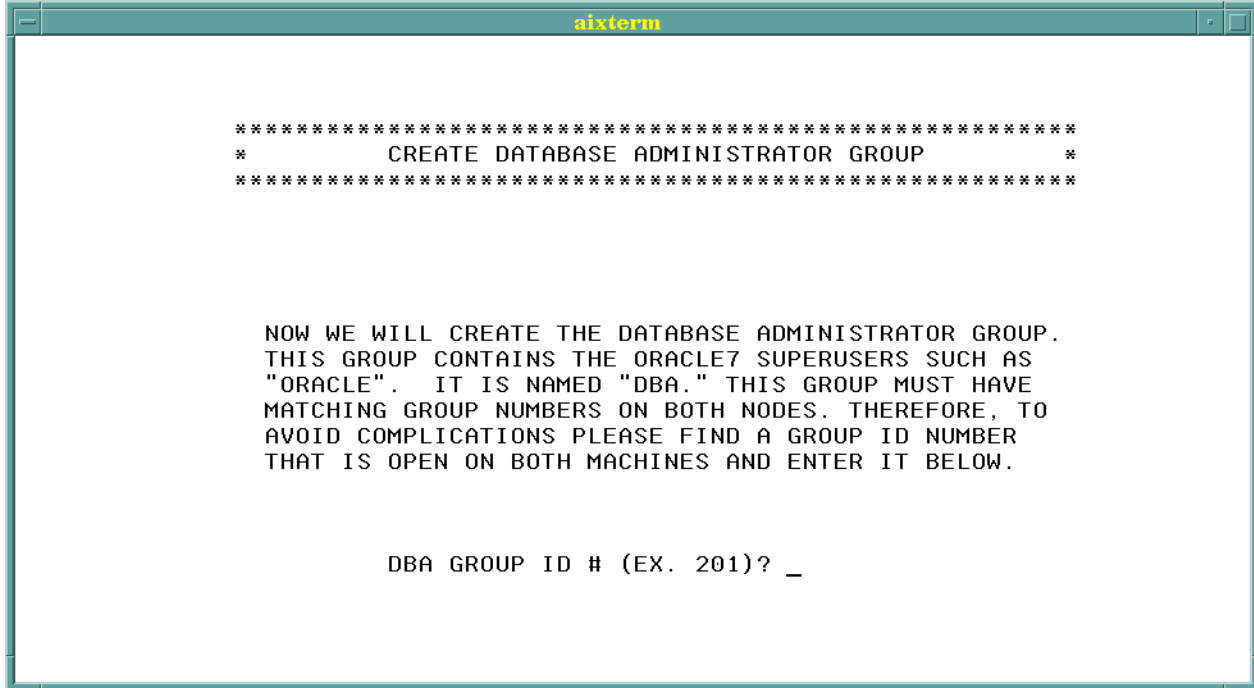
```
node1vg node1_scripts
```

When you have finished with the changes, enter **2** to continue on or **3** to quit the program.

6.8 Create DBA Group

This prompt asks you to enter the available group ID number that you found on your worksheet. This Database Administrator Group (dba) contains the database superusers. Once you enter the number (the default is 201), the user will be

created for you. If you get an error message, the number you entered is not available, or there is an existing dba group on your system.



```
aixterm

*****
*          CREATE DATABASE ADMINISTRATOR GROUP          *
*****

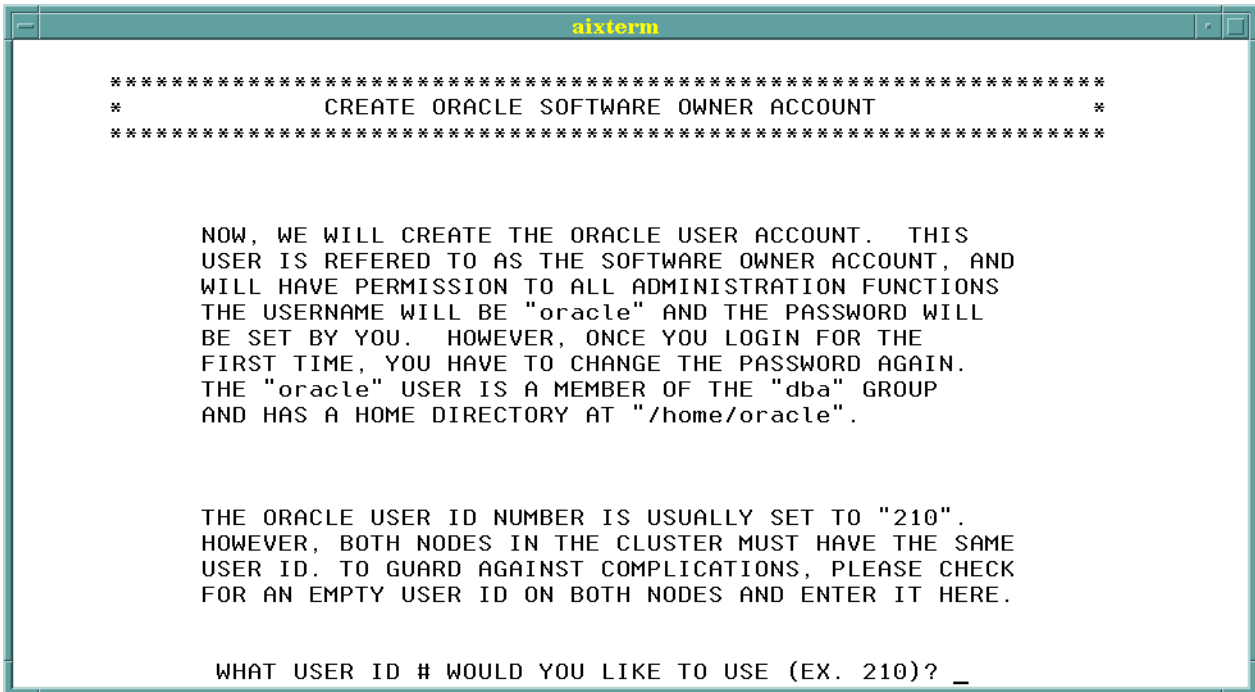
NOW WE WILL CREATE THE DATABASE ADMINISTRATOR GROUP.
THIS GROUP CONTAINS THE ORACLE7 SUPERUSERS SUCH AS
"ORACLE". IT IS NAMED "DBA." THIS GROUP MUST HAVE
MATCHING GROUP NUMBERS ON BOTH NODES. THEREFORE, TO
AVOID COMPLICATIONS PLEASE FIND A GROUP ID NUMBER
THAT IS OPEN ON BOTH MACHINES AND ENTER IT BELOW.

DBA GROUP ID # (EX. 201)? _
```

6.9 Create Oracle User Account

The Oracle user (oracle) is the superuser. This account has access to all areas of the database and permissions on all administrative functions. This user has a home directory at /home/oracle on the system, and all the local database files are located there. You are prompted for the user ID number. This number was found earlier, and should be written on your worksheet. The default is 210. After you enter this number, the user will be created. If you receive an error message, the

number you entered is not available, or an *oracle* user already exists on your system.



```
aixterm
*****
*          CREATE ORACLE SOFTWARE OWNER ACCOUNT          *
*****

NOW, WE WILL CREATE THE ORACLE USER ACCOUNT.  THIS
USER IS REFERED TO AS THE SOFTWARE OWNER ACCOUNT, AND
WILL HAVE PERMISSION TO ALL ADMINISTRATION FUNCTIONS
THE USERNAME WILL BE "oracle" AND THE PASSWORD WILL
BE SET BY YOU.  HOWEVER, ONCE YOU LOGIN FOR THE
FIRST TIME, YOU HAVE TO CHANGE THE PASSWORD AGAIN.
THE "oracle" USER IS A MEMBER OF THE "dba" GROUP
AND HAS A HOME DIRECTORY AT "/home/oracle".

THE ORACLE USER ID NUMBER IS USUALLY SET TO "210".
HOWEVER, BOTH NODES IN THE CLUSTER MUST HAVE THE SAME
USER ID.  TO GUARD AGAINST COMPLICATIONS, PLEASE CHECK
FOR AN EMPTY USER ID ON BOTH NODES AND ENTER IT HERE.

WHAT USER ID # WOULD YOU LIKE TO USE (EX. 210)? _
```

Once this user is created, you need to set the password for the account. This password is temporary. When you log in to the machine for the first time as *oracle*, you will be prompted to change the password again.

6.10 Creating Filesystems, Directories, and Setting Permissions

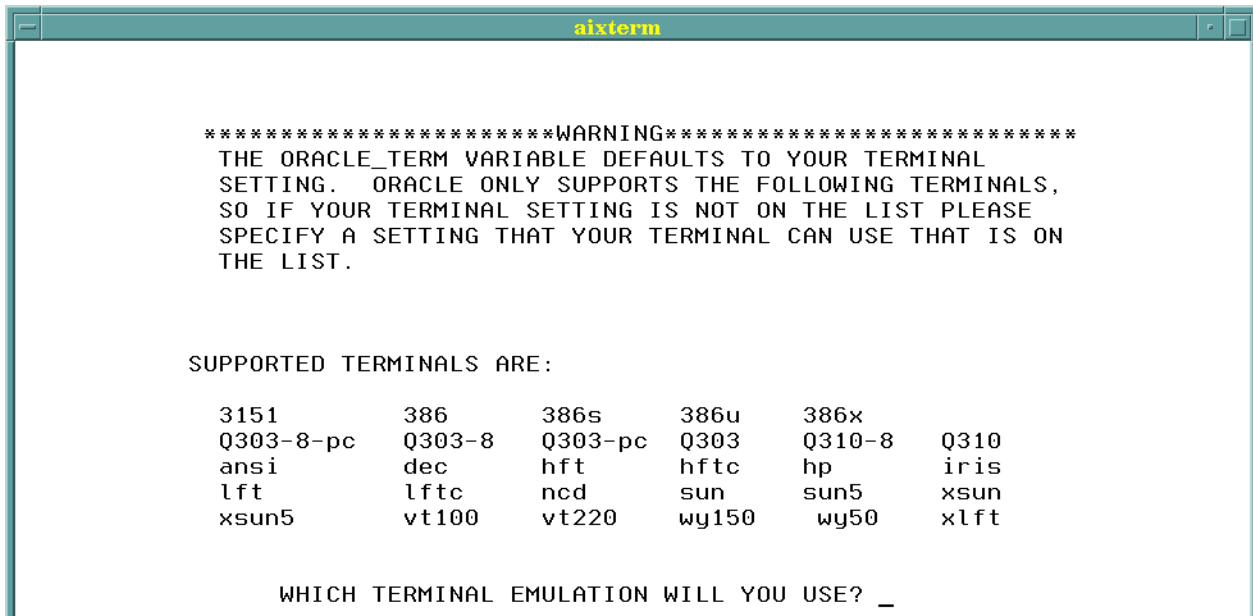
This area of the Quick Install Program requires no user input. There are status screens flashed on screen, as well as messages of success or error. The following is occurring during this time:

- First, the `/home/oracle` filesystem is created. This filesystem is created on `rootvg`, and it holds all the local Oracle 7 files and scripts. The volume group is checked for sufficient disk size (300 MB) and then created. Permissions are then set on this directory.
- Next, a local `bin` is created for the Oracle 7 database to store commands in. Permissions are set on this directory (`/usr/bin`).
- An entry for the Oracle Listener is created in the `/etc/services` file. The Listener handles communication between the users and the database.
- Lastly, another filesystem is created to be a temporary staging area for the Oracle 7 files. The CD-ROM files are copied and then linked to this area (`/home/oracle/stage`) for installation. `Rootvg` is checked for sufficient space before the filesystem is created. The proper permissions are then set on the `/home/oracle/stage` directory.

This directory may be deleted after installation is complete.

6.11 Setting Oracle User Environment

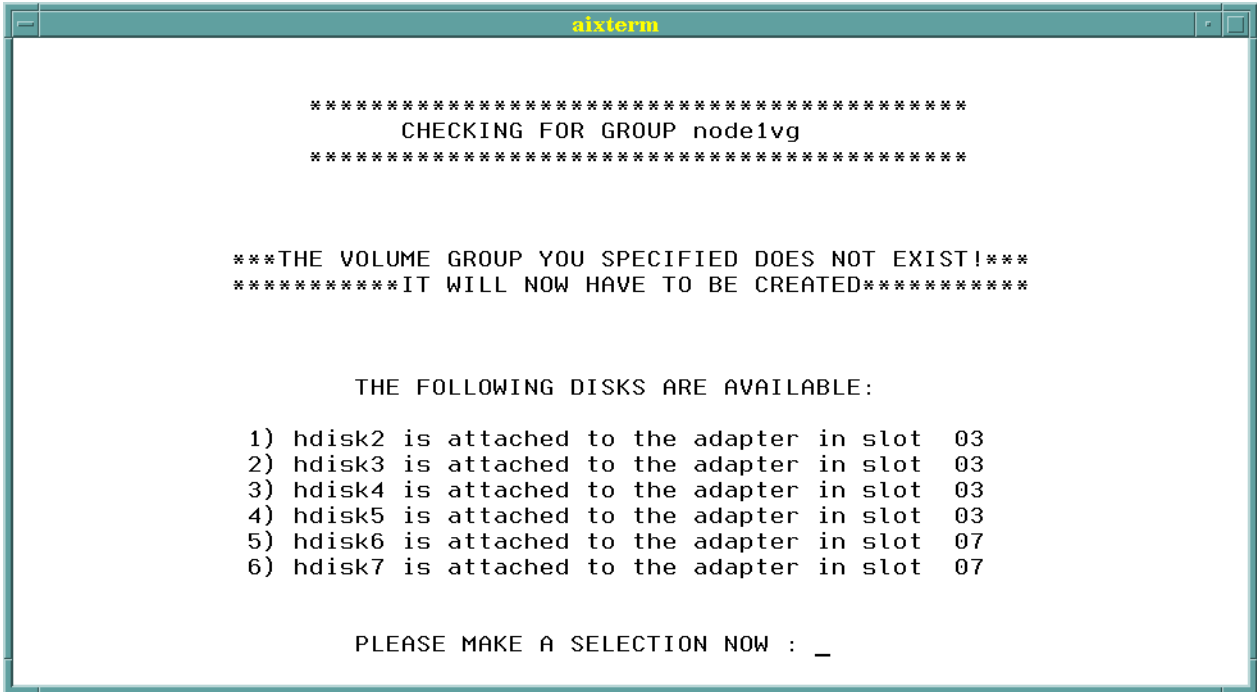
Oracle 7 requires several variables to be properly set in the superuser's environment. These variables, `ORACLE_HOME`, `ORACLE_TERM`, and `ORACLE_SID` are set for you. The only information you need to enter is the terminal type you found earlier and entered on your worksheet. This type must be compatible with Oracle 7 as well as with your hardware and must allow for you to backspace or delete text.



6.12 Creating Volume Groups

The volume group names you entered earlier must now be created if they do not already exist. This part of the Quick Install Program will check the name of the volume group and see if it exists. If so, it will check it for sufficient disk space and

continue on. If the volume group does not exist, you will be asked to create it, and you are shown a list of available disks. Select a disk by entering its line number.



```
*****
CHECKING FOR GROUP node1vg
*****

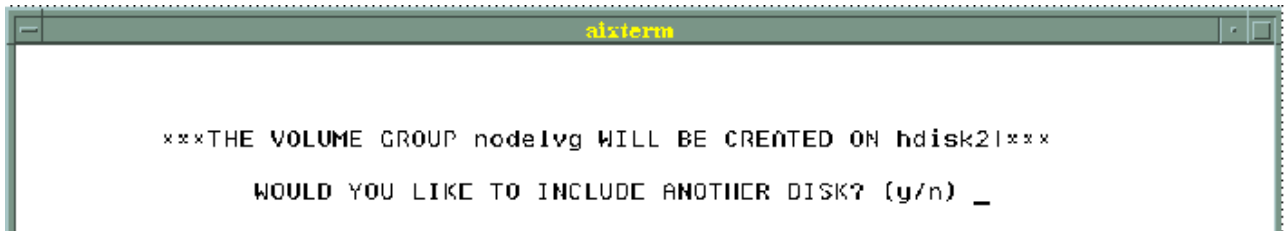
***THE VOLUME GROUP YOU SPECIFIED DOES NOT EXIST!***
*****IT WILL NOW HAVE TO BE CREATED*****

THE FOLLOWING DISKS ARE AVAILABLE:

1) hdisk2 is attached to the adapter in slot 03
2) hdisk3 is attached to the adapter in slot 03
3) hdisk4 is attached to the adapter in slot 03
4) hdisk5 is attached to the adapter in slot 03
5) hdisk6 is attached to the adapter in slot 07
6) hdisk7 is attached to the adapter in slot 07

PLEASE MAKE A SELECTION NOW : _
```

Once you have selected a disk, you are shown the disk you have already chosen and asked if you would like to use another disk in the creation of the volume group.



```
***THE VOLUME GROUP node1vg WILL BE CREATED ON hdisk2!***

WOULD YOU LIKE TO INCLUDE ANOTHER DISK? (y/n) _
```

If you select no (n), the volume group will be created. If you select yes (y), you are shown the available disks again and asked to choose. When you have selected all the disks you need, select **no**, and the volume group will be created.

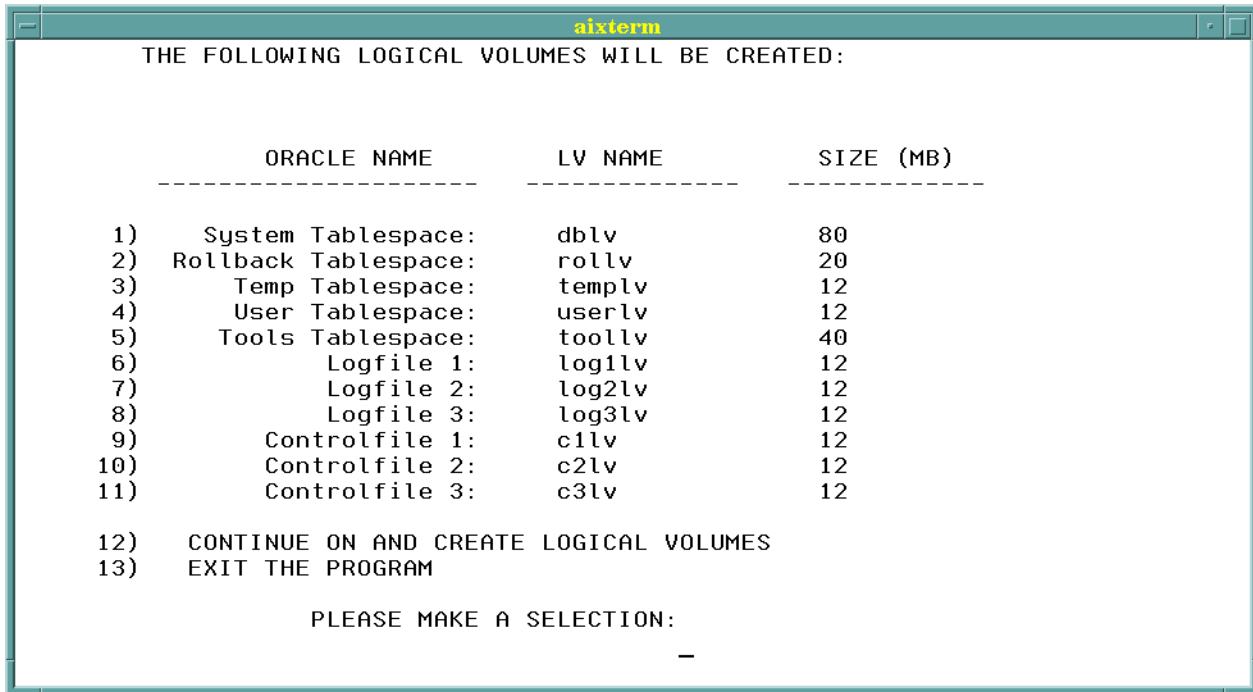
As the volume group is being created, the filesystem you named earlier is also being created. This filesystem will store the Oracle 7 control scripts for the database it is affiliated with.

If there is only one disk available, it will automatically be chosen and used to create the volume group.

NOTE: It is a good idea, when using two or more disks to create a volume group, to select disks attached to adapters in different slots to allow for mirroring. Mirroring will make two copies of the logical volumes on the volume group, making the cluster even more secure against failure. You will be prompted after creation as to whether or not you want mirroring.

6.13 Creation of Logical Volumes

The list of default logical volumes now appears.



```
THE FOLLOWING LOGICAL VOLUMES WILL BE CREATED:
```

	ORACLE NAME	LV NAME	SIZE (MB)
1)	System Tablespace:	dblv	80
2)	Rollback Tablespace:	rollv	20
3)	Temp Tablespace:	templv	12
4)	User Tablespace:	userlv	12
5)	Tools Tablespace:	toollv	40
6)	Logfile 1:	log1lv	12
7)	Logfile 2:	log2lv	12
8)	Logfile 3:	log3lv	12
9)	Controlfile 1:	c1lv	12
10)	Controlfile 2:	c2lv	12
11)	Controlfile 3:	c3lv	12
12)	CONTINUE ON AND CREATE LOGICAL VOLUMES		
13)	EXIT THE PROGRAM		

PLEASE MAKE A SELECTION:

—

You will notice that this list matches the list from your worksheet. Match up the sizes with those you have chosen for your database (the defaults are shown). If you have any changes, just select the line number of the logical volume you wish to change. If you want to continue enter **12**, and if you want to exit at this point, enter **13**.

When you continue on, you will see each logical volume being created one by one. This may take some time.

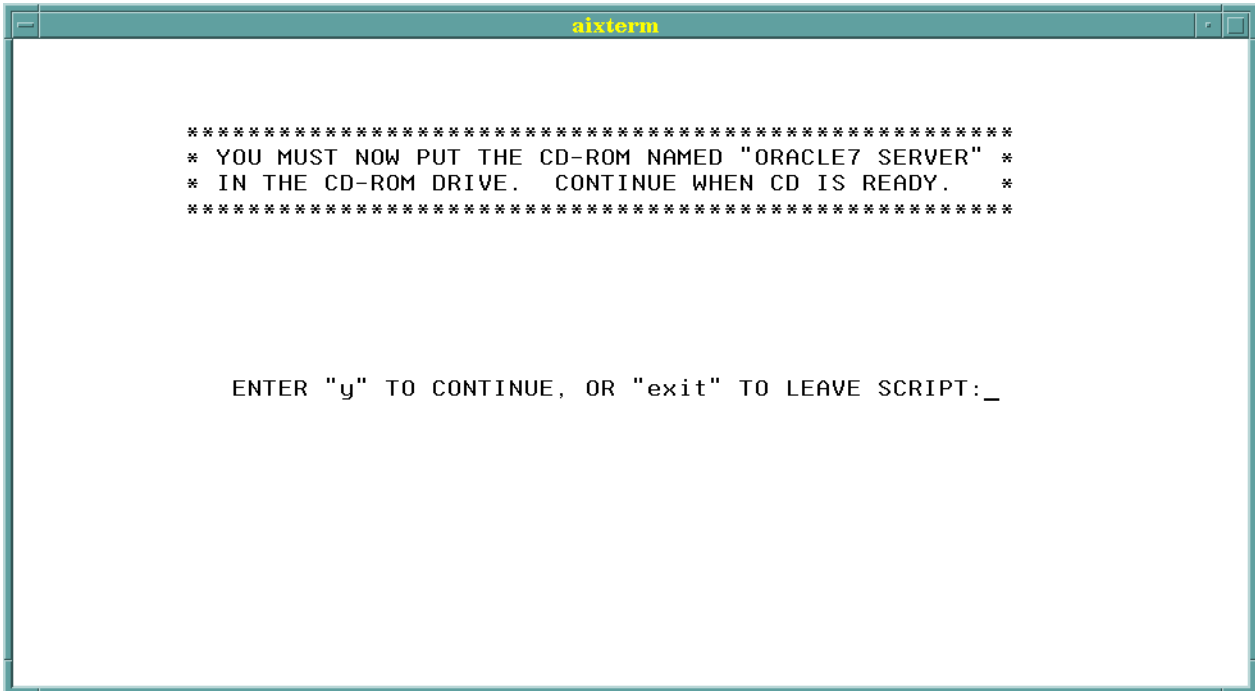
When all the logical volumes have been created, the permissions will be set on all of them.

NOTE: If you chose a Mutual Takeover Configuration, you have two databases; you therefore must have two volume groups, and two filesystems for storage. Therefore, you will now be taken back to the Create Volume Group prompt for a second time through.

NOTE: On node 2, the volume groups and logical volumes will be imported from node 1.

6.14 Mount CD-ROM

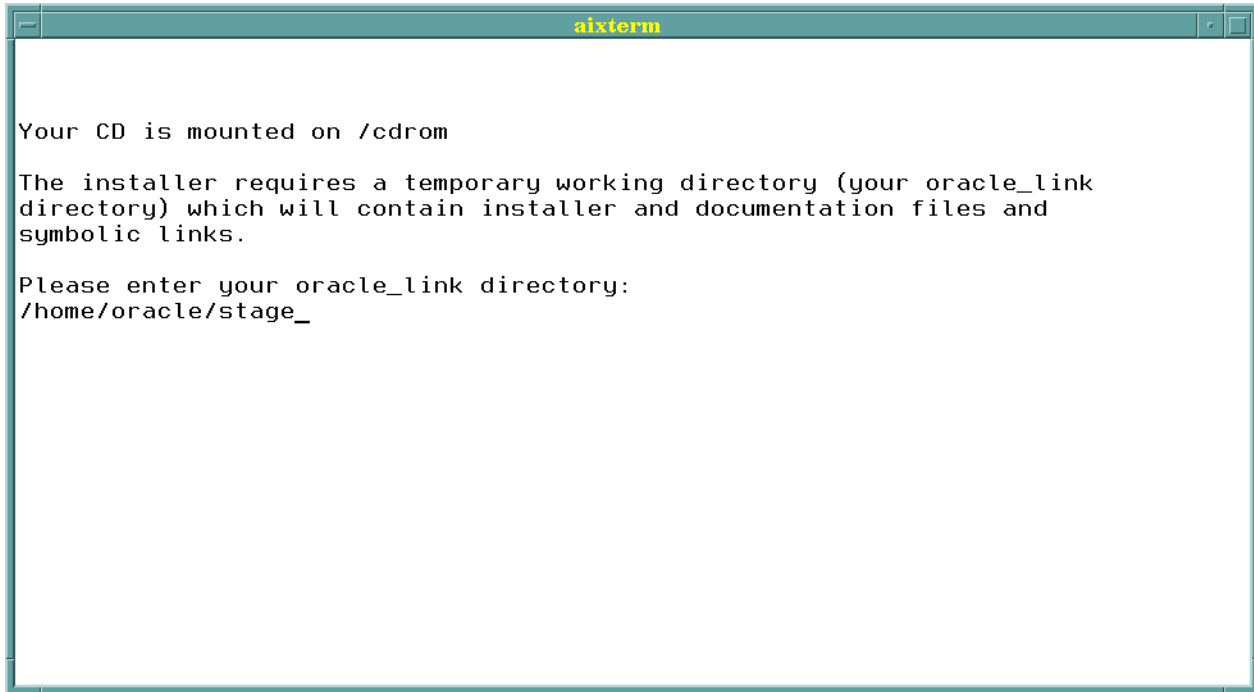
Make sure you have the CD-ROM marked "Oracle 7 Server" in the drive. When you are ready, enter **y** to continue.

The image shows a terminal window titled "aixterm". The window contains the following text:

```
*****  
* YOU MUST NOW PUT THE CD-ROM NAMED "ORACLE7 SERVER" *  
* IN THE CD-ROM DRIVE.  CONTINUE WHEN CD IS READY.  *  
*****  
  
ENTER "y" TO CONTINUE, OR "exit" TO LEAVE SCRIPT: _
```

6.15 Linking Oracle 7 Files to Staging Area

The install files will now be loaded into a staging area. You are prompted for the name of the staging area for oracle linking, which was created for you earlier. At the prompt, type: `/home/oracle/stage`.



```
aixterm

Your CD is mounted on /cdrom

The installer requires a temporary working directory (your oracle_link
directory) which will contain installer and documentation files and
symbolic links.

Please enter your oracle_link directory:
/home/oracle/stage_
```

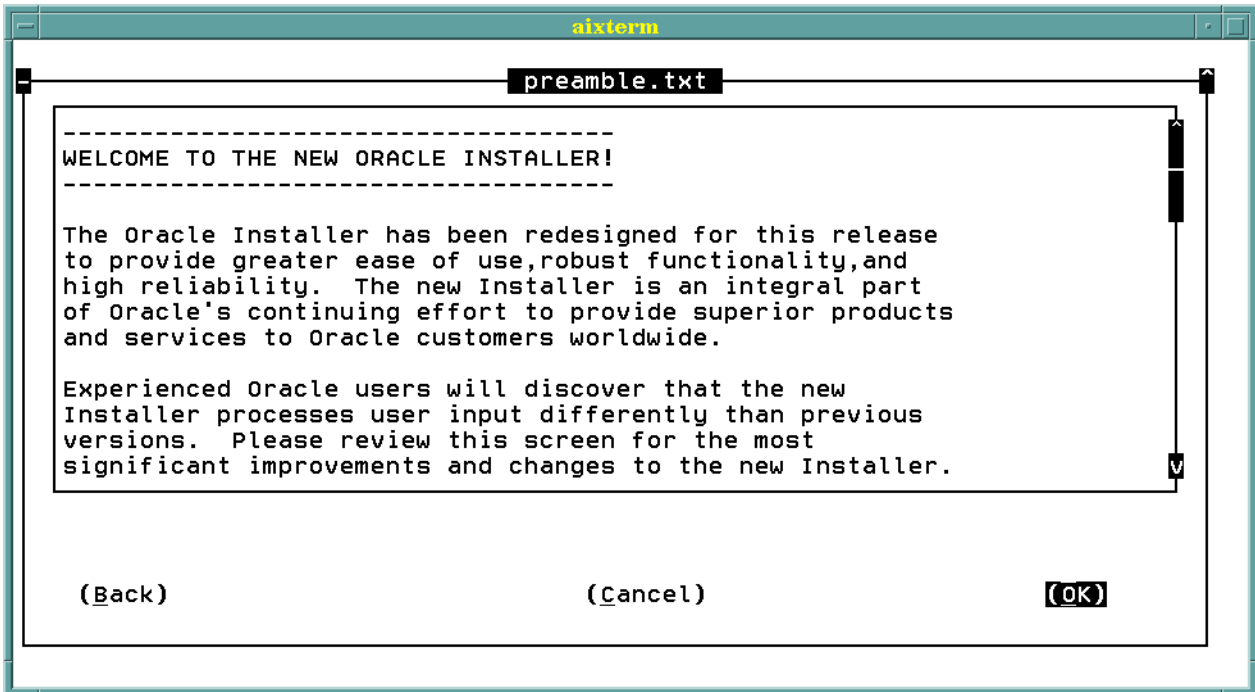
NOTE: This may take up to forty-five minutes.

6.16 Running The Preinstallation Script

When the linking has been completed, the Oracle 7 preinstallation script, `rootpre.sh` is run automatically. This file sets some environment variables.

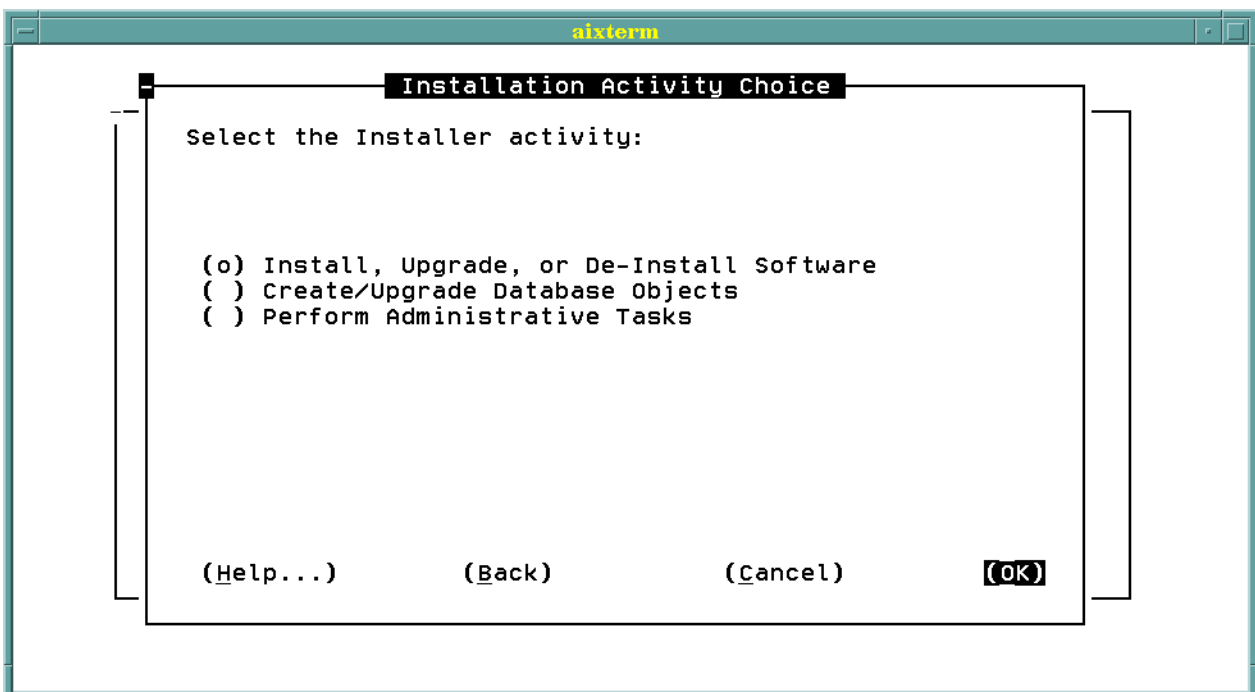
6.17 Oracle 7 Welcome

You are now launched into the Oracle 7 Install program. There are many prompts.



6.18 Enter Installer Activity

After you have started the Installer, you are asked to choose an Installer activity. You can choose to **Install, Upgrade, or De-Install Software**.



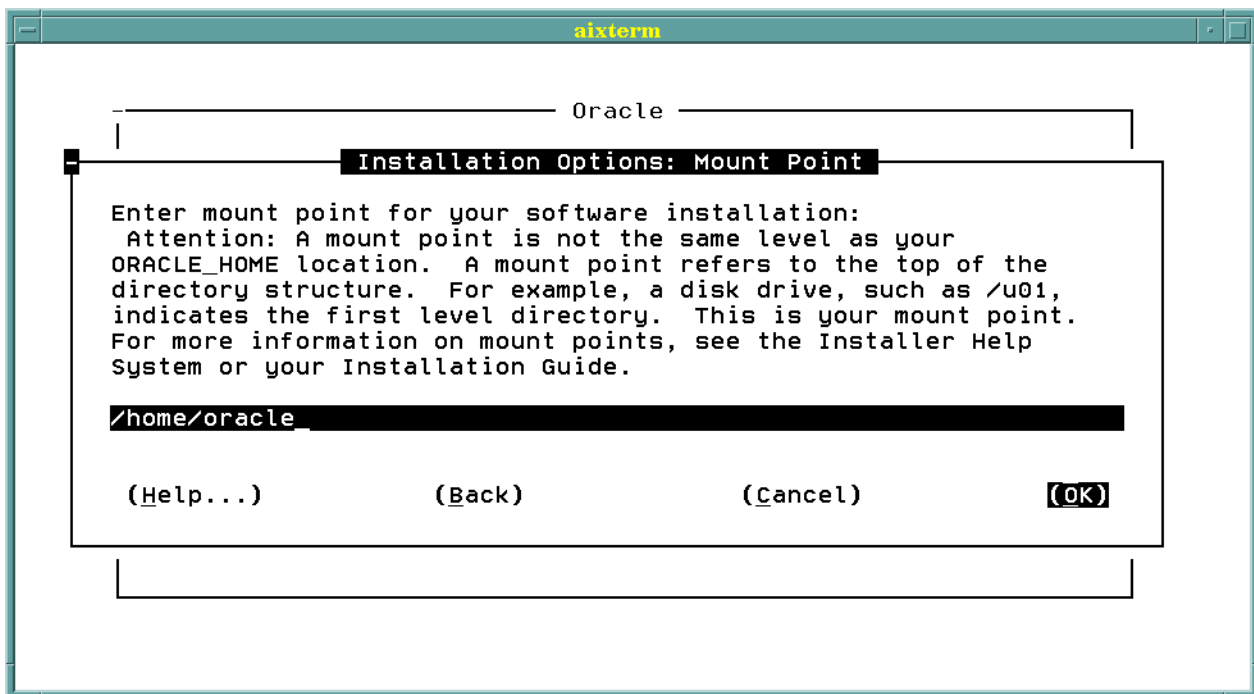
Note: This screen has radio buttons (the areas next to the options). To click the radio button, use the **Tab** key to move the highlighted cursor to the selectable fields, the up and down arrow keys to move the cursor to the option you want, and the **Space bar** to select that option.

6.19 Installation Options

At the Installation Options menu, you select **Install New Product**.

6.20 Entering the Mount Point

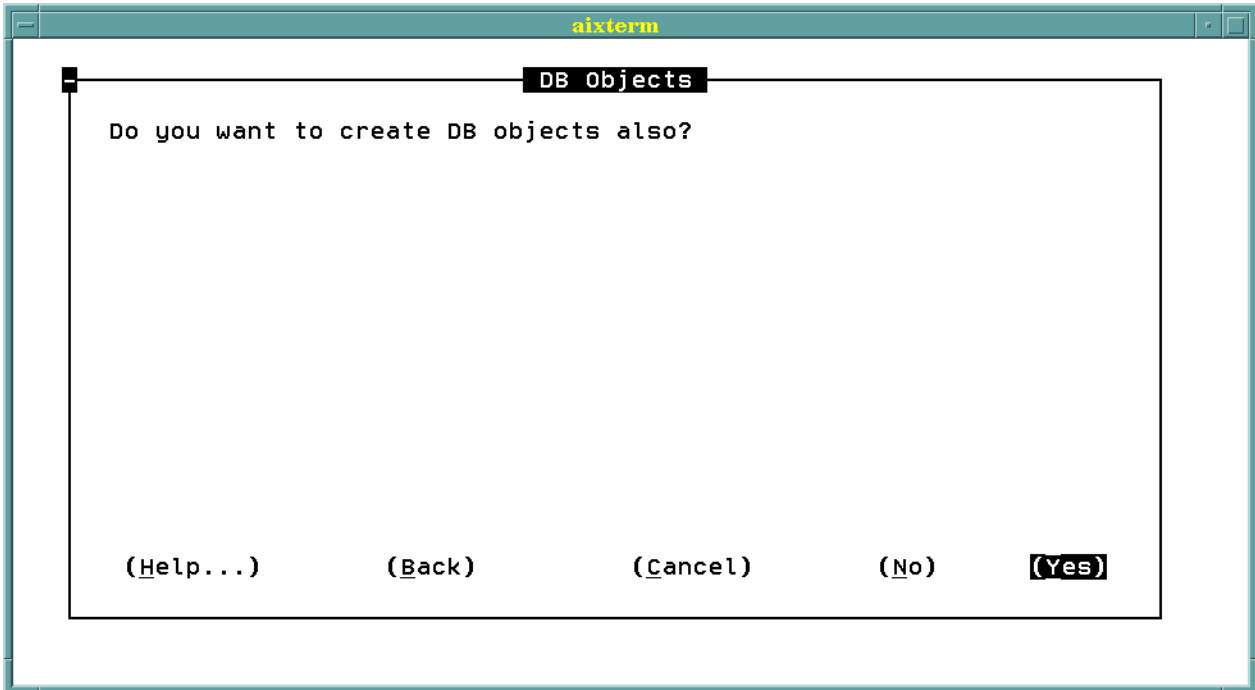
You must define the Mount Point for the software installation. This is different from the ORACLE_HOME pathname. The ORACLE_HOME variable is set to /home/oracle/app/oracle/product/7.3.2. The mount point is /home/oracle.



The next prompt will fill in the complete ORACLE_HOME location. Again, it should be /home/oracle/app/oracle/product/7.3.2. If that is correct, just enter **OK**.

6.21 Create Database Objects

Select **Yes** in the DB Objects menu to create the database. The database prompts will come later in the install. If you miss this step, the database creation prompts will be skipped.



6.22 Information Prompts

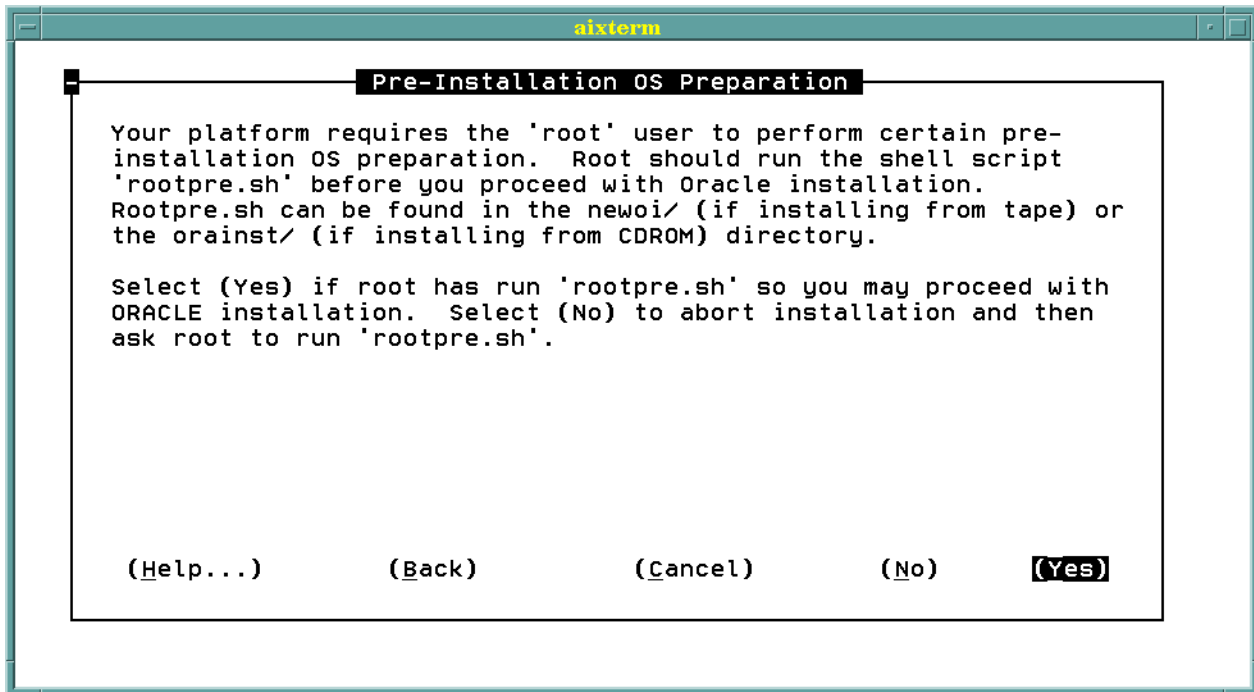
Continue on through the next series of information prompts. These provide you with information such as log locations and fixes. Just answer **OK** to these screens.

6.23 OPS Install

You do not wish to install the products on all the nodes that comprise the cluster for the Oracle 7 Parallel Server. Select **No** from the OPS Install prompt.

6.24 Preinstallation Preparation

The Quick Install Program has run the *rootpre.sh* preinstallation program for you; so select **Yes**.

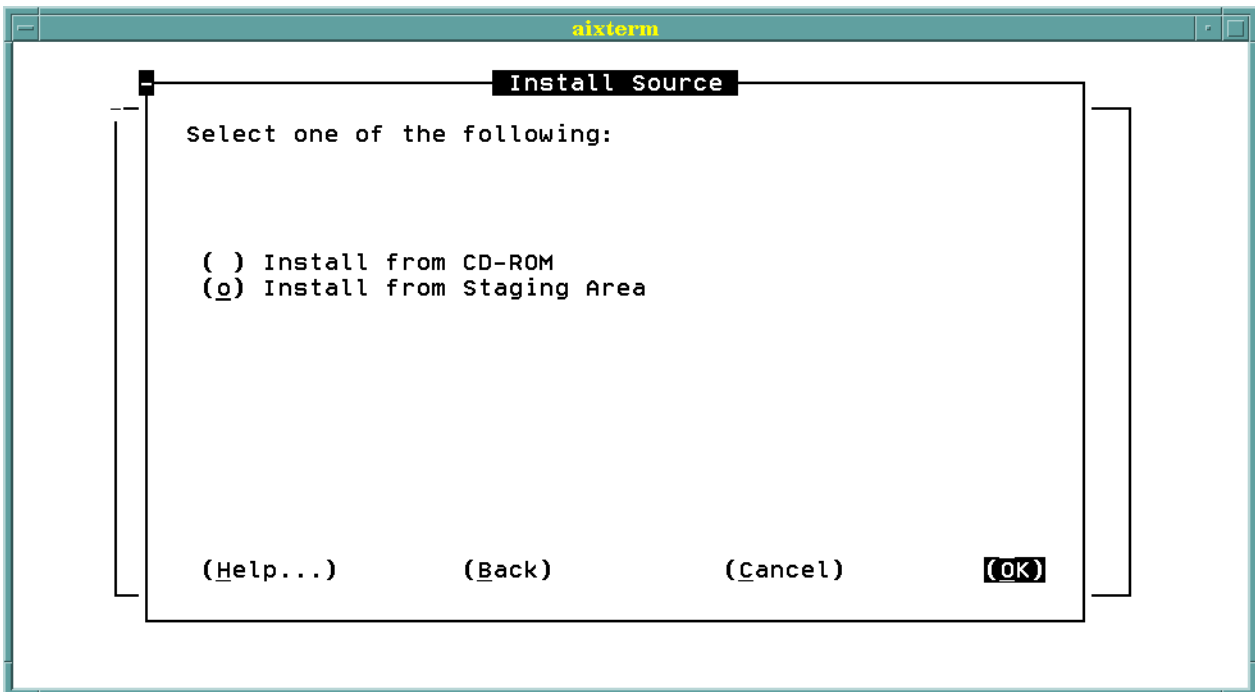


6.25 Install Products on All Nodes?

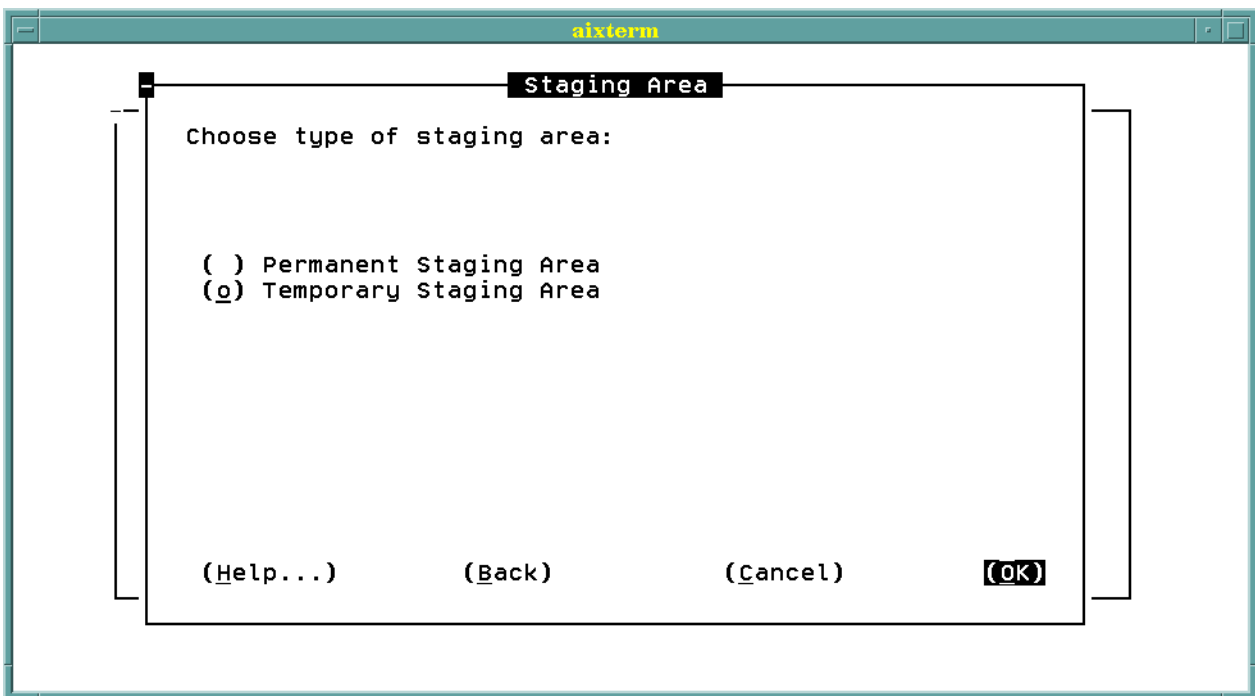
The next prompt asks if you want to install the products on all the nodes that comprise the cluster. Select **No**, because we only want to install the products on this node.

6.26 Installation Source

The next two prompts ask for the install source. We will be installing from a temporary staging area.

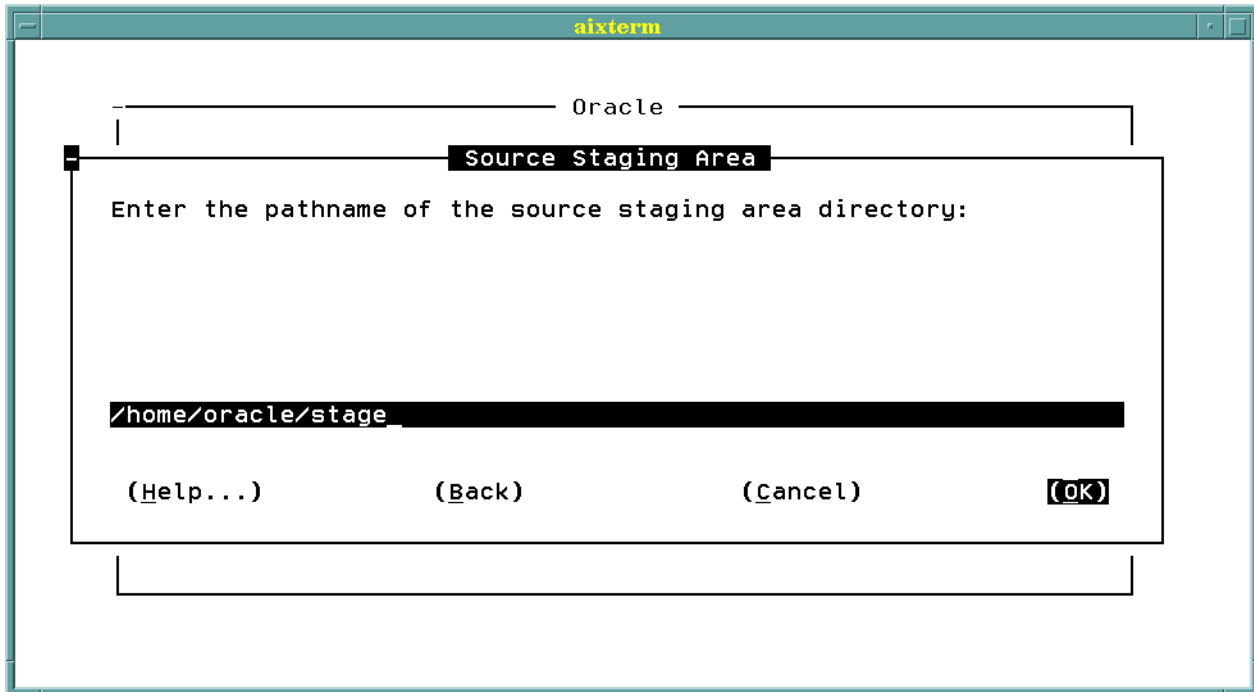


Note: Remember to delete the staging area when you have finished.



6.27 Source Staging Area

Enter the pathname of the source staging area. The staging area is located in `/home/oracle/stage`. This area will contain links to the CD-ROM.



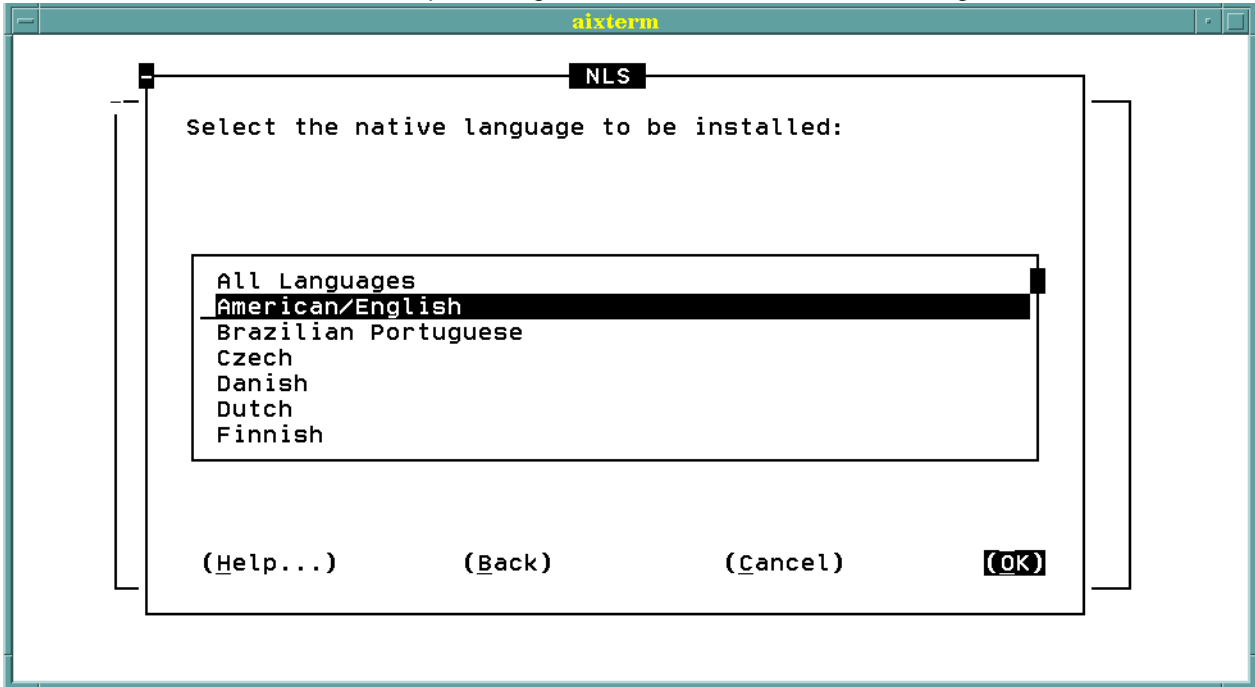
6.28 ORACLE_SID Variable

The next screen asks for your `ORACLE_SID` variable, which is the identifier for the database. This should be automatically read out of your environment (`/home/oracle/.profile`) area. For node 1, the SID is D1; for node 2 in the Mutual Takeover configuration, the SID is D2.

6.29 Specify National Language Support (NLS)

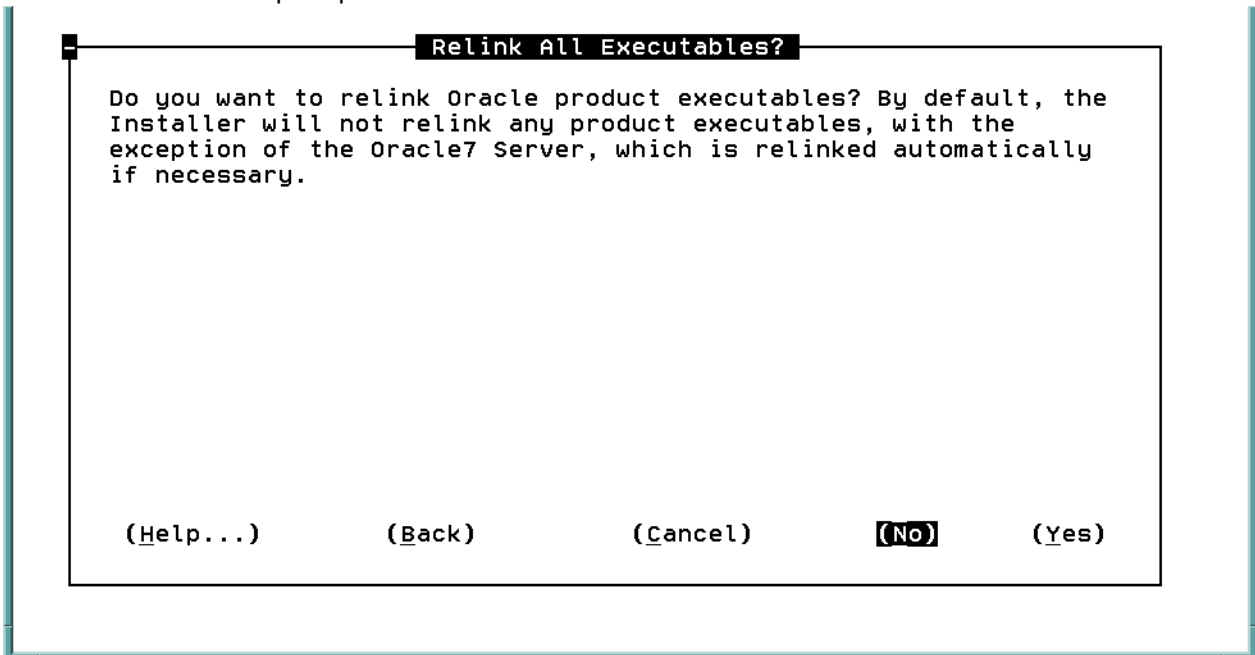
The Installer allows you to specify a language in which to receive screen messages for those utilities that support NLS.

For this sample configuration, the value is American/English.



6.30 Relink Executables?

The Installer offers you the option to relink the executables with the following prompt:

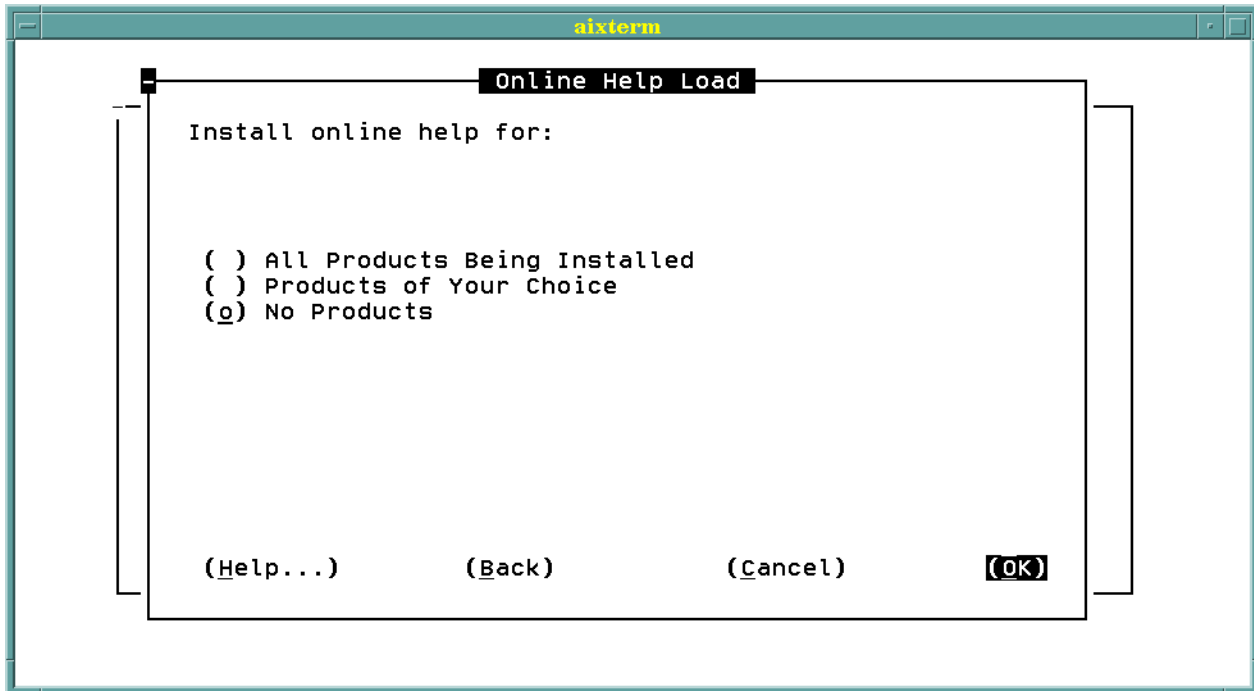


6.31 Location of Postinstallation File

The next screen tells you the location of the postinstallation file. This is always, `/home/oracle/app/oracle/product/7.3.2/orainst/root.sh`. Enter **OK** to continue.

6.32 Online Help and Documentation

In the sample configuration, we chose not to install any online help or documentation. The next three screens ask for the level of help and documentation you wish to install. If you choose any, make sure you consult the disk sizing worksheets in the previous chapter.



The next screen is for UNIX Documentation; we chose **No**.

The third screen is the Product Documentation Library, which is held on an additional CD-ROM. As in the screen shown above, we selected **No Products**.

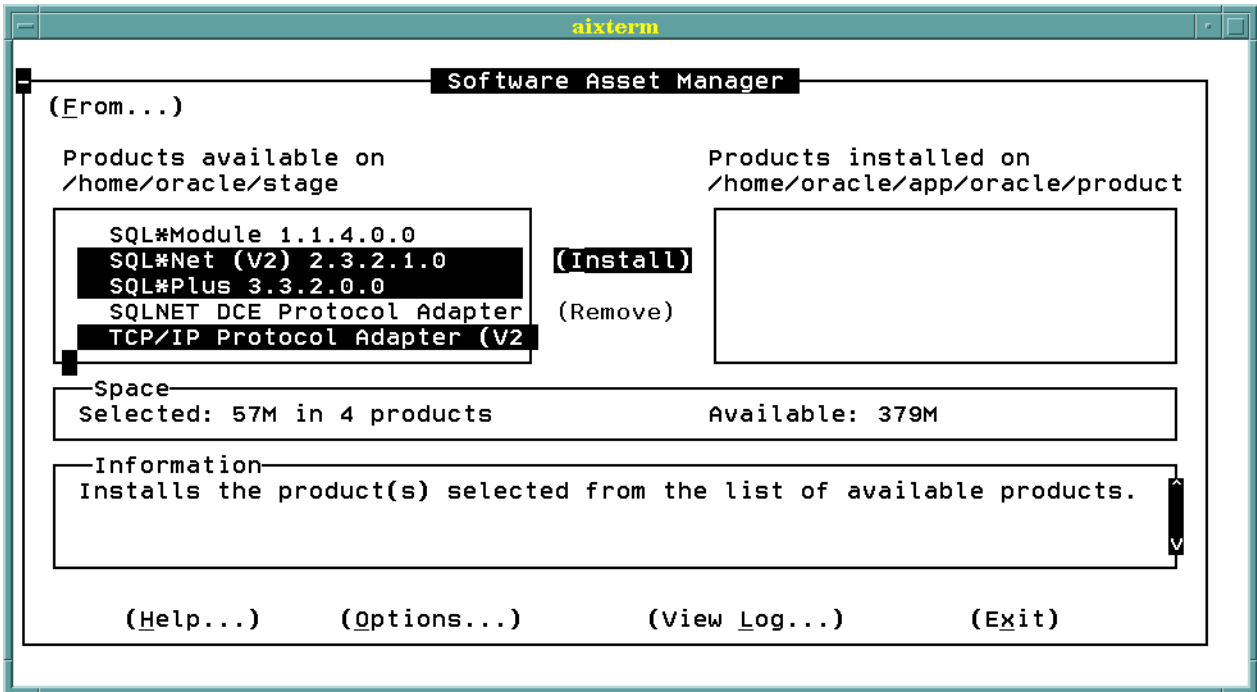
6.33 Products Menu

By now, you see that the Installer prompts you with the list of Oracle products that reside on your distribution media (tape or CD).

Use the arrow keys to scroll to, and the space bar to select, one or more Oracle products that you want to install.

After you have chosen the products, use the **Tab** key to navigate to the (Install) button, and press **Enter** or **Return** to activate the installation phase.

Note: For the cluster DAVE sample configuration, we only chose Oracle 7 Server (RDBMS) 7.3, SQL*Net V2, SQL*Plus 3.3, and the TCP/IP Protocol Adapter V2. (Not all shown in picture)



6.34 Official Hostname

You are prompted for the Official Hostname for the server. This refers to the label of the service adapter for the machine. Remember, that Rotating and Hot Standby configurations have the same database setup for node 1 and node 2; so both machines will have node 1's service adapter as an answer to this prompt.

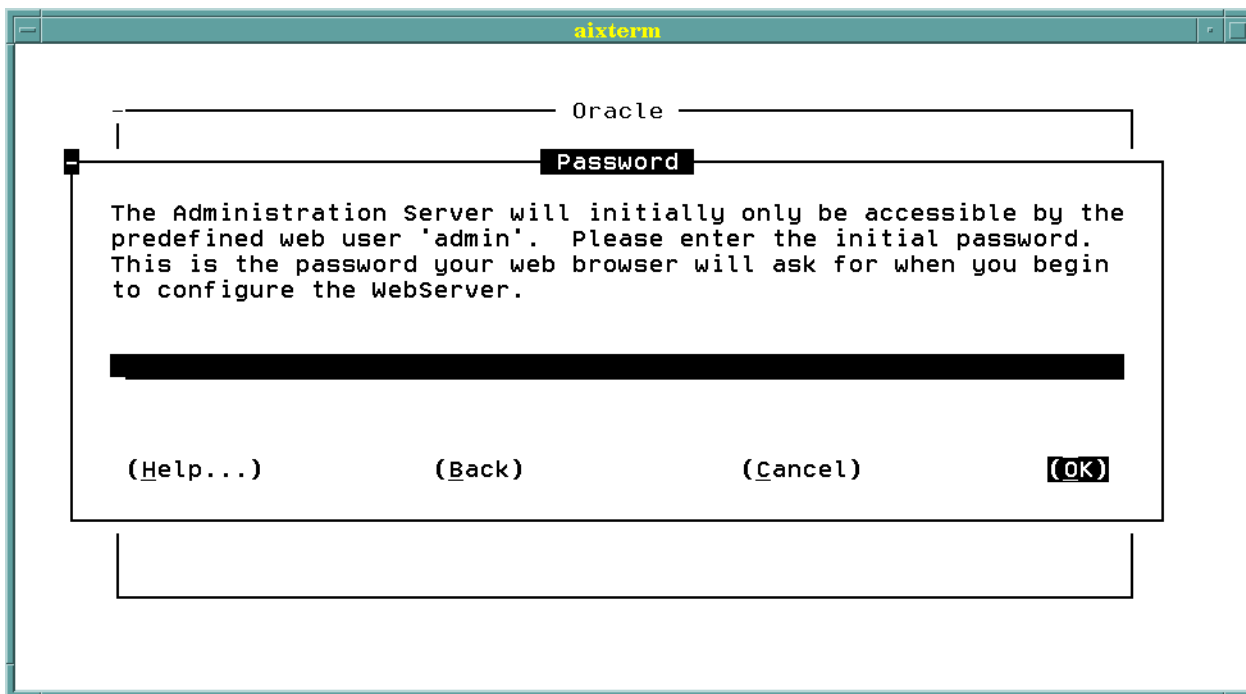


6.35 TCP Service Port

The TCP Service Port was entered into the `/etc/services` file earlier in the install. It is always set to 1521; so use the backspace to delete the **8888** and enter **1521**.

6.36 Passwords

Oracle 7 has many passwords, and you are asked to enter them all during this portion of the install. The following screen is the only one we will show you, but rest assured, there are many. Each is followed by a password-verification prompt. Make sure these are the same, or you will get an error.



6.37 Group Names

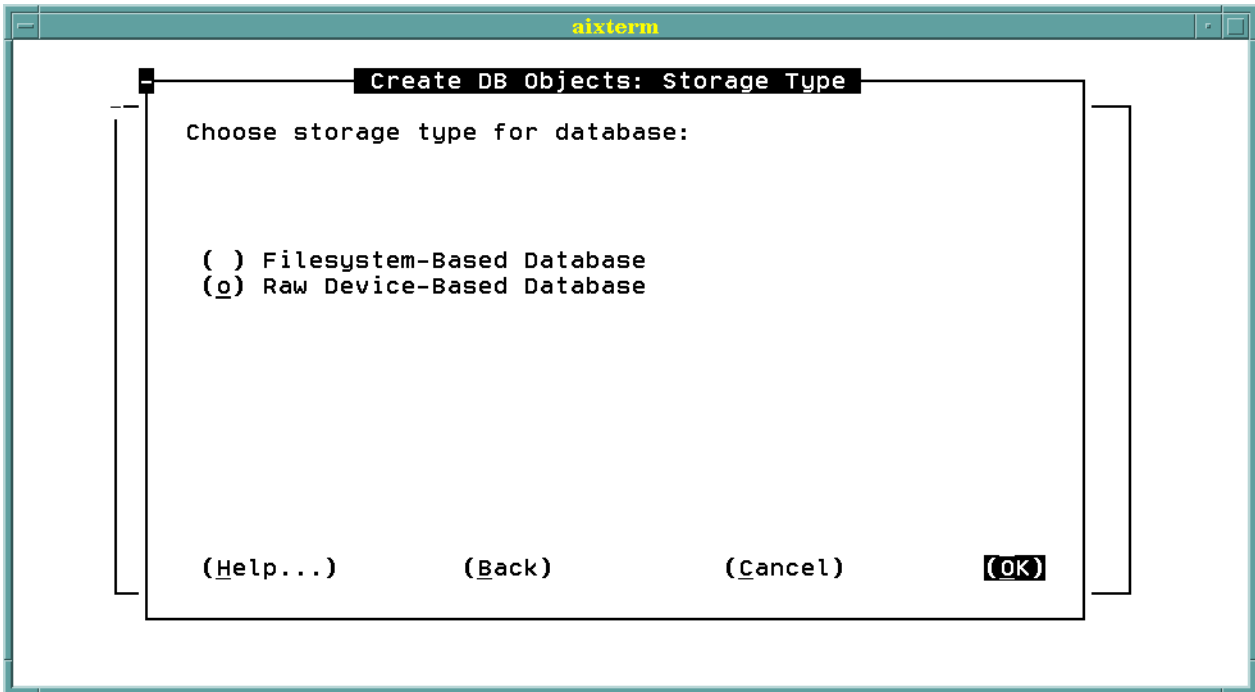
The next two prompts ask you for the Database Administrator Group and the OSOPER Group. For the first screen, select the radio button next to `dba`. For the second, just enter `dba`, because this was the group we created earlier for administrators.

6.38 More Than One Instance?

We do not want to allow more than one instance of the database; so we will enter **No**, to this next prompt to keep the maximum number of instances equal to one.

6.39 Choose Storage Type

The database storage is done on raw logical volumes.



6.40 Select the Character Set

You can choose a character set as listed from the Installer's prompt:

The default value is US7ASCII.

6.41 Enter the Password for the SYSTEM User and the SYS User

Note: The SYSTEM and SYS users are automatically created for you when you install Oracle Server. These users have dba privileges.

6.42 Do You Want To Set the Passwords for DBA and Operator?

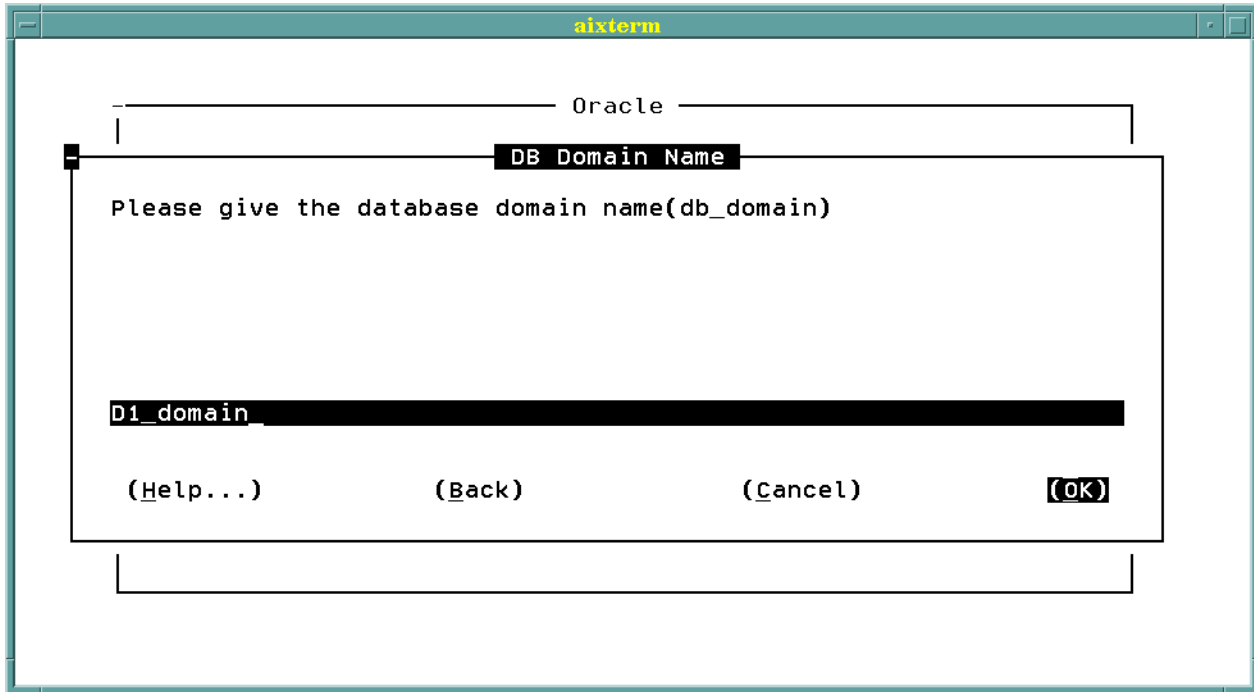
The default is **No**, but if you want to set passwords for these two internal users, allowing them to connect internally without passwords, then you must choose **YES**.

6.43 MTS Configured and the SQL*Net Listener Automatically Started?

The Multi-Threaded Server and the Listener should be on at all times, and therefore they should be configured at system startup. Select **Yes**.

6.44 Database Domain Name

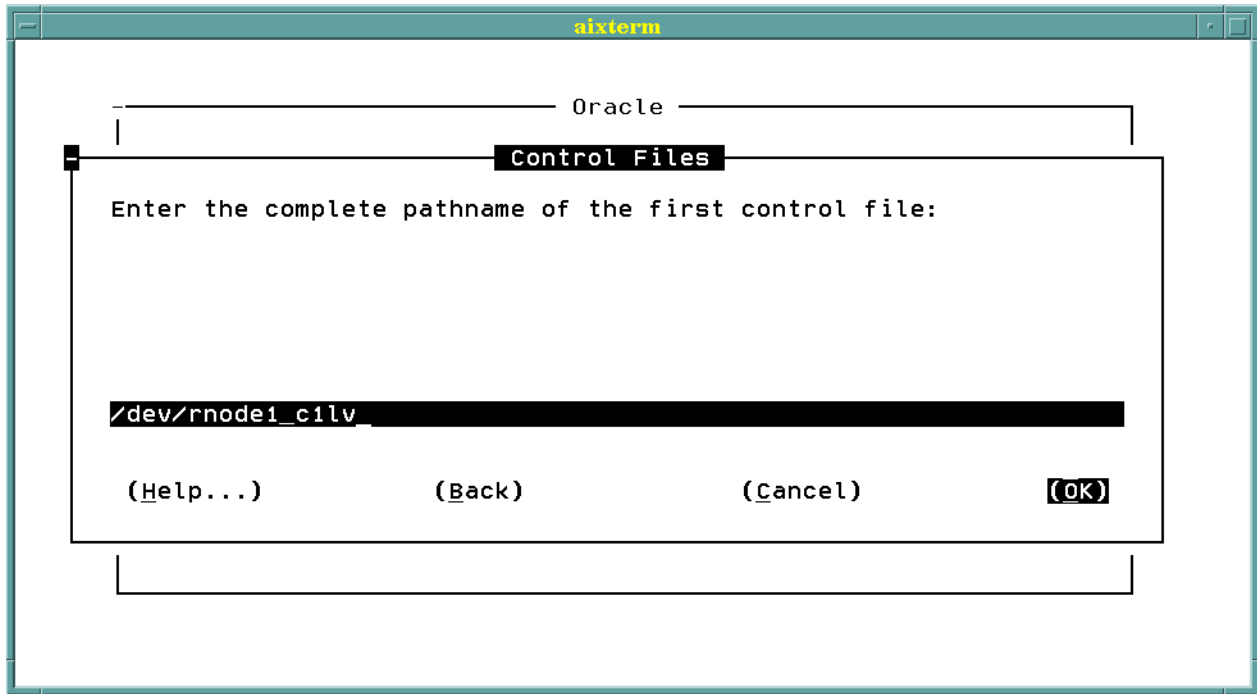
The DB domain name is just an identifier, but to keep it consistent, use the SID name (**D1** or **D2**), an underscore and the word **domain**. For node 2 on a Mutual Takeover configuration, use **D2_domain**.



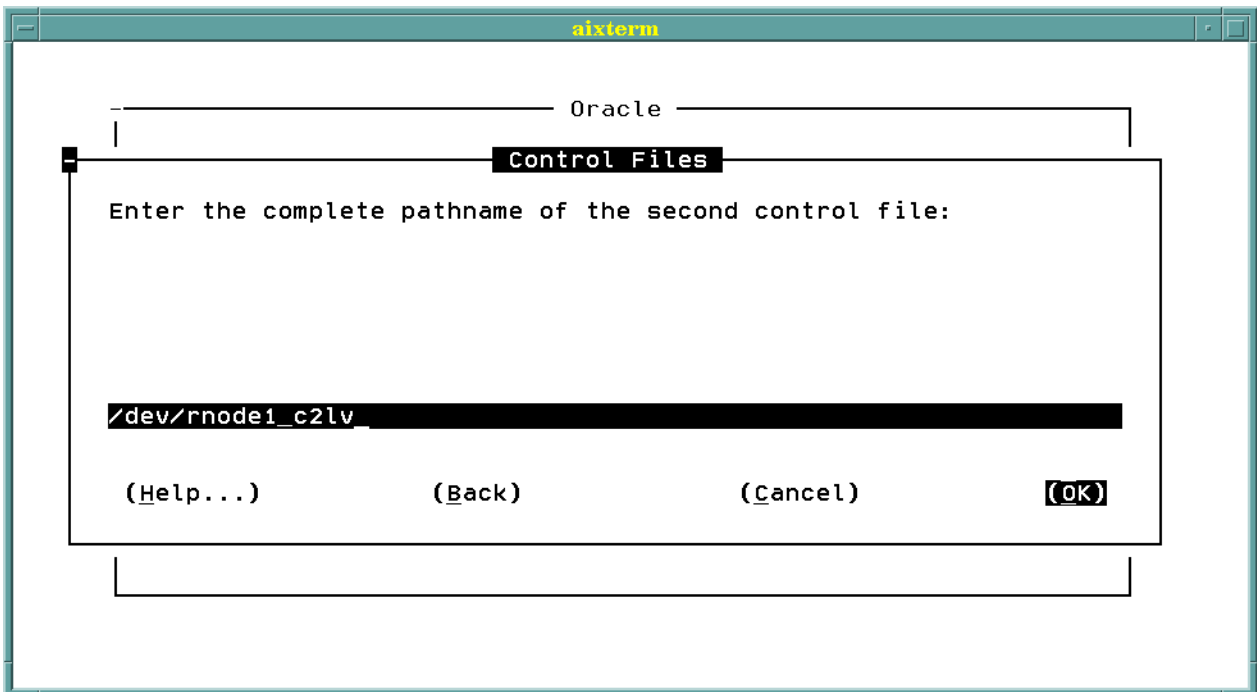
6.45 Name Control Files and Tablespaces

These prompts are asking for the names of the raw logical volumes that you created earlier. As stated before, this is where the database is stored. You should have these names on the Quick Install worksheet as well as on a table in the previous chapter.

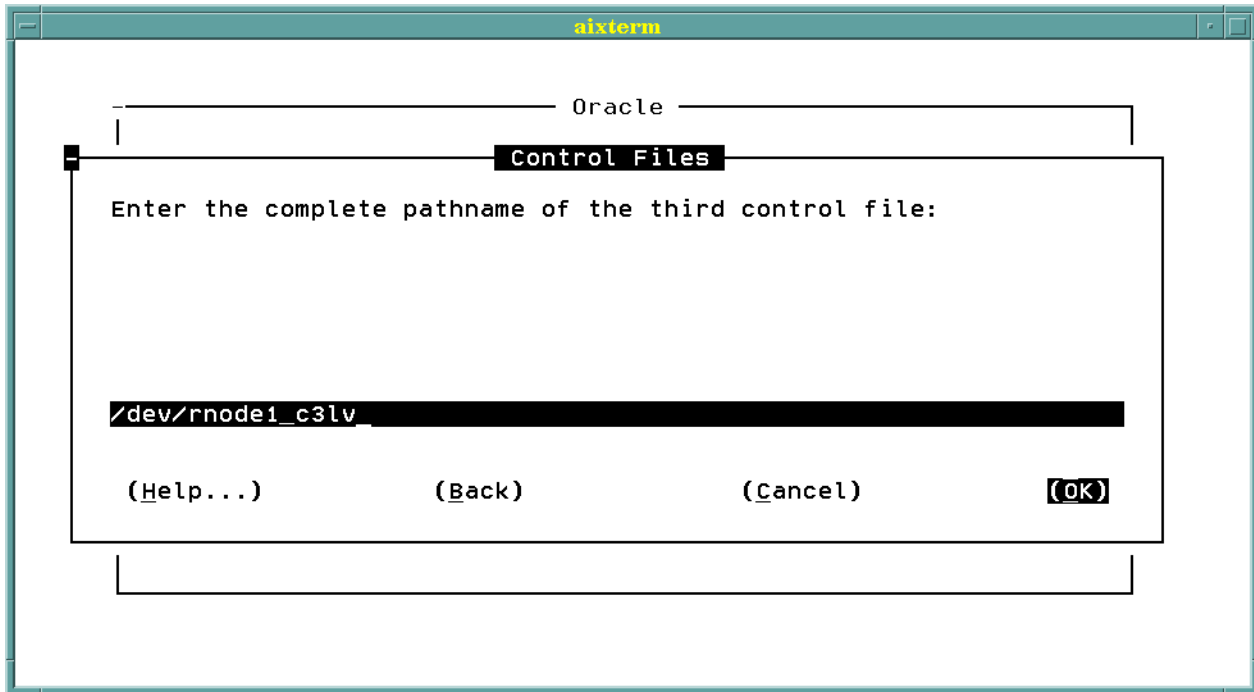
The first control file:



The second control file:



The third control file:



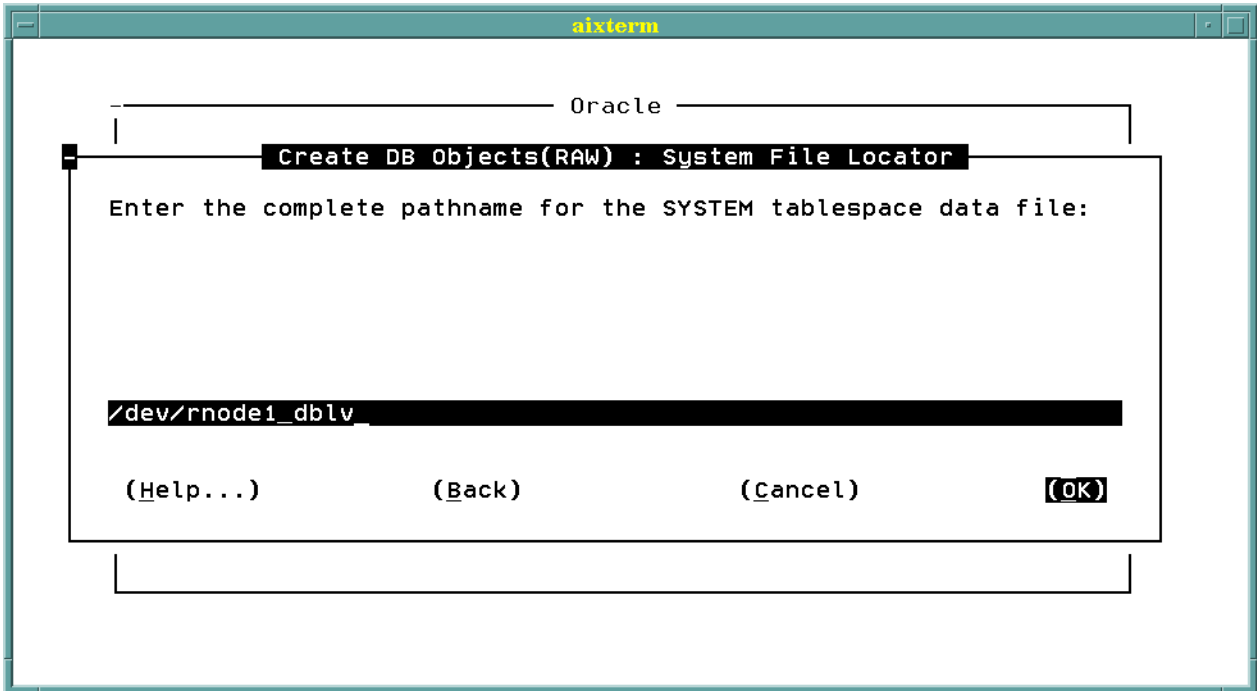
6.45.1 Name and Size Tablespaces and Logs:

The next few pages show the prompts with the answers highlighted in black. You will be shown the DB Defaults (two screens). These have the wrong names and sizes; so enter **OK** to get through both screens, and enter **No** at the last screen to go back and enter the correct names and sizes.

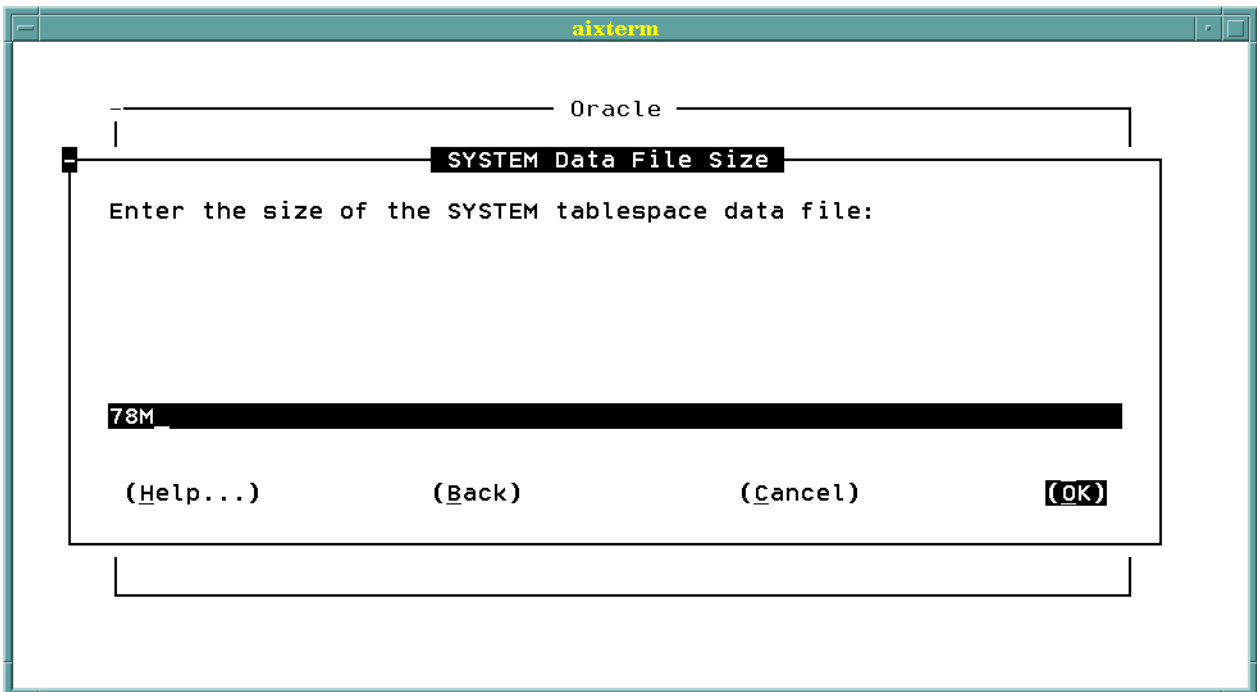
Note: Remember that the sizes entered will be one MB less than the actual size of the logical volume created earlier. This is to prevent any disk space problems later.

Here we go:

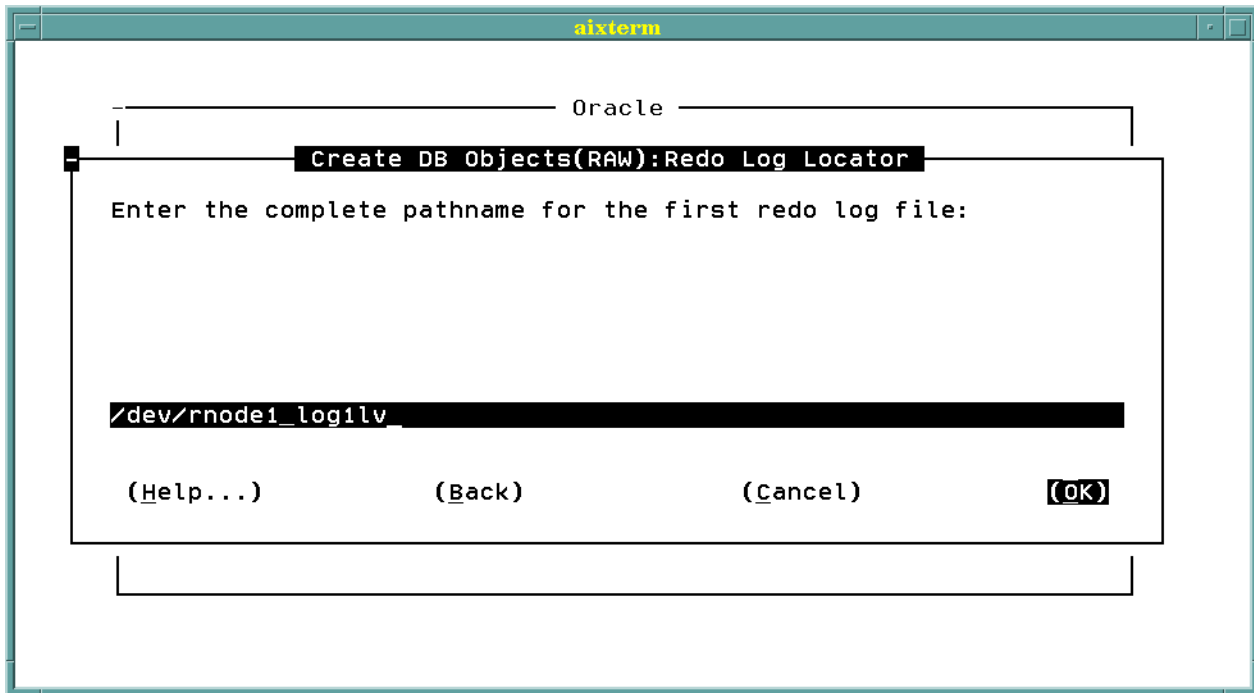
System Tablespace Data File:



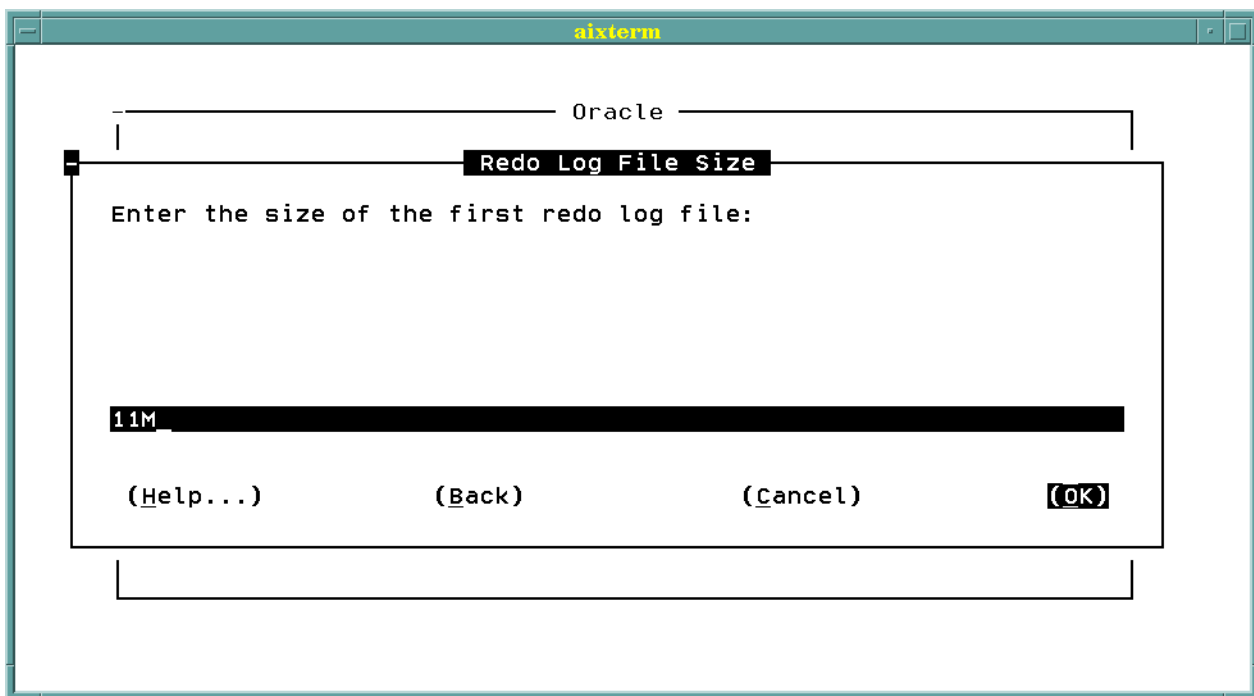
Size for System Tablespace:



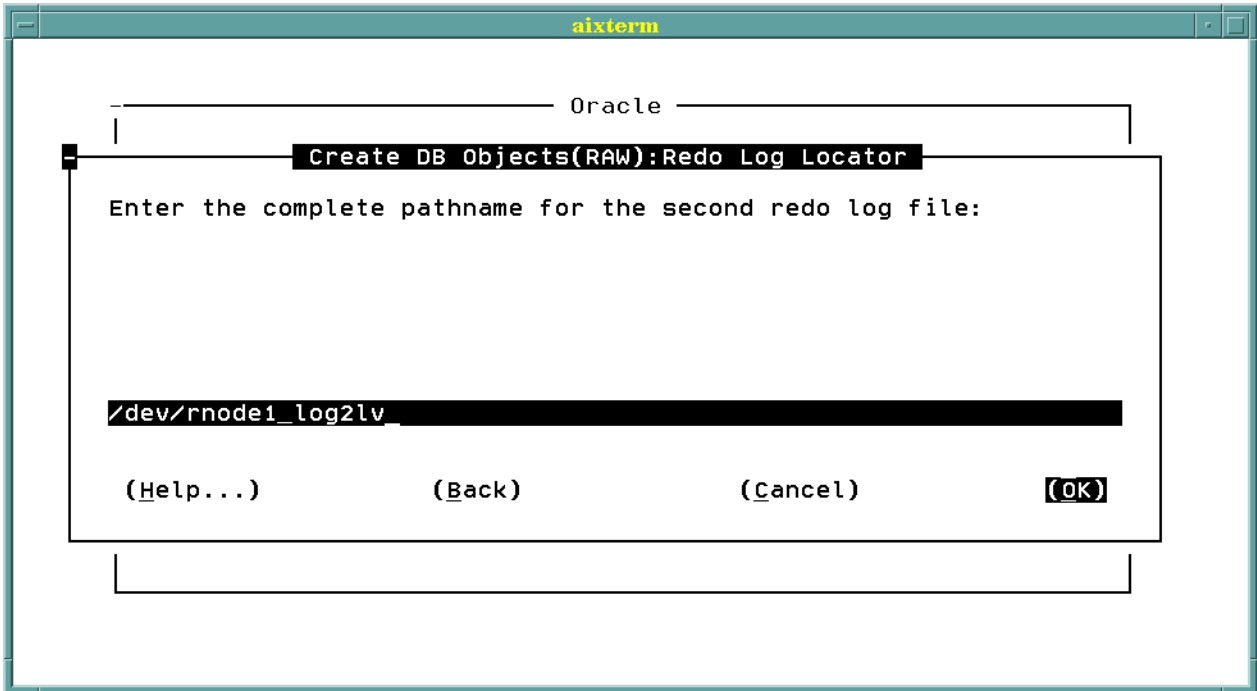
First Redo Log:



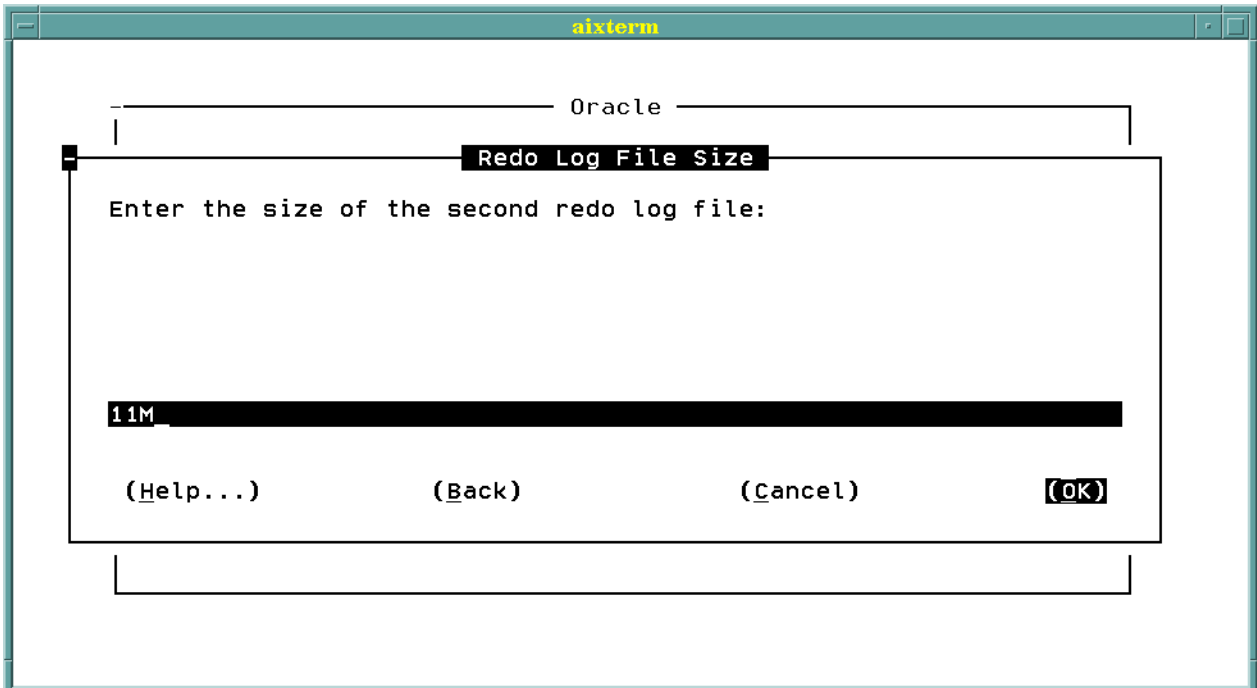
Size of First Redo Log:



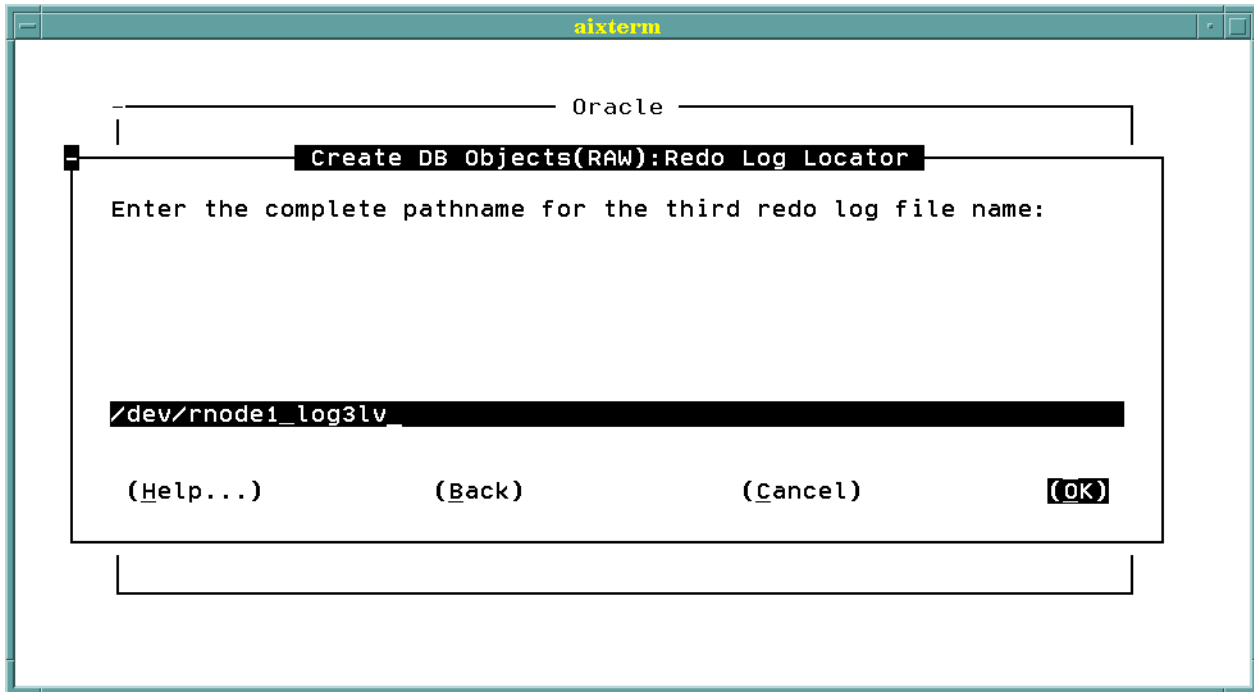
Second Redo Log:



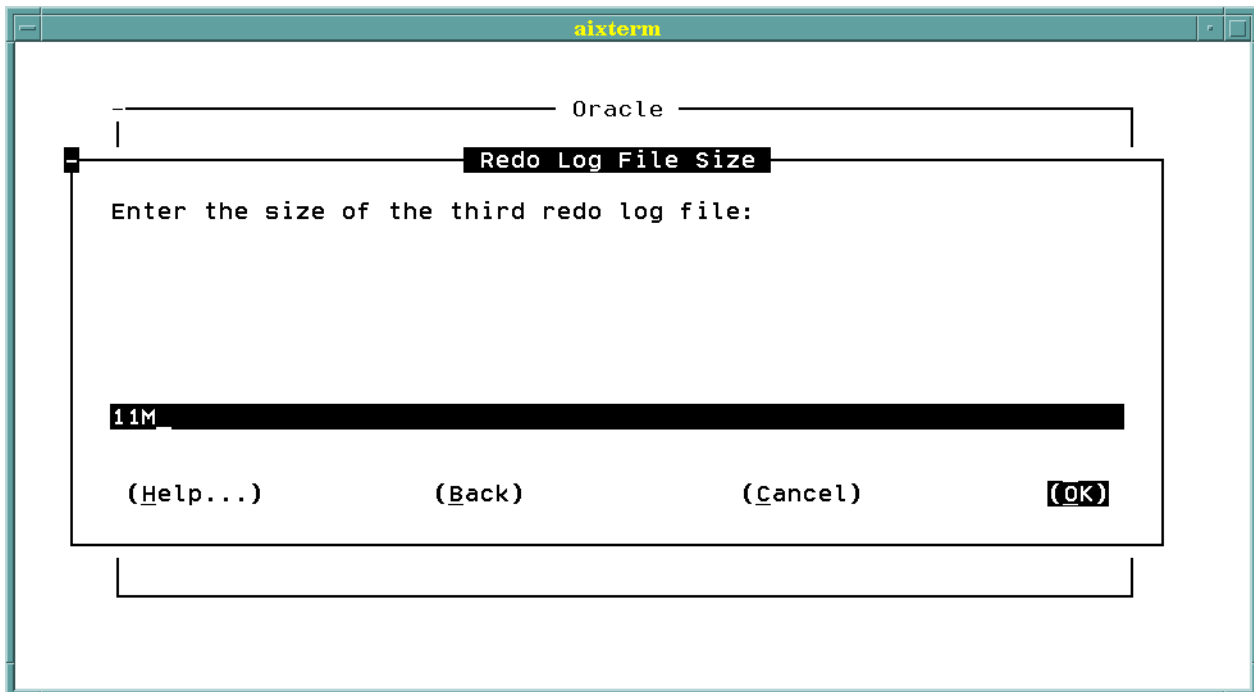
Size of Second Redo Log:



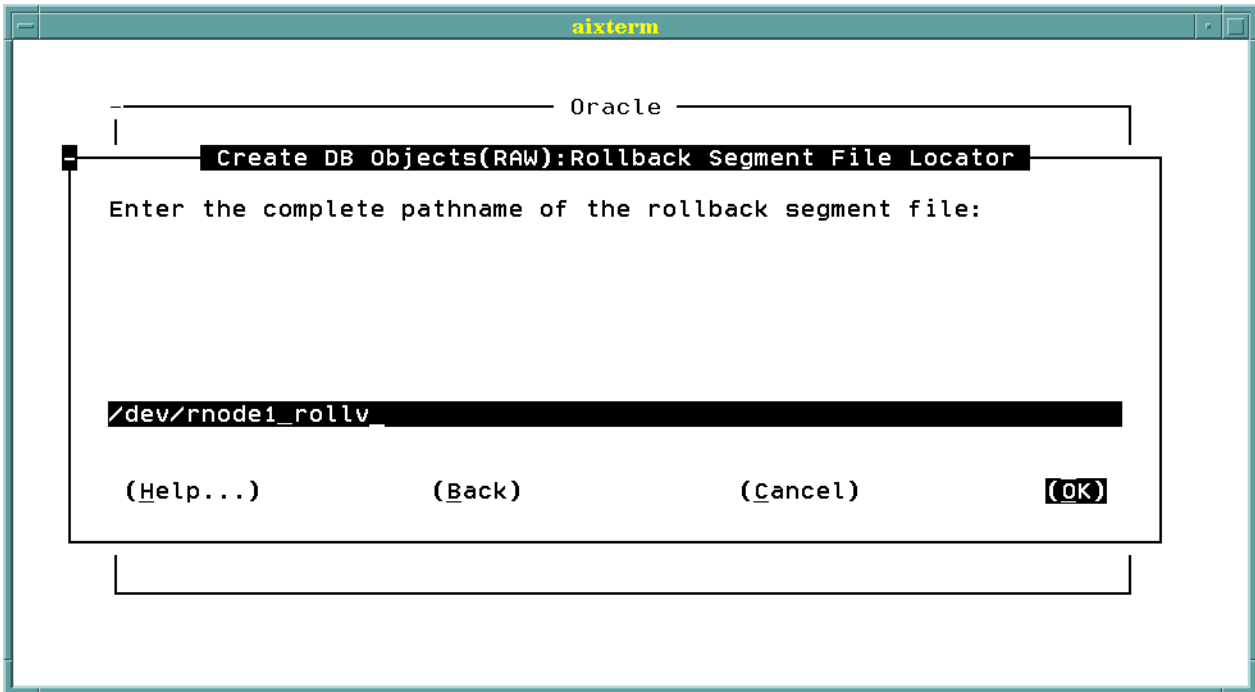
Third Redo Log:



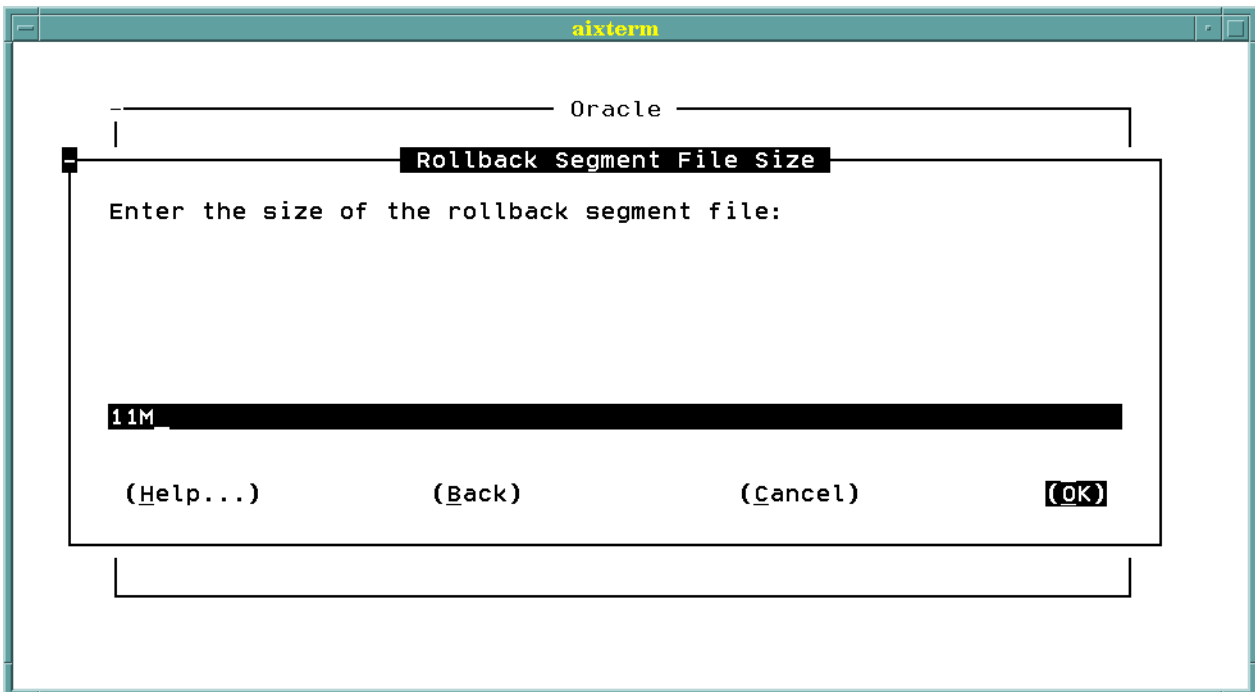
Size of Third Redo Log:



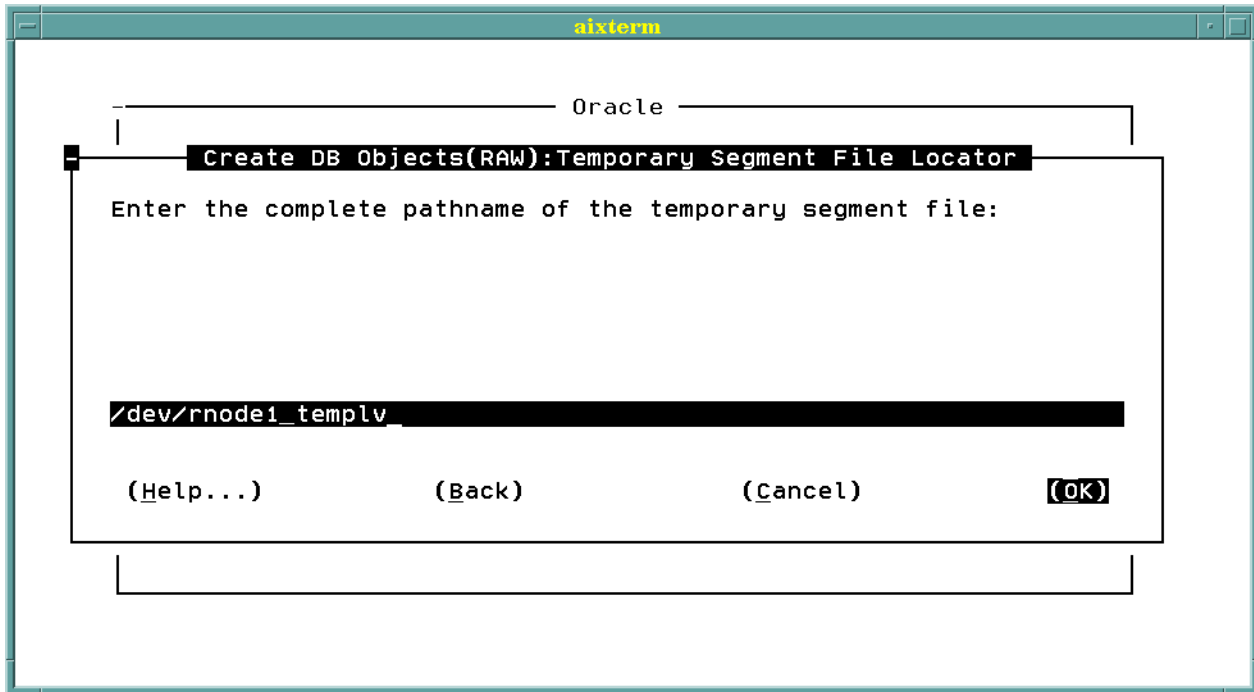
Rollback Segment File:



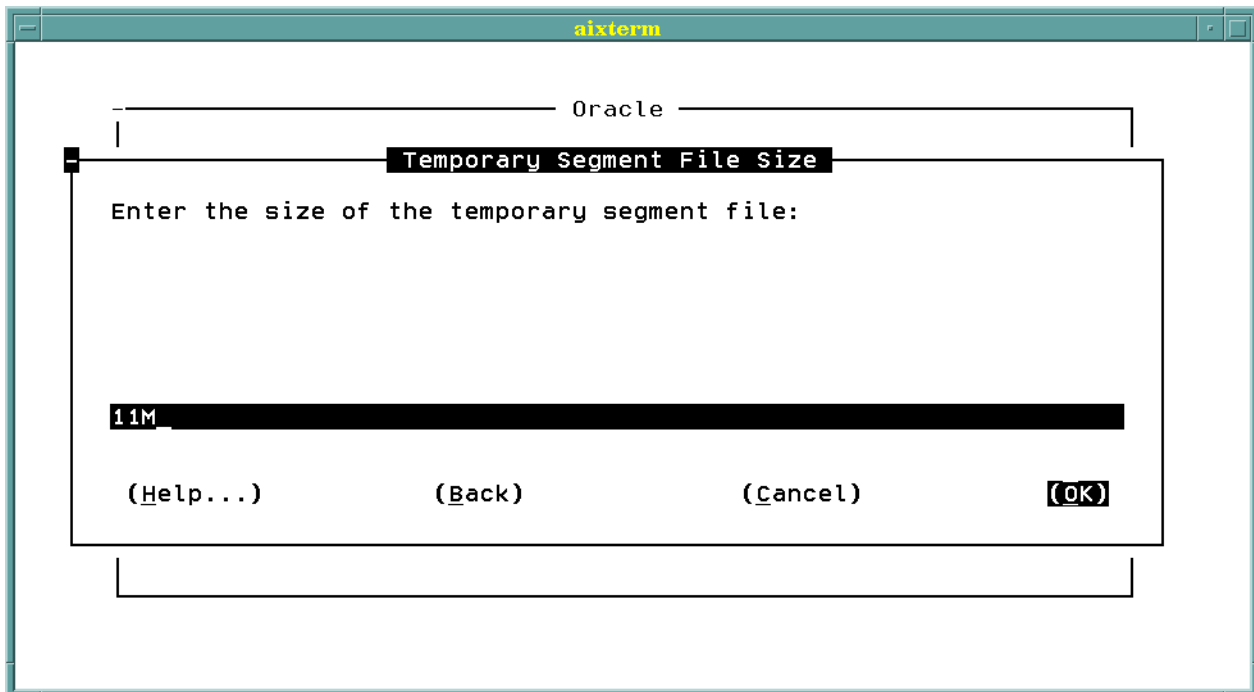
Size of Rollback Segment File:



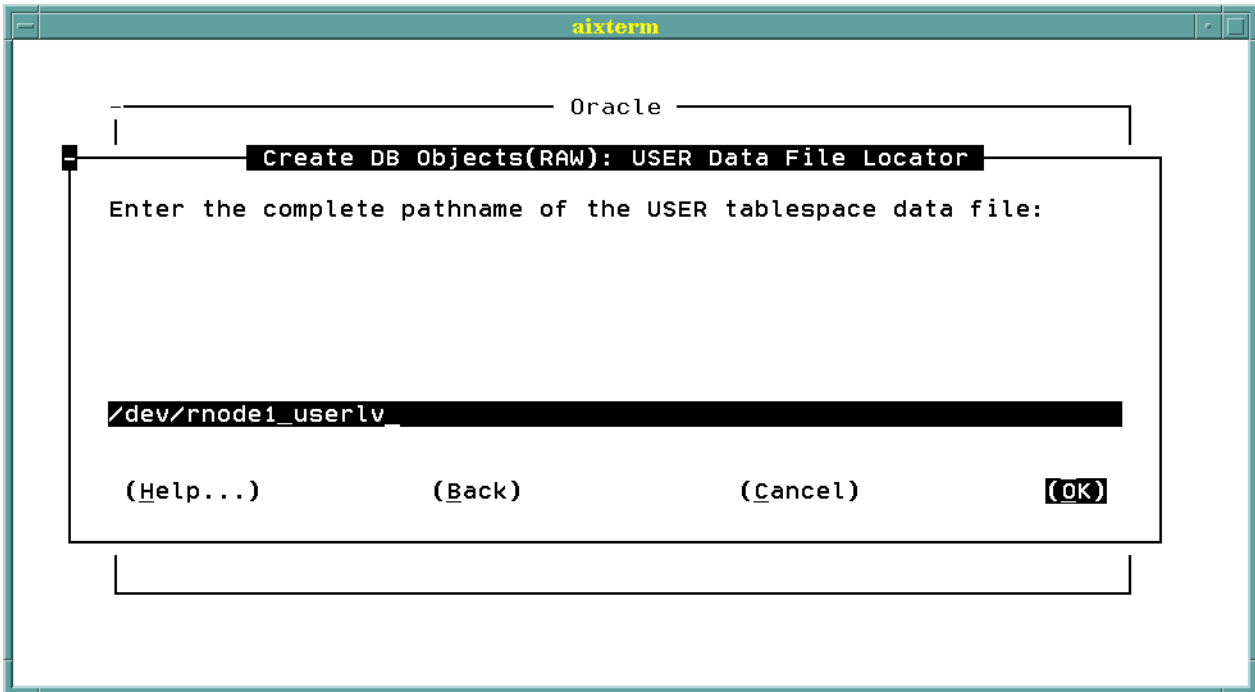
Temporary Segment File:



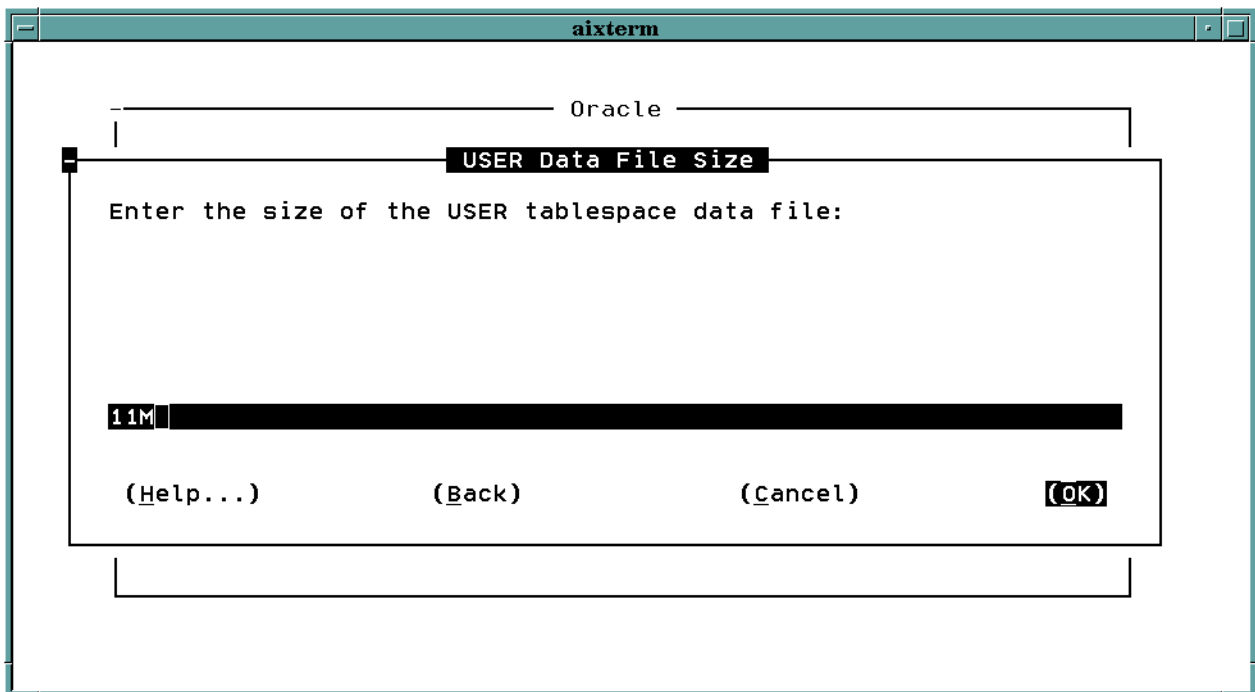
Size of Temporary Segment File:



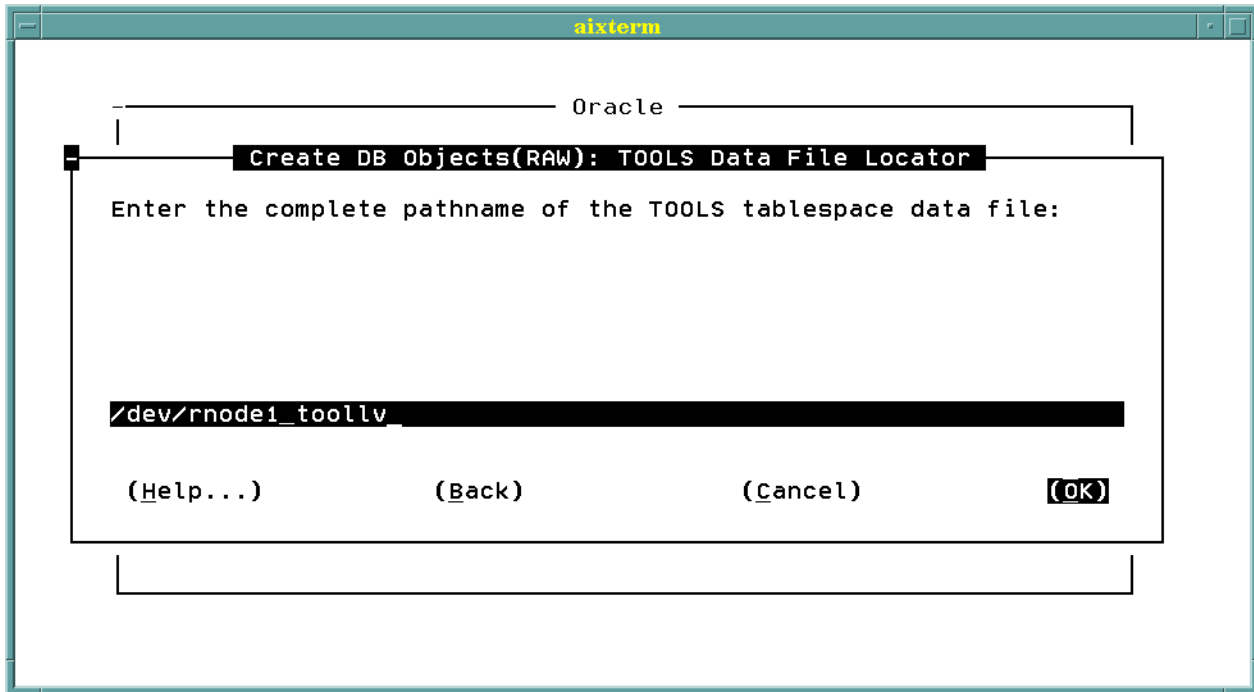
User Tablespace Data File:



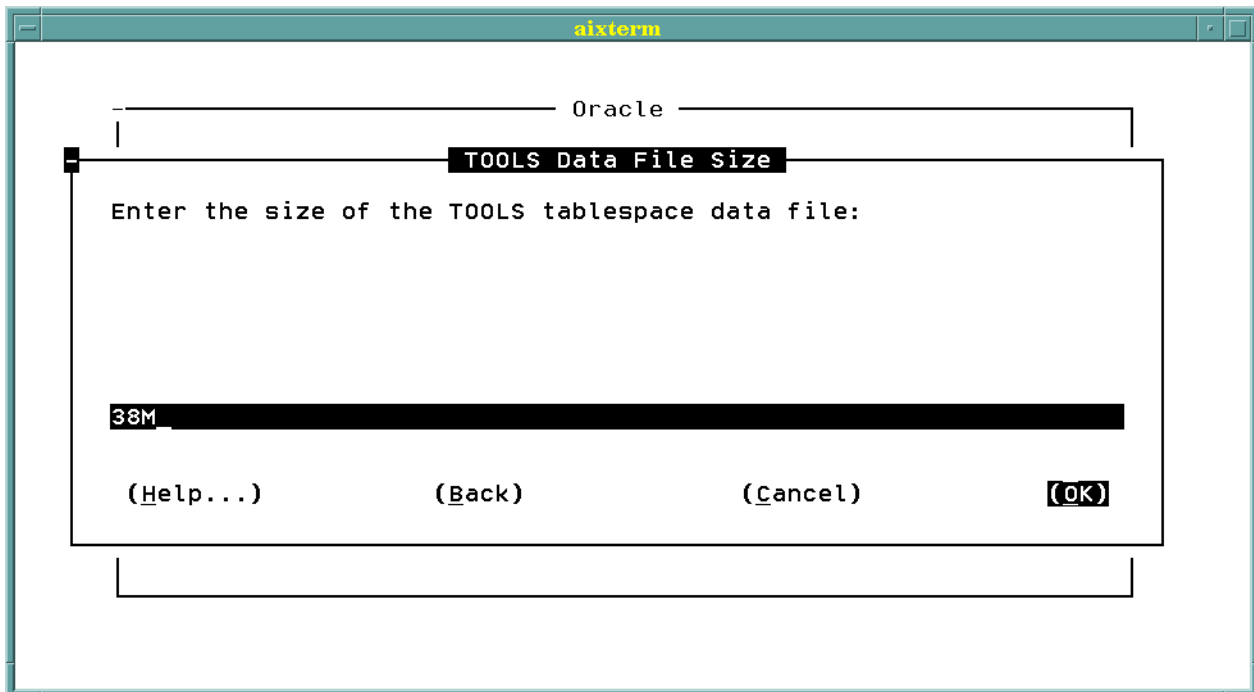
Size of User Tablespace:



Tools Tablespace Data File:



Size of Tools Tablespace:



You are now at the DB Defaults screens again. This time, however, they contain the data you just entered. Check them carefully. Make sure there are no mistakes.

Enter **OK** to get through both screens; then enter **Yes** at the last screen to continue on.

6.46 Optional Product Documentation

If you choose optional products you want to install (for example, in this configuration, we chose SQL*Plus V2), you are prompted for these by the Installer. For complete information about the product installation prompts, see the *Oracle 7 Server for AIX-Based Installation and Configuration Guide*.

Select **No** in the SQL Plus V2 Demo Tables and Help Facility Screens.

The Installer will do the rest of the work. This portion may take up to an hour depending on your hardware. Near the end of the install, you will see two or three messages built into the Installer. Read these, and then press **OK** to continue. When it is done, you will need to tab to **Exit** to exit the Installer.

6.47 Post-Installation Program

The output of the root.sh script will be something like the following example:

```
Running ORACLE7 root.sh script...
The following environment variables are set as :
ORACLE_HOME= /home/oracle/app/oracle/product/7.3.2
ORACLE_SID= D1 (or D2)
Are these settings correct (Y/N)? [Y]:
```

These settings should be fine; so enter **Y** to continue.

```
Enter the full pathname of the local bin directory [/usr/lbin]:
```

This local directory was created earlier in the install process. Just press the **Enter** key since the default is the correct directory.

You should get a warning message about the \$ORACLE_HOME directory not being consistent. This message is expected. When you are prompted, just enter **Yes** to continue.

Depending on which products you have selected to install, messages are displayed to alert you to the progress of the root.sh actions.

Attention: Because you log in as root, you may have to set your ORACLE_HOME, ORACLE_SID, and ORACLE_OWNER variables by issuing the following commands on the command line according to your environment. Below are the commands we issued for the sample configuration:

```
# ORACLE_HOME= /home/oracle/app/oracle/product/7.3.2 ; export ORACLE_HOME
# ORACLE_SID= D1 ; export
```

Note: D1 is the ORACLE_SID for node 1, and D2 is the ORACLE_SID for node 2.

6.48 Copying and Linking Necessary Files

After the copying of necessary files, you are finished with the setup of the first node of the cluster.

You must now shut down the machine. At the prompt enter:

```
# shutdown -Fr
```

**YOU HAVE FINISHED WITH THE QUICK INSTALL SCRIPT FOR NODE 1.
PLEASE RETURN TO THE TOP AND REPEAT FOR NODE 2. THERE IS LESS
USER INPUT IN NODE 2, AND THUS A FASTER INSTALL.**

Chapter 7. Post-Installation of the Oracle Server Database

The Quick Install Program takes you through the post-installation tasks, updates the control files with the appropriate variables, and moves them to the required directories. When you are done with the Quick Install, you are ready to test the database. However, this is an outline of the files that have been manipulated from their defaults by the Quick Install Program.

7.1 Verify Database Security and Integrity

To maintain integrity of the Oracle code, all executables in the \$ORACLE_HOME/bin directory must be writable by the Oracle owner only.

You must verify that these files have file modes 755 or 777. Also, to maintain discretionary access to data, all databases, redo log and control files that may be put in the external disks must be readable by the Oracle owner and group only, including the root user for highly available purposes.

7.2 Create Oracle Server User Logins

For each Oracle user, create a login with the following properties:

Login name	The user's user name
User ID number	A number between 3 and 32767
Primary Group	A number between 3 and 32767, not the dba group
Home directory	/home/oracle
Login shell	/bin/ksh

7.3 Check your Filesystems

The filesystems you created on the external disks (node1_scripts in our example) must contain the necessary control scripts and data files. Below is a copy of the files stored in the node1_scripts directory by the Quick Install Program for the Mutual Takeover configuration and for the Rotating/Hot Standby configurations:

Mutual Takeover Configuration:

```
aixterm
# ls
get_HACMP_service  ora_stop          oratabD2          stop_D1
listener.ora       oracle.vars       start_D1          stop_D2
node.vars          oratabD1          start_D1.LNODE    stop_appli
ora_start          oratabD12        start_D2          tnsnames.ora
# _
```

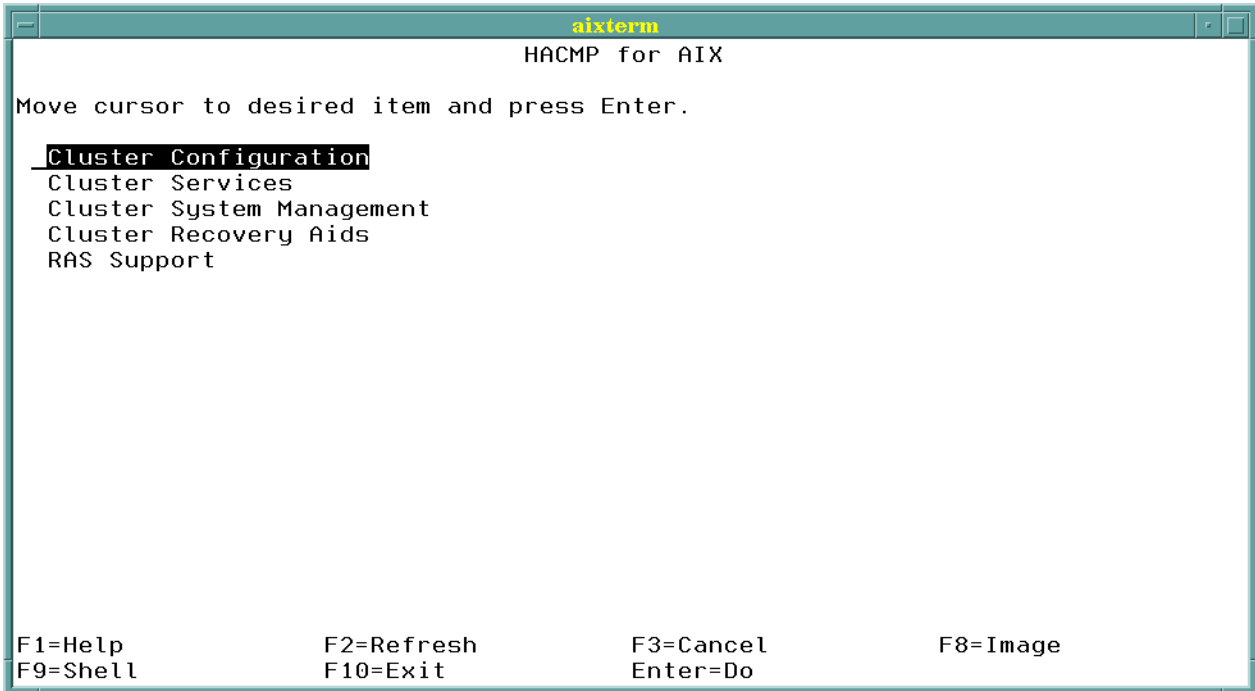
Rotating/Hot Standby Configuration:

```
aixterm
# ls
get_HACMP_service* node.vars         oracle.vars       stop_D1*
iffer              ora_start*       oratabD1          stop_appli
listener.ora       ora_stop*        start_D1*         tnsnames.ora
# _
```

If any of these files are missing, please copy them from the `cascc_scripts` (for Mutual Takeover) or `rot_scripts` (for Rotating/Hot Standby) directory on the accompanying disk. The content of these files will be checked later in this chapter.

7.4 Testing Cluster Synchronization

Testing the HACMP portion of the install can be done with HACMP's cluster verification utility. To get to this utility, enter `smit hacmp` from the command prompt.



Select **Cluster Configuration**. From the next menu, select **Cluster Verification**. At the prompt, just press **Enter** until the process begins. Run this tool on both nodes of the cluster.

7.5 HACMP Cluster Services

When are ready to begin failover testing, you must first start the HACMP services on both nodes. One node at a time, from the command prompt, enter `smit hacmp`. Select **Cluster Services**. The next menu will allow you to either start, stop, or check on the cluster services.

After selecting **Start Cluster Services**, press **Enter** until the process begins. When the process is finished, and you get the **OK** prompt, you can exit out of SMIT by pressing the **F10** key.

To see the progress of the cluster, from the command prompt, type:

```
# tail -f /tmp/hacmp.out
```

This will monitor the progress of the startup. When you want exit out of this, press **Ctrl-C**.

When you are done, bring up the second node in the same way. To check if the database is running, type:

```
# ps -eaf |grep ora
```

from the command prompt. You should see something similar to Figure 11.

```

oracle 7374      1  0  Oct 13      -  0:34 ora_pmon_D1
oracle 9434      1  0  Oct 13      -  0:18 ora_smon_D1
oracle 10206     1  0  Oct 13      -  0:00 ora_s000_D1
oracle 10670     1  0  Oct 13      -  0:00 /home/oracle/app/oracle/product/7
3.2/bin/tnslsnr lisD1 -inherit
  root 10866 15204  0 14:39:13 pts/2  0:00 grep ora
oracle 11478     1  0  Oct 13      -  0:33 ora_lgwr_D1
oracle 12258     1  0  Oct 13      -  0:08 ora_d000_D1
oracle 12754     1  0  Oct 13      -  0:51 ora_dbwr_D1

```

Figure 11. Oracle Processes Running on Node with Database

The Listener process is lisD1 (PID 10670).

To test a failover, simply run the Stop Cluster Services SMIT screen and select the **graceful with takeover** option, or power-off the machine running the database. With Rotating or Hot Standby, this will be one of the nodes in the cluster. With Mutual Takeover, you can power-off either machine because each is serving its own database. On the machine in the cluster that is still running, execute the following command to watch the takeover activity:

```
tail -f /tmp/hacmp.out
```

Chapter 8. Common Errors, Problems and Recommendations

This chapter discusses common Oracle error messages that you might encounter and problems you might have when you do the installation and setup. Also, we will give some recommendations concerning the highly available database.

All of the discussions here are based on our experiences while preparing and doing all the tasks described in previous chapters. Therefore, this chapter is not intended to explain all errors and problems in Oracle. You might find errors or problems that differ from this chapter when you do your installation. However, if you run into some of the same problems that we did, perhaps this chapter will be helpful to you.

8.1 Common Errors and Problems

This section describes several common errors that occur during the installation and during testing of database access from the client to the server.

For complete information about error messages, please refer to the *Oracle 7 Server Messages and Codes Manual* and to the *SQL*Net Administration Guide*. Another good reference is the Appendix B discussion about “*Troubleshooting Your Installer Session*” in the *Oracle 7 for AIX-Based Installation and Configuration Guide* for handling problems with the Installer.

8.1.1 Errors and Problems During the Installation Process

When you want to install the Oracle products, always remember the three stages of the process: preinstallation, installation, and postinstallation.

Failure to complete one of these three tasks will end in error, and you will not be able to run your database server and other Oracle products.

8.1.1.1 Installation-Stage Common Errors and Problems

This table shows the common errors that might happen while installing and helps you to identify the possible reasons for the problem.

Table 18. Common Errors During the Installation Stage

Stage of Installation	Common Error
Installing the Installer	System requirement does not meet environment setup
Initial Installer prompts	Environment setup
Loading the product file	Not enough space in staging area File permissions problem
Installing the product files	Not enough space in \$ORACLE_HOME System requirement not met File permissions problem
Relinking product executables	Not enough space System requirement not met
Creating database objects	System configuration (such as shared memory) Database access problems

For example, you may see the following error message in your Installer menu while creating the database objects:

```
ORA-00210: cannot open control file /oradatafs2/ora_ctl/ctl1.ora
```

From this, you know there is something wrong with your control file. Maybe the full pathname you specified is not correct, or your directory is inaccessible. You need to look at your directory permissions.

For complete information about error messages (cause and action), please refer to the *Oracle 7 Server Messages and Codes Manual*.

8.1.1.2 Postinstallation-Stage Common Errors and Problems

Remember that every time you finish an installation, you must do the postinstallation tasks. If you forget to complete these tasks, an error situation will result, and you won't be able to run your installed products.

You need to run `root.sh` after you finish the installation.

```
# ./root.sh
```

You might find that the program always says:

```
Unmatched variables for ORACLE_SID or ORACLE_HOME.
```

If this occurs, you have to do the following:

```
# ORACLE_SID= <your oracle_sid> ; export ORACLE_SID
# ORACLE_HOME= <your oracle_home> ; export ORACLE_HOME
```

and as root, and run the `root.sh` again.

8.1.2 SQL*Net Errors and Problems

When you try to run the Listener on the server by typing:

```
$ lsnrctl start
```

you might receive this message:

```
TNS-01151: Missing listener name, lisD1 or lisD12, in LISTENER.ORA
```

The message above means that you do not have a Listener name in your Listener.ora file, or that you have a Listener name, but did not name it *LISTENER* or the default name. You have to check your Listener.ora file and see if you have specified your Listener name (see the example in the appropriate Appendix to your configuration). If you have, you should run the Listener by typing the Listener name after the `lsnrctl start` command.

The other common error message received when you run the Listener is shown below:

```
TNS-12154: TNS: could not resolve service name
```

This error points to a problem related to the `tnsnames.ora` file. The service name specified cannot be located in the `tnsnames.ora` file. If you get this error message, do the following:

- Verify that your `tnsnames.ora` exists and is in the proper directory. Please refer to the previous chapter sections that discuss postinstallation tasks on the server and postinstallation tasks on the client.

- Check to see that the service name is mapped to a `connect` descriptor in the `tnsnames.ora` file and add or correct it if necessary.
- Make sure there are no syntax errors in the `tnsnames.ora` file; check to see if it has unmatched parentheses or stray characters.

TNS-12203: TNS: unable to connect to destination

This error indicates that the client is not able to find the desired database. If you get this error message, do the following:

- Be sure you have correctly entered the service name you want to reach.
- Look at the `tnsnames.ora` file to see that the `ADDRESS` parameters in the connect descriptor for the service name are correct.
- Check and verify if the Listener on the remote node is up and listening. You can verify this by typing `lsnrctl status <listener_name>`:

TNS-12545: TNS: name lookup failure

This error occurs when the Listener on the remote node cannot be contacted. The `ADDRESS` in the `tnsnames.ora` file or the `Listener.ora` file may be incorrect. This message also appears if the Listener on the remote node has not been started.

Note: The error messages in SQL*Net or client-to-server connection errors are usually caused by errors in the `tnsnames.ora` or `Listener.ora` files. Therefore, these files must be rechecked when you receive any of these errors.

For more information about all SQL*Net error messages, please refer to the *Oracle Network Products Message Manual*.

8.2 Recommendations

We recommend strongly that you put the control files, redo log files, and all the database object files, such as system tablespace, rollback tablespace, temporary tablespace, and user tablespace, in the shared external disk in the HACMP environment.

The reason is that the location of these files is one of the critical points that determines whether or not your database is highly available.

We also recommend that before you do all the installation tasks, carefully plan your database and prepare your system. For example, make sure your cluster functions properly without the database application. Carefully planning your database means planning all the names and variables for your database, the size of your database, and making sure you have the prerequisite hardware and software. For parallel databases, you also have to create all the raw logical volumes needed for all of the database object files prior to the installation process.

Appendix A. Quick Installation Worksheets

This Appendix provides worksheets to fill out in preparation for running the Quick Install Program.

A.1 Mutual Takeover Worksheet

This worksheet is to help you in gathering the necessary information to smoothly complete the install of the Oracle 7 database on a highly available cluster using the Quick Install Program. This worksheet is for the Mutual Takeover configuration. Please write the information in the areas provided, and keep it with you during the installation and setup process.

Installation Information:

The following tables hold the adapter information for the cluster. In the Mutual Takeover configuration, each cluster has a service, standby and boot adapter. Please fill in the labels and IP addresses for each adapter. As an example, the boot adapter for node 1 in our sample configuration would have a label of `node1_boot` and an IP address of `10.0.1.11`.

Table 19. Node1 Adapters for Mutual Takeover Configuration

Node1 Adapters	Adapter Label	Adapter Address
Boot Adapter		
Service Adapter		
Standby Adapter		

Table 20. Node2 Adapters for Mutual Takeover Configuration

Node2 Adapters	Adapter Label	Adapter Address
Boot Adapter		
Service Adapter		
Standby Adapter		

You need to find an available Group ID and User ID for both node1 and node2. These ID #'s must be the same on both machines. The defaults used in our sample configuration are Group #201, and User #210. To find this information, enter `smit group` or `smit user`, and choose the **list** option.

Available Group ID # _____

Available User ID # _____

Pick a terminal emulation that is supported by Oracle and accepted by your hardware. If you have an incorrect terminal emulation, you will not be able to install the database. The following are supported by Oracle:

3151, 386*, Q303*, Q310*, ansi, dec, hft*, hp, iris, lft, sun, vt100, vt220,wy50, wy150, x1ft, xsun, xsun50 (*Multiple versions of these emulations are available.)

What terminal emulation do you choose? _____

The following tables represent the raw logical volumes necessary for Oracle 7 database storage and scripts. The logical volumes for the database on node1 are in Table 21 on page 90. The logical volumes for the database on node2 are in Table 22 on page 90. The names of these logical volumes are set. The numbers represent the recommended sizes of each file. Please fill in the desired sizes.

Table 21. Database Logical Volumes (Node 1)

Name	LV Name	Default Size	Your Size
Control File 1	node1_c1lv	12 MB	
Control File 2	node1_c2lv	12 MB	
Control File 3	node1_c3lv	12 MB	
System Tablespace	node1_dblv	80 MB	
Logfile 1	node1_log1lv	12 MB	
Logfile 2	node1_log2lv	12 MB	
Logfile 3	node1_log3lv	12 MB	
Rollback Tablespace	node1_log1lv	20 MB	
Temp Tablespace	node1_tmplv	12 MB	
User Tablespace	node1_userlv	12 MB	
Tools Tablespace	node1_toollv	40 MB	

Table 22. Database Logical Volumes (Node 2)

Name	LV Name	Default Size	Your Size
Control File 1	node2_c1lv	12 MB	12M
Control File 2	node2_c2lv	12 MB	12M
Control File 3	node2_c3lv	12 MB	12M
System Tablespace	node2_dblv	80 MB	
Logfile 1	node2_log1lv	12 MB	
Logfile 2	node2_log2lv	12 MB	
Logfile 3	node2_log3lv	12 MB	
Rollback Tablespace	node2_log1lv	20 MB	

Name	LV Name	Default Size	Your Size
Temp Tablespace	node2_templv	12 MB	
User Tablespace	node2_userlv	12 MB	
Tools Tablespace	node2_toollv	40 MB	

A.2 Rotating Standby Worksheet

This worksheet is to help you in gathering the necessary information to smoothly complete the install of the Oracle7 Database on a Highlay Available Cluster using the Quick Install Program. This worksheet is for the Rotating Configuration. Please write the information in the areas provided, and keep it with you during the installation and setup process.

Installation Information:

The following tables hold the adapter information for the cluster. In the Rotating configuration each cluster has a standby and boot adapter, but both share a service adapter. Please fill in the labels and IP addresses for each adapter. As an example, the shared adapter for node 1 and node 2 in our sample configuration would have a label of **node1** and an IP address of **10.0.1.1**.

Table 23. Node 1 Adapters for Rotating Configuration

Node 1 Adapters	Adapter Label	Adapter Address
Boot Adapter		
Service Adapter		
Standby Adapter		

Table 24. Node 2 Adapters for Rotating Configuration

Node 2 Adapters	Adapter Label	Adapter Address
Boot Adapter		
Service Adapter	Node1 Shared Adapter	Node1 Shared Adapter
Standby Adapter		

You need to find an available Group ID and User ID for both node1 and node2. These ID #'s must be the same on both machines. The defaults used in our sample configuration are Group #201, and User #210. To find this information, enter **smit group** or **smit user**, and choose the **list** option.

Available Group ID # _____
Available User ID # _____

Pick a terminal emulation that is supported by Oracle, and accepted by your hardware. If you have an incorrect terminal emulation, you will not be able to install the database. The following are supported by Oracle:

3151, 386*, Q303*, Q310*, ansi, dec, hft*, hp, iris, lft, sun, vt100, vt220,wy50, wy150, xlft, xsun, xsun50 (* Multiple versions of these emulations are available.)

What terminal emulation do you choose? _____

The following tables represent the raw logical volumes necessary for Oracle 7 database storage and scripts. The logical volumes for the database on node1 and node2 are in Table 25 on page 93. Because we only require one database in a Rotating configuration, the logical volumes on node2 are exactly the same as those on node1. The names of these logical volumes are set. The numbers represent the recommended sizes of each file. Please fill in the desired sizes.

Table 25. Database Logical Volumes for Both Nodes in Rotating Configuration

Name	LV Name	Default Size	Your Size
Control File 1	node1_c1lv	12 MB	
Control File 2	node1_c2lv	12 MB	
Control File 3	node1_c3lv	12 MB	
System Tablespace	node1_dblv	80 MB	
Logfile 1	node1_log1lv	12 MB	
Logfile 2	node1_log2lv	12 MB	
Logfile 3	node1_log3lv	12 MB	
Rollback Tablespace	node1_log1lv	20 MB	
Temp Tablespace	node1_templv	12 MB	
User Tablespace	node1_userlv	12 MB	
Tools Tablespace	node1_toollv	40 MB	

A.3 Hot Standby Worksheet

This worksheet is to help you in gathering the necessary information to smoothly complete the install of the Oracle7 Database on a Highlay Available Cluster using the Quick Install Program. This worksheet is for the Hot Standby Configuration. Please write the information in the areas provided, and keep it with you during the installation and setup process.

Installation Information:

The following tables hold the adapter information for the cluster. In the Hot Standby configuration each cluster has a service and standby adapter, but node 2 has no boot adapter. Please fill in the labels and IP addresses for each adapter. As an example, the standby adapter for node 1 in our sample configuration would have a label of `node1_standby` and an IP address of `10.0.2.1`.

Table 26. Node1 Adapters for Hot Standby Configuration

Node 1 Adapters	Adapter Label	Adapter Address
Boot Adapter		
Service Adapter		
Standby Adapter		

Table 27. Node2 Adapters for Hot Standby Configuration

Node 2 Adapters	Adapter Label	Adapter Address
Boot Adapter		
Service Adapter		
Standby Adapter		

You need to find an available Group ID and User ID for both node1 and node2. These ID #'s must be the same on both machines. The defaults used in our sample configuration are Group #201, and User #210. To find this information, enter **smit group** or **smit user**, and choose the **list** option.

Available Group ID # _____

Available User ID # _____

Pick a terminal emulation that is supported by Oracle, and accepted by your hardware. If you have an incorrect terminal emulation, you will not be able to install the database. The following are supported by Oracle:

3151, 386*, Q303*, Q310*, ansi, dec, hft*, hp, iris, lft, sun, vt100, vt220,wy50, wy150, xlft, xsun, xsun50 (* Multiple versions of these emulations are available.)

What terminal emulation do you choose?

What terminal emulation do you choose? _____
--

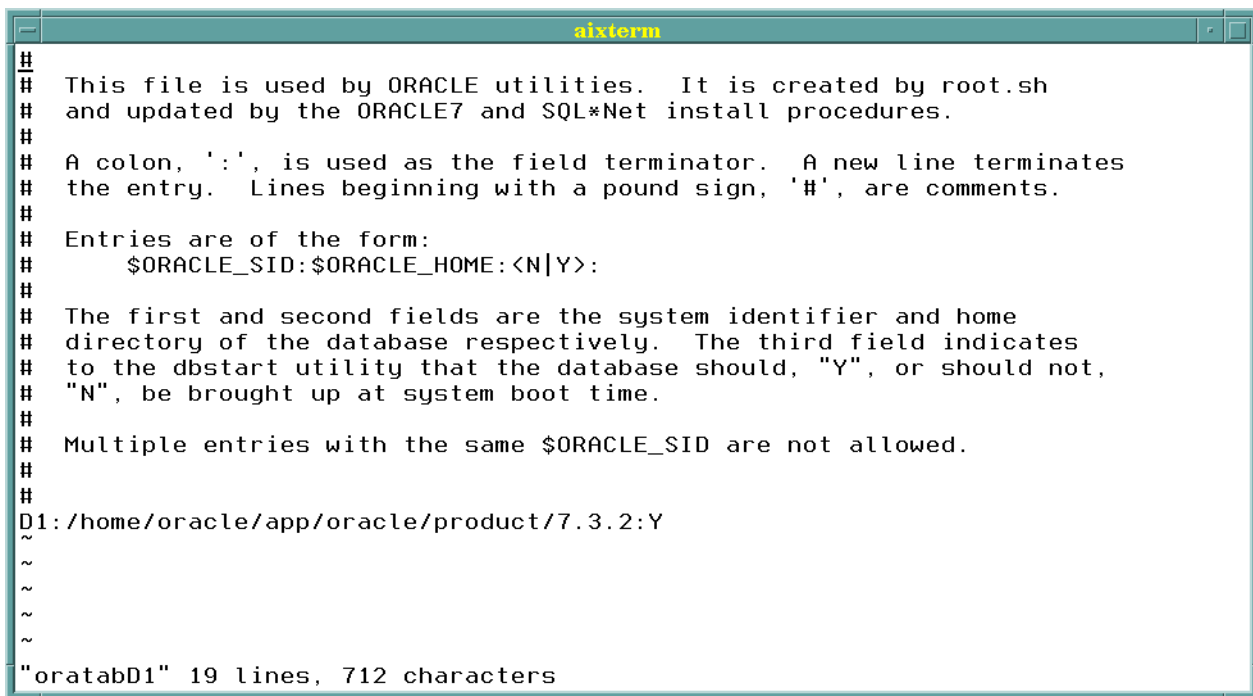
The following tables represent the raw logical volumes necessary for Oracle 7 database storage and scripts. The logical volumes for the database on node 1 and node 2 are in Table 28. Because we only require one database in a Hot Standby configuration, the logical volumes on node2 are exactly the same as those on node1. The names of these logical volumes are set. The numbers represent the recommended sizes of each file. Please fill in the desired sizes.

Table 28. Database Logical Volumes for Both Nodes in Hot Standby Configuration

Name	LV Name	Default Size	Your Size
Control File 1	node1_c1lv	12 MB	
Control File 2	node1_c2lv	12 MB	
Control File 3	node1_c3lv	12 MB	
System Tablespace	node1_dblv	80 MB	
Logfile 1	node1_log1lv	12 MB	
Logfile 2	node1_log2lv	12 MB	
Logfile 3	node1_log3lv	12 MB	
Rollback Tablespace	node1_log1lv	20 MB	
Temp Tablespace	node1_templv	12 MB	
User Tablespace	node1_userlv	12 MB	
Tools Tablespace	node1_toollv	40 MB	

B.1.3 File: /etc/oratab

Below is the /etc/oratab file for nodes in a Rotating or Hot Standby configuration.



```
#
# This file is used by ORACLE utilities.  It is created by root.sh
# and updated by the ORACLE7 and SQL*Net install procedures.
#
# A colon, ':', is used as the field terminator.  A new line terminates
# the entry.  Lines beginning with a pound sign, '#', are comments.
#
# Entries are of the form:
#   $ORACLE_SID:$ORACLE_HOME:<N|Y>:
#
# The first and second fields are the system identifier and home
# directory of the database respectively.  The third field indicates
# to the dbstart utility that the database should, "Y", or should not,
# "N", be brought up at system boot time.
#
# Multiple entries with the same $ORACLE_SID are not allowed.
#
#
D1:/home/oracle/app/oracle/product/7.3.2:Y
~
~
~
~
"oratabD1" 19 lines, 712 characters
```

Figure 14. /etc/oratab File for Rotating or Hot Standby Resources.

B.1.4 File: //node1_scripts/listener.ora

The figure below shows a Listener .ora file for a system running a single database.


```
listener
lisD1=
  (ADDRESS_LIST=
    (ADDRESS=
      (PROTOCOL=TCP)
      (HOST=node1)
      (PORT=1521)
    )
  )

STARTUP_WAIT_TIME_LISTENER = 0
CONNECT_TIMEOUT_LISTENER = 10

sid_list_lisD1=
  (SID_LIST=
    (SID_DESC=
      (SID_NAME=D1)
      (ORACLE_HOME=/home/oracle/app/oracle/product/7.3.2)
    )
  )

TRACE_LEVEL_LISTENER = OFF
USE_CKPFIL_LISTENER = true
~
~
"listener.ora" 22 lines, 328 characters
```

Figure 15. listener.ora File for Rotating or Hot Standby Resources

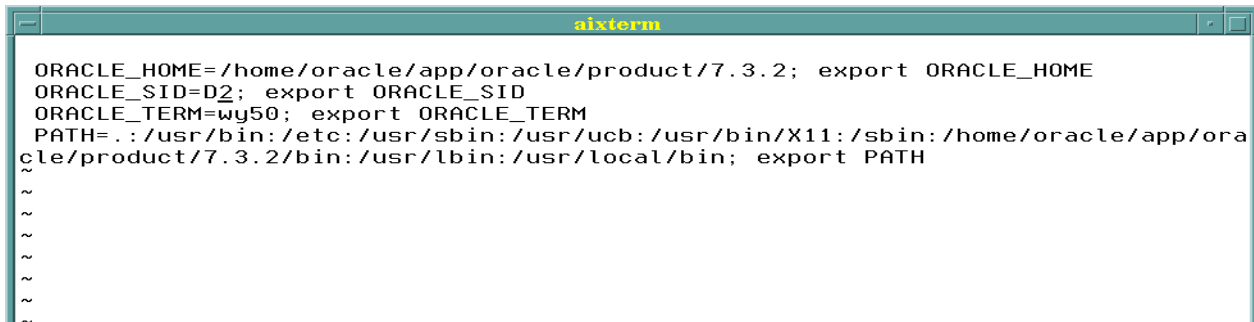
B.1.5 File: /node1_scripts/tsnames.ora

This is the /home/oracle/tns_admin/tsnames.ora file for Rotating or Hot Standby resources.

```
listener
D1=
  (DESCRIPTION=
    (ADDRESS_LIST=
      (ADDRESS=
        (PROTOCOL=TCP)
        (HOST =node1)
        (PORT =1521)
      )
    )
  )
  (CONNECT_DATA=
    (SID=D1)
    (GLOBAL_NAME=D1)
  )
~
~
~
~
~
~
~
~
~
```

Figure 16. tnsnames.ora File for Rotating or Hot Standby Resources

This is the /home/oracle/.profile file for node 2.



```
aixterm
ORACLE_HOME=/home/oracle/app/oracle/product/7.3.2; export ORACLE_HOME
ORACLE_SID=D2; export ORACLE_SID
ORACLE_TERM=wy50; export ORACLE_TERM
PATH=./usr/bin:/etc:/usr/sbin:/usr/ucb:/usr/bin/X11:/sbin:/home/oracle/app/oracle/product/7.3.2/bin:/usr/sbin:/usr/local/bin; export PATH
~
~
~
~
~
~
~
```

Figure 19. /home/oracle/.profile File on Second Node

The only difference we had between the Oracle user .profiles was the different ORACLE_SID for each database.

You will probably need to change the following variable in the .profile files above:

- ORACLE_TERM

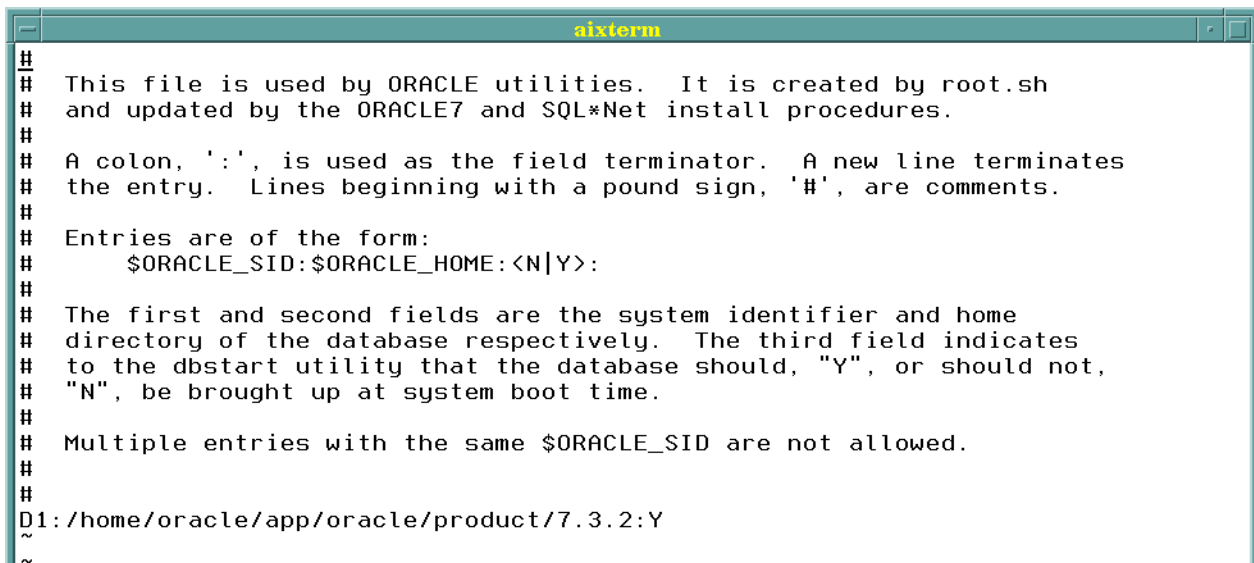
C.1.3 File: /etc/oratab

Below is the /etc/oratab file for nodes in a Mutual Takeover configuration.

A copy of the relevant oratab file is kept in our scripts directory. At cluster startup, we issue:

```
cp -p /node1_scripts/oratabD1 /etc/oratab
```

This is the oratab entry for node 1.



```
aixterm
#
# This file is used by ORACLE utilities.  It is created by root.sh
# and updated by the ORACLE7 and SQL*Net install procedures.
#
# A colon, ':', is used as the field terminator.  A new line terminates
# the entry.  Lines beginning with a pound sign, '#', are comments.
#
# Entries are of the form:
#   $ORACLE_SID:$ORACLE_HOME:<N|Y>:
#
# The first and second fields are the system identifier and home
# directory of the database respectively.  The third field indicates
# to the dbstart utility that the database should, "Y", or should not,
# "N", be brought up at system boot time.
#
# Multiple entries with the same $ORACLE_SID are not allowed.
#
#
D1:/home/oracle/app/oracle/product/7.3.2:Y
~
~
```

Figure 20. Node 1 oratab Entry for the Database

This is the oratab entry for node 2. We no longer show the comments contained in the file.

```
D2:/home/oracle/app/oracle/product/7.3.2:Y
```

Figure 21. Node 2 oratab Entry for the Database

When there has a been a failover, we use a file called oratabD12. This contains both database entries, as shown below.

```
D1:/home/oracle/app/oracle/product/7.3.2:Y
D2:/home/oracle/app/oracle/product/7.3.2:Y
```

Figure 22. oratab Entries for Both Databases Running on One Node

We copied this file into /etc/oratab. Make sure the file is copied with the *oracle* user ID as its owner.

C.1.4 File: /node1_scripts/listener.ora

The figure below shows a listener.ora file for a system running more than one database. We had two databases and thus created a Listener capable of working with each database separately or with both databases simultaneously.

This is the listener.ora file for node 1.



```
lisD12=
  (ADDRESS_LIST=
    (ADDRESS=
      (PROTOCOL=TCP)
      (HOST=node1)
      (PORT=1521)
    )
  )

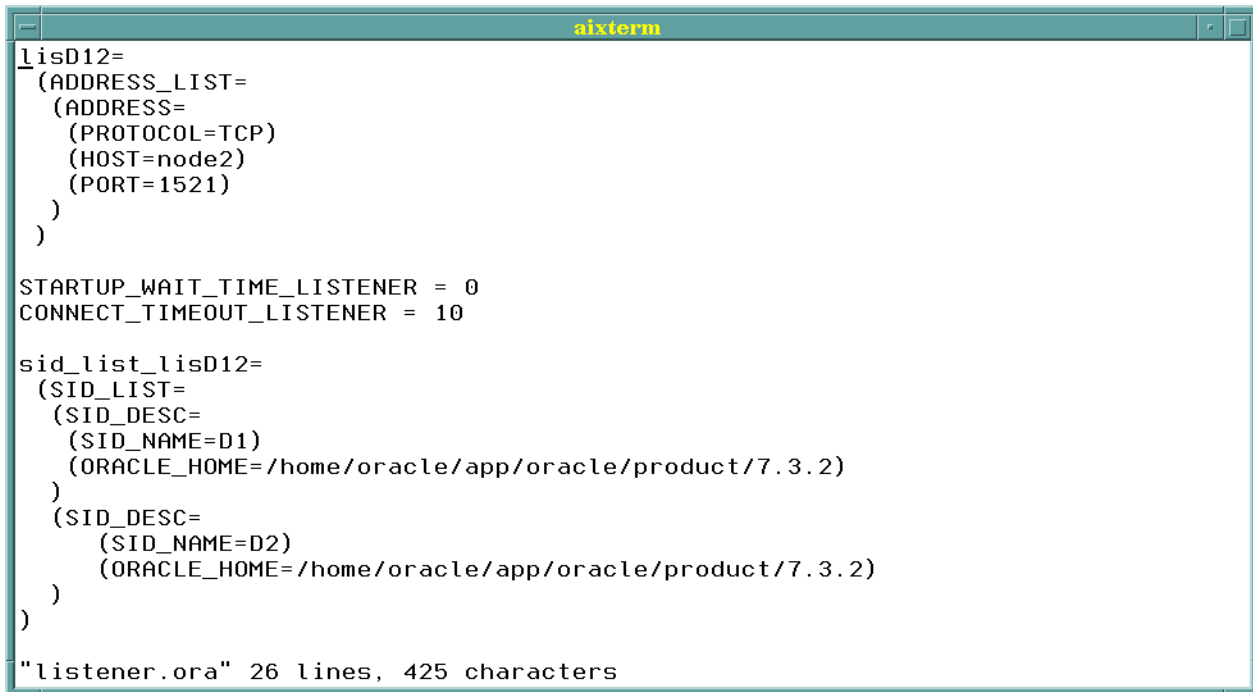
STARTUP_WAIT_TIME_LISTENER = 0
CONNECT_TIMEOUT_LISTENER = 10

sid_list_lisD12=
  (SID_LIST=
    (SID_DESC=
      (SID_NAME=D1)
      (ORACLE_HOME=/home/oracle/app/oracle/product/7.3.2)
    )
    (SID_DESC=
      (SID_NAME=D2)
      (ORACLE_HOME=/home/oracle/app/oracle/product/7.3.2)
    )
  )

"listener.ora" 26 lines, 425 characters
```

Figure 23. listener.ora File for Mutual Takeover Resources, Node 1

This is the listener.ora file for node 2.



```
lisD12=
  (ADDRESS_LIST=
    (ADDRESS=
      (PROTOCOL=TCP)
      (HOST=node2)
      (PORT=1521)
    )
  )

STARTUP_WAIT_TIME_LISTENER = 0
CONNECT_TIMEOUT_LISTENER = 10

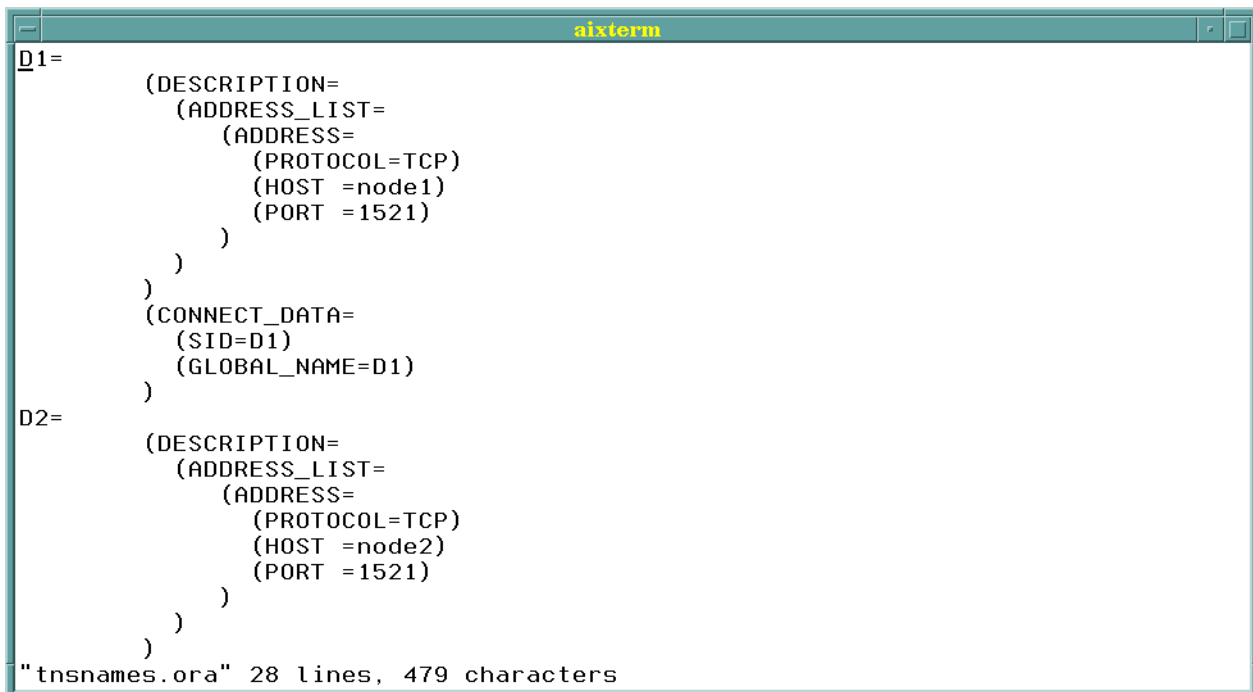
sid_list_lisD12=
  (SID_LIST=
    (SID_DESC=
      (SID_NAME=D1)
      (ORACLE_HOME=/home/oracle/app/oracle/product/7.3.2)
    )
    (SID_DESC=
      (SID_NAME=D2)
      (ORACLE_HOME=/home/oracle/app/oracle/product/7.3.2)
    )
  )
)

"listener.ora" 26 lines, 425 characters
```

Figure 24. listener.ora File for Mutual Takeover Resources, Node 2

C.1.5 File: //node1_scripts/tnsnames.ora

The tnsnames.ora file is the same on both nodes.



```
D1=
  (DESCRIPTION=
    (ADDRESS_LIST=
      (ADDRESS=
        (PROTOCOL=TCP)
        (HOST =node1)
        (PORT =1521)
      )
    )
  )
)
(CONNECT_DATA=
  (SID=D1)
  (GLOBAL_NAME=D1)
)
D2=
  (DESCRIPTION=
    (ADDRESS_LIST=
      (ADDRESS=
        (PROTOCOL=TCP)
        (HOST =node2)
        (PORT =1521)
      )
    )
  )
)
)

"tnsnames.ora" 28 lines, 479 characters
```

Figure 25. tnsnames.ora File for Mutual Takeover Resources

Appendix D. Special Notices

This publication is intended to help network administrators to install and set up highly available Oracle databases. The information in this publication is not intended as the specification of any programming interfaces that are provided by HACMP for AIX or for the Oracle database products. See the PUBLICATIONS section of the IBM Programming Announcement for HACMP for AIX for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood, NY 10594 USA.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The information about non-IBM ("vendor") products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including these reference numbers is to alert IBM customers to specific information relative

to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

The following terms, which are denoted by an asterisk (*) in this publication, are trademarks of the International Business Machines Corporation in the United States and/or other countries:

AIX ®	AS/400 ®
IBM ®	PROFS ®
RISC System/6000 ®	

The following terms are trademarks of other company:

UNIX is a registered trademarks in the United States and other countries licensed exclusively through X/Open Company Limited.

Other company, product, and service names may be trademarks or service marks of others.

Appendix E. Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

E.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see “How To Get ITSO Redbooks” on page 109.

- *High Availability on the RISC System/6000 Family*, SG24-4551
- *An HACMP Cookbook*, SG24-4553

E.2 Redbooks on CD-ROMs

Redbooks are also available on CD-ROMs. **Order a subscription** and receive updates 2-4 times a year at significant savings.

CD-ROM Title	Subscription Number	Collection Kit Number
System/390 Redbooks Collection	SBOF-7201	SK2T-2177
Networking and Systems Management Redbooks Collection	SBOF-7370	SK2T-6022
Transaction Processing and Data Management Redbook	SBOF-7240	SK2T-8038
AS/400 Redbooks Collection	SBOF-7270	SK2T-2849
RISC System/6000 Redbooks Collection (HTML, BkMgr)	SBOF-7230	SK2T-8040
RISC System/6000 Redbooks Collection (PostScript)	SBOF-7205	SK2T-8041
Application Development Redbooks Collection (available soon)	SBOF-7290	SK2T-8037
Personal Systems Redbooks Collection (available soon)	SBOF-7250	SK2T-8042

E.3 Other Publications

These publications are also relevant as further information sources:

- *HACMP for AIX, Version 4.2.1: Concepts and Facilities*, SC23-1938
- *HACMP for AIX, Version 4.2.1: Planning Guide*, SC23-1939
- *HACMP for AIX, Version 4.2.1: Installation Guide*, SC23-1940
- *HACMP for AIX, Version 4.2.1: Administration Guide*, SC23-1941
- *HACMP for AIX, Version 4.2.1: Troubleshooting Guide*, SC23-1942
- *HACMP for AIX, Version 4.2.1: Programming Locking Applications*, SC23-1943
- *HACMP for AIX, Version 4.2.1: Programming Client Applications*, SC23-1944
- *HACMP for AIX, Version 4.2.1: Master Index and Glossary*, SC23-1945
- *Official Oracle product documentation, shipped with the database product.*

How To Get ITSO Redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, CD-ROMs, workshops, and residencies. A form for ordering books and CD-ROMs is also provided.

This information was current at the time of publication, but is continually subject to change. The latest information may be found at URL <http://www.redbooks.ibm.com>.

How IBM Employees Can Get ITSO Redbooks

Employees may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **PUBORDER** – to order hardcopies in United States
- **GOPHER link to the Internet** – type `GOPHER.WTSCPOK.ITSO.IBM.COM`
- **Tools disks**

To get LIST3820s of redbooks, type one of the following commands:

```
TOOLS SENDTO EHONE4 TOOLS2 REDPRINT GET SG24xxxx PACKAGE
TOOLS SENDTO CANVM2 TOOLS REDPRINT GET SG24xxxx PACKAGE (Canadian users only)
```

To get lists of redbooks:

```
TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET ITSOCAT TXT
TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET LISTSERV PACKAGE
```

To register for information on workshops, residencies, and redbooks:

```
TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ITSOREGI 1996
```

For a list of product area specialists in the ITSO:

```
TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ORGCARD PACKAGE
```

- **Redbooks Home Page on the World Wide Web**

<http://w3.itso.ibm.com/redbooks>

- **IBM Direct Publications Catalog on the World Wide Web**

<http://www.elink.ibm.com/pbl/pbl>

IBM employees may obtain LIST3820s of redbooks from this page.

- **REDBOOKS category on INEWS**
- **Online** – send orders to: `USIB6FPL` at `IBMMAIL` or `DKIBMBSH` at `IBMMAIL`
- **Internet Listserver**

With an Internet E-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an E-mail note to `announce@webster.ibm.com` with the keyword `subscribe` in the body of the note (leave the subject line blank). A category form and detailed instructions will be sent to you.

How Customers Can Get ITSO Redbooks

Customers may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **Online Orders** (Do not send credit card information over the Internet) – send orders to:

	IBMMAIL	Internet
In United States	usib6fpl at ibmmail	usib6fpl@ibmmail.com
In Canada	caibmbkz at ibmmail	lmannix@vnet.ibm.com
Outside North America	dkibmbsh at ibmmail	bookshop@dk.ibm.com

- **Telephone orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	(long distance charges apply)
(+45) 4810-1320 - Danish	(+45) 4810-1020 - German
(+45) 4810-1420 - Dutch	(+45) 4810-1620 - Italian
(+45) 4810-1540 - English	(+45) 4810-1270 - Norwegian
(+45) 4810-1670 - Finnish	(+45) 4810-1120 - Spanish
(+45) 4810-1220 - French	(+45) 4810-1170 - Swedish

- **Mail Orders** – send orders to:

IBM Publications Publications Customer Support P.O. Box 29570 Raleigh, NC 27626-0570 USA	IBM Publications 144-4th Avenue, S.W. Calgary, Alberta T2P 3N5 Canada	IBM Direct Services Sortemosevej 21 DK-3450 Allerød Denmark
--	--	--

- **Fax** – send orders to:

United States (toll free)	1-800-445-9269
Canada	1-800-267-4455
Outside North America	(+45) 48 14 2207 (long distance charge)

- **1-800-IBM-4FAX (United States) or (+1) 408 256 5422 (Outside USA)** – ask for:

Index # 4421 Abstracts of new redbooks
Index # 4422 IBM redbooks
Index # 4420 Redbooks for last six months

- **Direct Services** – send note to softwareshop@vnet.ibm.com

- **On the World Wide Web**

Redbooks Home Page	http://www.redbooks.ibm.com
IBM Direct Publications Catalog	http://www.elink.ibm.link.ibm.com/pbl/pbl

- **Internet Listserver**

With an Internet E-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an E-mail note to announce@webster.ibm.link.ibm.com with the keyword `subscribe` in the body of the note (leave the subject line blank).

List of Abbreviations

<i>ANSI</i>	American National Standards Institute
<i>APA</i>	All Points Addressable
<i>ARCH</i>	Archiver Process
<i>ASCII</i>	American Standard Code for Information Interchange
<i>AS/400</i>	Application System/400
<i>CASE</i>	Computer Assisted Software Engineering
<i>CKPT</i>	Checkpoint
<i>DBWR</i>	Database Writer
<i>DNS</i>	Domain Name Server
<i>DSMIT</i>	Distributed System Management Interface Tool
<i>GODM</i>	Global Object Data Manager
<i>HACMP</i>	High Availability Cluster Multi-Processing
<i>HANFS</i>	High Availability Network File System
<i>IBM</i>	International Business Machines Corporation
<i>ITSO</i>	International Technical Support Organization
<i>LGWR</i>	Log Writer
<i>LOCK</i>	Parallel Server Lock Processor
<i>ODM</i>	Object Data Manager
<i>OLTP</i>	On Line Transaction Processing
<i>PGA</i>	Program Global Area
<i>PMON</i>	Process Monitor
<i>PROFS</i>	Professional Office System
<i>RECO</i>	Recover Process
<i>SGA</i>	System Global Area or Shared Global Area
<i>SMIT</i>	System Management Interface Tool
<i>SMON</i>	System Monitor
<i>VGDA</i>	Volume Group Descriptor Area
<i>VGSA</i>	Volume Group Status Area

Index

Symbols

/etc/oratab 98, 102
/etc/services 101
/home/oracle/.profile 97, 101
/home/oracle/tns_admin/listener.ora 98, 103
/home/oracle/tns_admin/tnsnames.ora 104

A

abbreviations 113
adapter configuration 42
archived redo log 9
archiver (ARCH) process 14

B

bibliography 107

C

cascading cluster 32
checkpoint (CKPT) 13
client/server architecture 4
context area 3
control file 9
control files 67
cursor 3

D

data buffer cache 10
data dictionary 7
data files 7
data segment 9
database domain name 67
database writer (DBWR) 13
DBA group 48
decision support database 16
dedicated serve 11
development database 16
disk space requirements 19

E

errors 85
extent 8

H

hot standby configuration 31

I

in doubt transaction 14
index segment 9
init.ora file 10
instance 11

L

listener process 84
log writer (LGWR) 13
logical volumes for Oracle 7 storage 40

M

mixed database 16
Multi-Threaded Server 66
multithreaded server 11
mutual takeover 32
mutual takeover configuration 32

N

network type 42

O

OLTP database 16
online redo log 9
Oracle configuration files for cascading resources 101
Oracle user 49
Oracle user environment 51
ORACLE_HOME 57
ORACLE_SID variable 61
ORACLE_TERM 39

P

parallel server lock processes (LCKn) 14
parsing 3
PGA 10
process monitor (PMON) 13
Program Global Area 10

R

recover (RECO) process 14
redo log buffer 10
redo log files 9
rollback segment 8
rotating Ccluster 31
rotating standby configuration 30

S

segment 8
SGA 10
shared pool 10
sizing 19
SQL 2
SQL language 2
system monitor (SMON) 13
system tablespace 8

T

tablespaces 8, 67
TCP service port 65

temporary segment 9
testing the Oracle server database 81

U

user tablespace 8

W

worksheets 89

ITSO Redbook Evaluation

Oracle Cluster POWERsolution Guide
SG24-2019-00

Your feedback is very important to help us maintain the quality of ITSO redbooks. **Please complete this questionnaire and return it using one of the following methods:**

- Use the online evaluation form found at <http://www.redbooks.com>
- Fax this form to: USA International Access Code + 1 914 432 8264
- Send your comments in an Internet note to redeval@vnet.ibm.com

Please rate your overall satisfaction with this book using the scale:
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction _____

Please answer the following questions:

Was this redbook published in time for your needs? Yes___ No___

If no, please explain:

What other redbooks would you like to see published?

Comments/Suggestions: (THANK YOU FOR YOUR FEEDBACK!)
