COMPAQ

Oracle Workgroup Server 7.3.3 Service Failover Using Microsoft Cluster Server

TechNote

First Edition (December 1997)
Part Number 297325-001
Compaq Computer Corporation

Notice

The information in this publication is subject to change without notice.

COMPAQ COMPUTER CORPORATION SHALL NOT BE LIABLE FOR TECHNICAL OR EDITORIAL ERRORS OR OMISSIONS CONTAINED HEREIN, NOR FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES RESULTING FROM THE FURNISHING, PERFORMANCE, OR USE OF THIS MATERIAL. THIS INFORMATION IS PROVIDED “AS IS” AND COMPAQ COMPUTER CORPORATION DISCLAIMS ANY WARRANTIES, EXPRESS, IMPLIED OR STATUTORY AND EXPRESSLY DISCLAIMS THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR PARTICULAR PURPOSE, GOOD TITLE AND AGAINST INFRINGEMENT.

The software described in this guide is furnished under a license agreement or nondisclosure agreement. The software may be used or copied only in accordance with the terms of the agreement.

Compaq, Deskpro, Compaq Insight Manager, ProLiant, ROMPaq, SmartStart, NetFlex, registered United States Patent and Trademark Office.

Netelligent is a trademark and/or service mark of Compaq Computer Corporation.

Microsoft, MS-DOS, Windows, and Windows NT are registered trademarks of Microsoft Corporation.

Other product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

Oracle Workgroup Server 7.3.3 Service Failover Using Microsoft Cluster Server

First Edition (December 1997)
Part Number 297325-001

Contents

Chapter 1
About This TechNote

Introduction

The purpose of this document is to provide a practical and effective way of implementing a successful Compaq ProLiant Cluster service failover for Oracle Workgroup Server 7.3.3 using Microsoft Cluster Server.

Expected Audience

This document is intended to help in the installation and configuration of Microsoft Cluster Server with Oracle Workgroup Server 7.3.3 and assumes that the reader has knowledge of the following:

- Installing and configuring Compaq Servers
- Installing and configuring Microsoft Windows NT Server Enterprise
- Microsoft Windows NT Server interface and setting up and administering user accounts in a domain
- Installing Microsoft Cluster Server
- Installing Oracle Workgroup Server 7.3.3

A supplemental document, *Introduction to ProLiant Clusters*, is available to help you install and configure the Microsoft Cluster Server on a Compaq ProLiant Cluster.

Text Conventions

This document uses the following conventions to distinguish elements of text:

Table 1-1
Text Conventions

Convention	Use
OK, CANCEL	Window command button labels appear in bold caps.
CTRL + DEL	Keyboard keys appear in bold caps.
c:\dirname\filename.exe	Path Names of items such as files, directories, resources, groups, and services appear in bold italics.
Select **Item** → **Item**	Items you select from a pull down menu appear in bold initial caps, separated by arrows for each submenu item
USER INPUT	Information to be entered by the user is shown in uppercase italics

Chapter 2
Overview of a Failover

Types of Failover

In the first release of Microsoft Cluster Server, there are two types of failover: service and resource. Both types of failover allow for high levels of server and application availability.

In a service failover configuration, a service is defined to Microsoft Cluster Server as a resource. Once defined, its Failover Manager Process ensures that the service is running on only one member of the cluster at a given time. Although the service is running on only one member of the cluster at a time, the resource group where the service is defined uses a common name. Therefore, all services running on that resource group are available to all connected network client machines using the common name.

A more comprehensive model is the resource failover, which takes advantage of Independent Software Vendor's (ISV) cluster-aware resource DLLs to execute the failover. In a resource failover configuration, it is assumed that the service is running on both members of the cluster and that some resource such as a database or a website fails over, and not the service.

Since most ISV applications do not yet have resources to make them cluster-aware, service failovers are the most common forms of implementing the failover in Microsoft Cluster Server. Microsoft provides a generic service resource DLL to allow these applications to function under Microsoft Cluster Server.

What are the Advantages of a Failover?

Both service and resource failovers provide a highly available computing environment with relative simplicity and can be used for both planned and unplanned service outages. A planned outage occurs when one of the members in the cluster is taken offline for upgrades or maintenance. An unplanned outage occurs when one of the members or the application on that member fails. During both types of outages, the services and/or resources can be moved to the other member of the cluster. Since the clients do not communicate to the server directly but through a virtual server, the failover can be transparent to the clients.

Note: This depends on the mechanism that allows the client and the clusters to re-establish a live connection to continue the services. For example, developing client software that automatically re-establishes a connection after loosing them provides near transparent failover.

What are the Disadvantages of a Failover?

One of the main disadvantages of a Microsoft Cluster Server Failover is that in a failover condition, the performance of the server taking over the work load of the failed server may degrade since it is now doing the work of both servers. However, by manually balancing the load and selecting the appropriate hardware, performance degradation can be minimized.

Another major disadvantage of a failover environment is that the state information of the application is not monitored. As a result when the application fails over, an entire new instance of the application is started. Any work in progress is lost and data recovery must occur.

Chapter 3
Model Failover Environment

Introduction

The model environment presented here consists of two Compaq ProLiant servers sharing a ProLiant Storage Unit and its drives, networked to each other through a LAN and a dedicated server interconnect.

In order to enhance availability, Compaq strongly recommends the use of at least two PCI network cards per server. One of these cards should be connected to its counterpart in the other node of the cluster via a private hub or a special crossover cable, such as the one supplied in the cluster kit. These cards should be configured for cluster communication only. The other card in each server should be configured for both client access and cluster communication. This configuration allows the client network to provide an alternative path for the heartbeat should the dedicated connection fail.

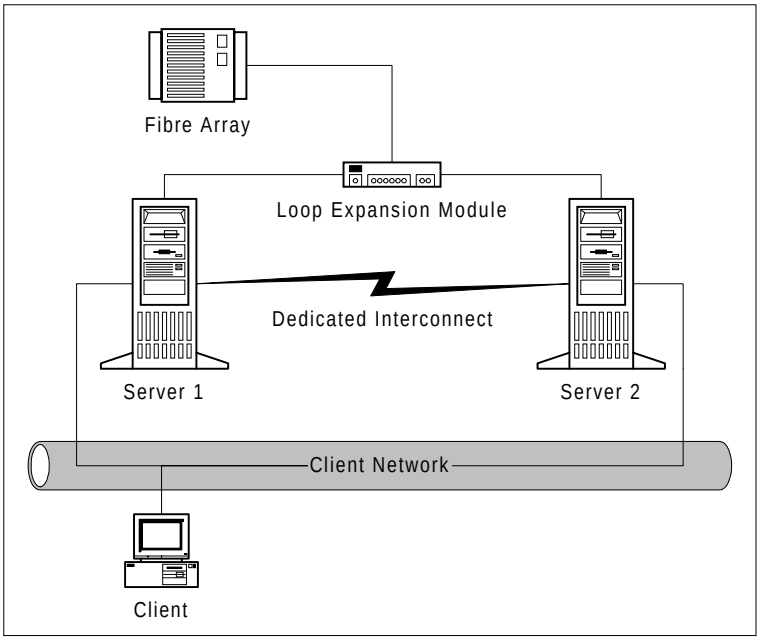

Figure 3-1. Standard Compaq ProLiant Cluster Configuration

Compaq recommends the use of duplexed network cards configured for client access and cluster communication for even higher availability, in addition to using a dedicated connection as described above.

Note: To create a cluster, a Domain environment is required using TCP/IP as the network protocol. Both servers as a cluster should be configured to be servers in the domain. It is highly recommended that you do not use either server in a cluster as a Primary or Backup domain controller.

Model Environment Naming Conventions

Table 3-1
System and Cluster Specifics

Operating System	Windows Enterprise Server 4.0 w/ Service Pack 3
Domain Name	OR7DOMAIN
Name of Server 1	OR7NODE1
IP Address of Server 1	110.12.120.111 (subnet mask 255.255.0.0)
IP Address of Interconnect	110.14.120.111 (subnet mask 255.255.255.0)
Name of Server 2	OR7NODE2
IP Address of Server 2	110.12.120.112 (subnet mask 255.255.0.0)
IP Address of Interconnect	110.14.120.112 (subnet mask 255.255.255.0)
Name of Cluster	OR7CLUSTER
IP Address of Cluster	110.12.120.120 (subnet mask 255.255.0.0)
Location of Quorum Disk	Disk E:
Location of Cluster Files	C:\WinNT\system32

Table 3-2
Oracle Database Specifics

Service	ORACLE WORKGROUP SERVER 7.3.3
Location of Executables	C:\Orant\
Location of Shared Disk	Disk F:
Location of Shared Data Files	F:\DBORCL\DATA
Location of Shared Control Files	F:\DBORCL\DATA
Location of Shared Log Files	F:\DBORCL\LOG
Location of Shared Trace Files	F:\DBORCL\TRACE
Service ID	ORCL
Instance Name	ORACLE

Table 3-3
Oracle Database Resource Group Specifics

IP Address	110.12.120.121 (subnet mask 255.255.0.0)
Network Name	OR7GROUP
TNS Listener Service	OracleTNSListener
Oracle Server Service	OracleServiceORCL
Oracle Instance Service	OracleStartORCL

Hardware Requirements

Server Configuration

Two Compaq ProLiant 2500 Servers, each with:

- 1 processor / 128-Mbytes RAM
- Compaq Integrated Netelligent 10/100 TX UTP Controller
- Compaq Netelligent 10/100 TX PCI UTP Network Interface Controller
- Compaq Fibre Channel Host PCI Controller

Compaq ProLiant Fibre Channel Disk Storage System F2 with:

- Compaq ProLiant Fibre Channel Array Controller
- 4 SCSI hard disk drives (RAID 1, 2 logical disks)

Client Configuration

Compaq Professional Workstation 5100 with:

- Embedded 10/100 TX PCI UTP Network Interface Controller
- 32-Mbytes RAM

For detailed configuration information, see the Appendix.

Software Requirements

Server Configuration

- Compaq SmartStart and Support Software Release 3.4 or later
- Compaq Array Configuration Utility 1.20 or later
- Microsoft Windows NT Enterprise Server 4.0
- Microsoft Windows NT Service Pack 3
- Compaq NT Support Software Utility 2.02 or later
- Microsoft Cluster Server 1.0
- Oracle Workgroup Server 7.3.3

Client Configuration

- Microsoft Windows NT Workstation 4.0
- Microsoft Windows NT Service Pack 3
- Compaq NT Support Software Utility 2.02 or later
- Microsoft Cluster Administrator 1.0
- Oracle 7.3.3 Client Products

Note: To create a cluster, a domain environment is required using TCP/IP as the network protocol. Both servers as a cluster should be configured to be a server in the domain. It is highly recommended that you do not use either server in a cluster as a Primary or Backup domain controller.

Chapter 4
Installing Microsoft Cluster Server

Microsoft Cluster Server must be installed on both servers before installing the application. Detailed installation instructions are be provided in this document. Refer to *Introduction to Compaq ProLiant Clusters* for complete instructions. This document also includes a section on troubleshooting. The following table details the necessary information to complete the installation of Microsoft Cluster Server on both servers.

Table 4-1
Microsoft Cluster Server Installation Specifics

Domain Name	OR7DOMAIN
Cluster Name	OR7CLUSTER
Cluster IP address	110.12.120.120 (subnet mask 255.255.0.0)
Shared Disks	Disk E:, F:
Cluster Admin User / Password	Administrator / <password>
Location of Cluster Executables	Default (C:\WinNT\System32)
Quorum Disk	Disk E:

Keep this information and everything from the model environment section ready. You will use all of these in installing the Microsoft Cluster Server.

Preliminary Tasks

Before moving forward with the Oracle Cluster configuration, the following tasks should already be completed.

- Hardware configuration for each server
- Array Configuration for the Fiber Channel shared storage
- Windows NT Enterprise Server 4.0 is installed
- Service Pack 3 for Windows NT 4.0 is applied
- Compaq NT SSD 2.02 or later (Support Software Disk) is applied
- Fiber Channel Loop driver is installed and both servers can access the shared drives.

Chapter 5
Creating the Oracle Cluster Group Resources

Use the following table to create the listed resources for this group.

Table 5-1
Resource Group (Oracle Database)

Resource Type	Resource Name	Value
Physical Disk	Disk F:	Disk F:
IP Address	IPOR7	110.12.120.121
Network Name	NMOR7	OR7GROUP
Generic Service	LSNROR7	OracleTNSListener
Generic Service	SRVCOR7	OracleServiceORCL
Generic Service	INSTOR7	OracleStartORCL

Dependencies

The following dependencies should be followed when creating the resources:

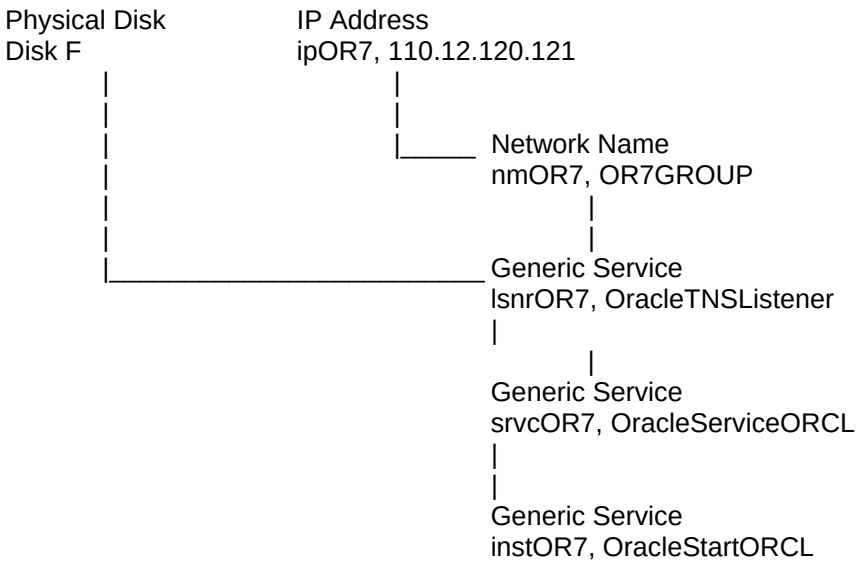

Complete the following steps before creating the IP Address and Network Name Resources:

1. From Node1, Launch the Cluster Administrator
2. When prompted, enter the Cluster Name OR7CLUSTER
3. Rename the "**Disk Group 1**" to "**MSCS Quorum Disk**"
4. Rename the "**Disk Group 2**" to "**Oracle Database**"

IP Address Resource

Use the following information to setup your IP Address resource.

Table 5-2
Oracle Database (IP Address) Properties

Resource Name	IPOR7
Resource Type	IP Address
IP address	110.12.120.121 (subnet mask 255.255.0.0)
Dependencies	None

1. Enter **IPOR7** as the resource name. Select **IP Address** as the resource type. Click **Next**.

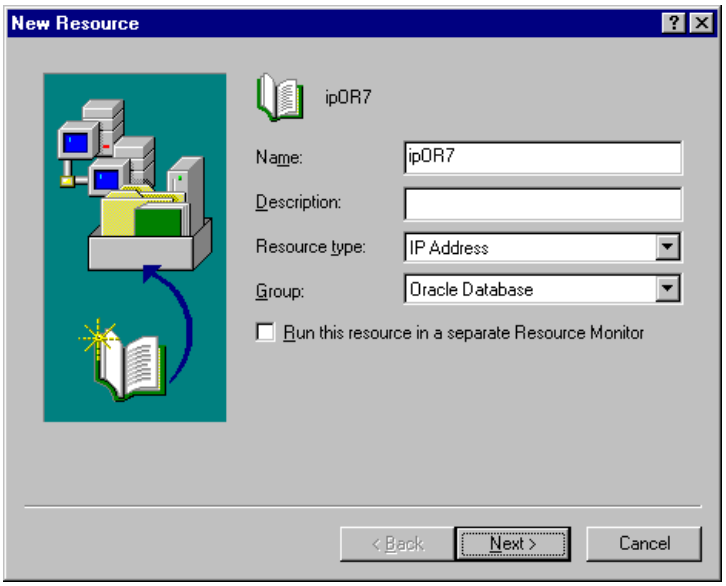

Figure 5-1. New Resource Wizard

2. Enter both nodes as possible hosts for the resource. Click **Next**.

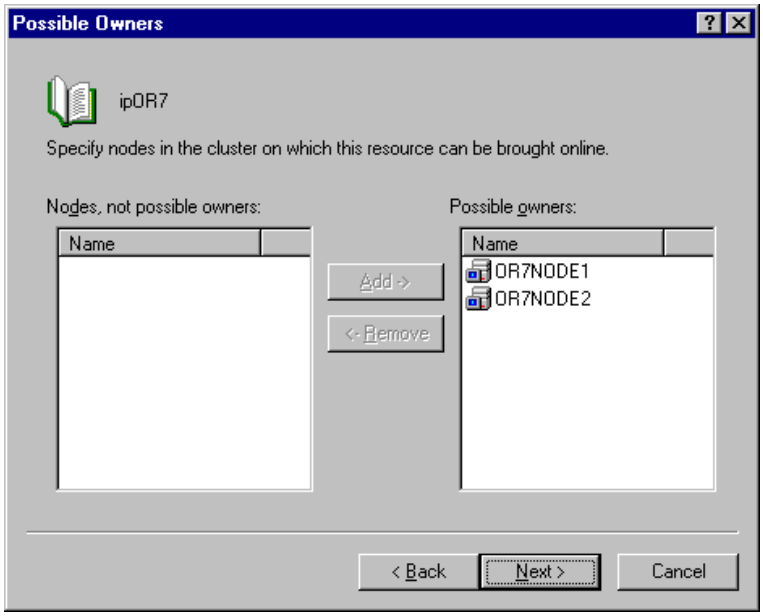

Figure 5-2. Possible Owners Dialog Box

3. There are no dependencies. Click **Next** to continue.

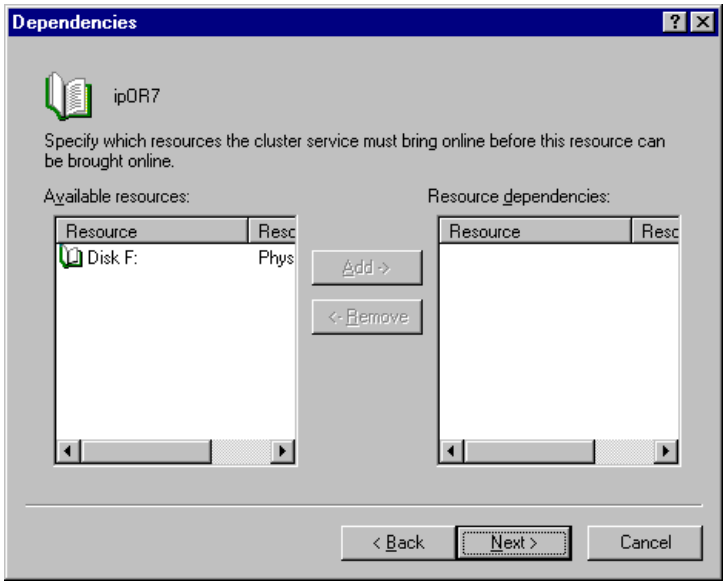

Figure 5-3. Dependencies Dialog Box

4. Enter the IP Address **110.12.120.121** with a subnet mask **255.255.0.0**. Make sure to select the network that will be used by the clients to connect to the cluster. Click **Finish**.

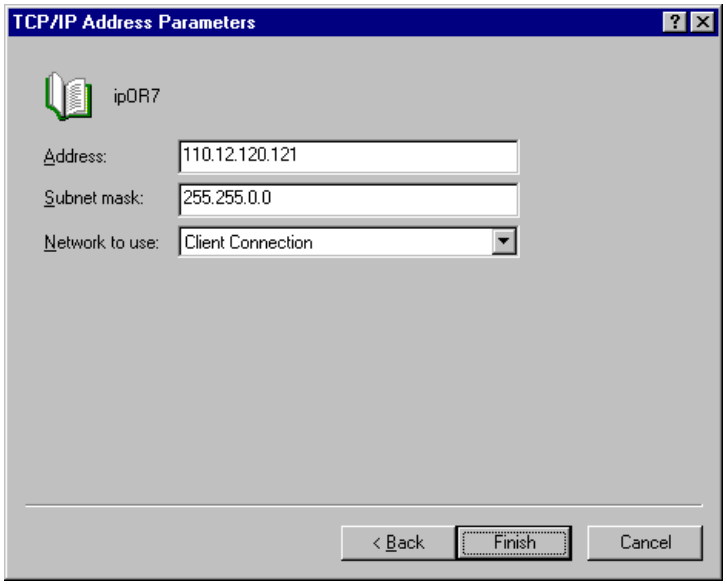

Figure 5-4. IP Address Dialog Box

Your Cluster Administrator should appear similar to the following:

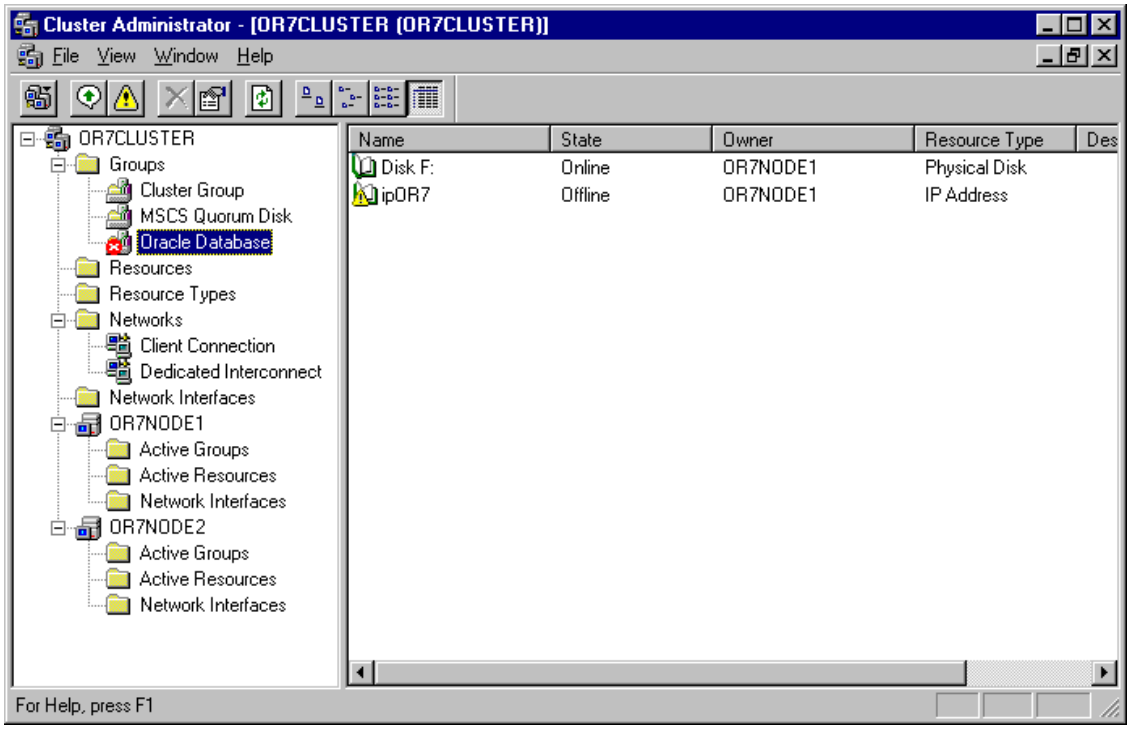

Figure 5-5. Cluster Administrator

Network Name Resource

Use the following information to setup your Network Name resource.

Table 5-3
Oracle Database (Network Name)

Resource Name	NMOR7
Resource Type	Network Name
Network Name	OR7GROUP
Dependencies	IPOR7

1. Enter **NMOR7** as the resource name. Select **Network Name** as the resource type. Click **Next**.

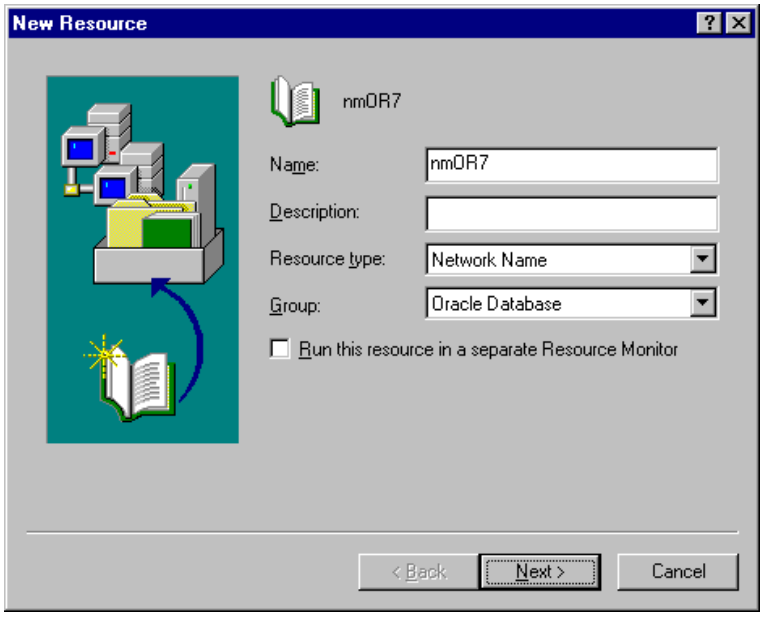

Figure 5-6. New Resource Dialog Box

2. Enter both nodes as possible hosts for the resource. Click **Next**.

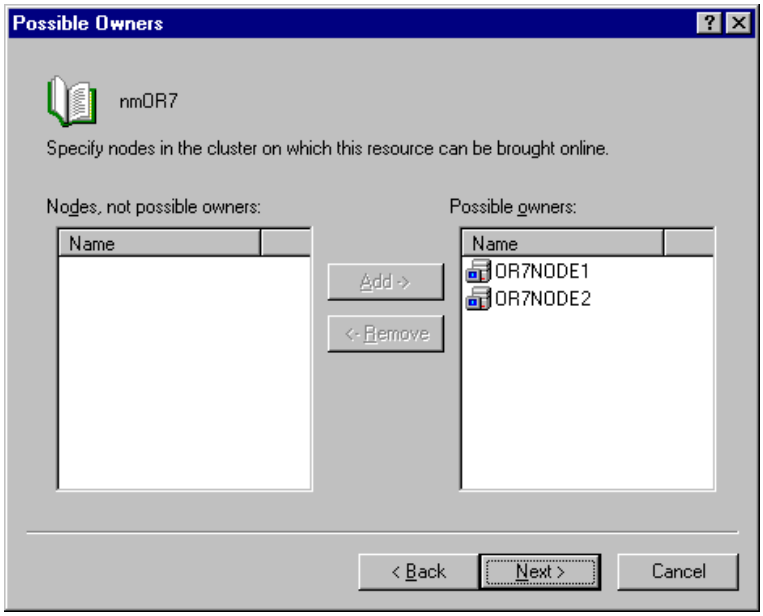

Figure 5-7. Possible Owners Dialog Box

3. Select the **IPOR7** resource as a dependency. Click **Next**.

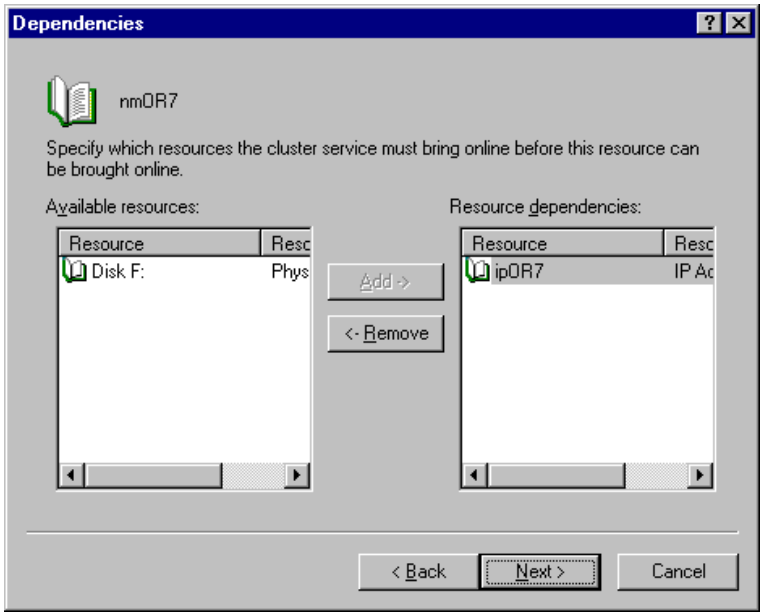

Figure 5-8. Dependencies Dialog Box

4. Enter **OR7GROUP** as the host name used by the clients to connect to the cluster. Click **Finish**.

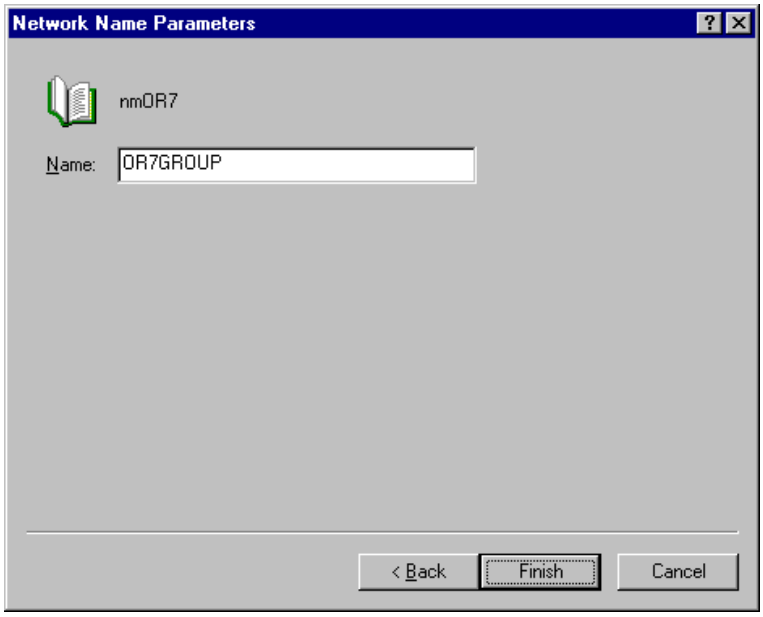

Figure 5-9. Network Name Parameters Dialog Box

5. Bring the **Oracle Database** resource group **On-line**. Your Cluster Administrator should look similar to the following:

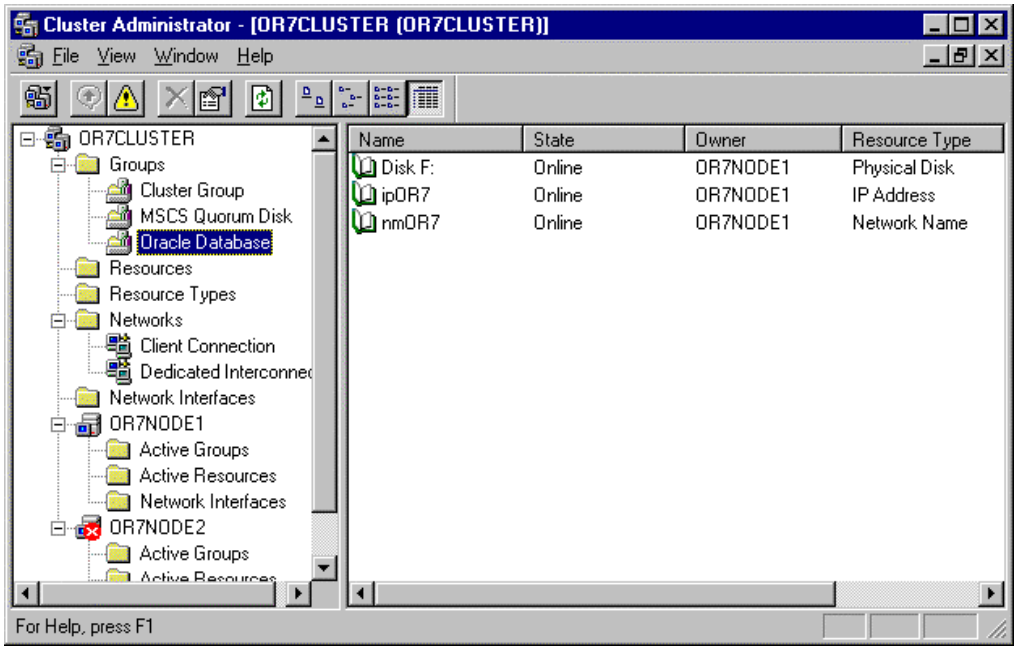

Figure 5-10. Cluster Administrator

Chapter 6
Oracle Installation

The following sections provide a step-by-step installation guide for Oracle7 Workgroup Server 7.3.3 for both cluster nodes. You will learn how to install the Oracle client products on the workstation. If you need additional information on Oracle installation, please refer to the documentation included in the Oracle Server CD.

Installing Oracle Workgroup Server on both Servers

Important: This installation must be done on both servers.

1. Insert the CD into the first server, and click **Yes** when prompted if you want to install Oracle.

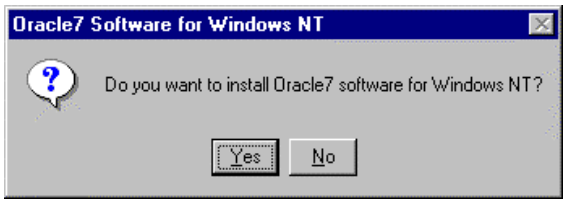

Figure 6-1. Oracle Installation Message

2. Select the language in which you want to run Oracle.

Figure 6-2. Language Selection Dialog Box

3. Enter your Company name and a **local** home directory, such as C:\ORANT.

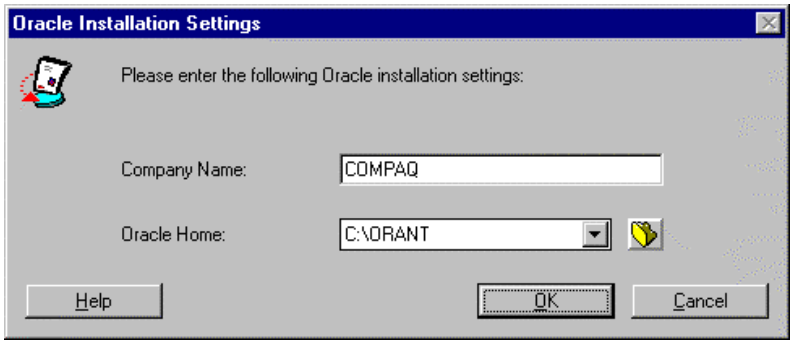

Figure 6-3. Installation Settings Dialog Box

4. Click **Yes** when prompted to restart the installation and update the path, otherwise proceed to Step 9.

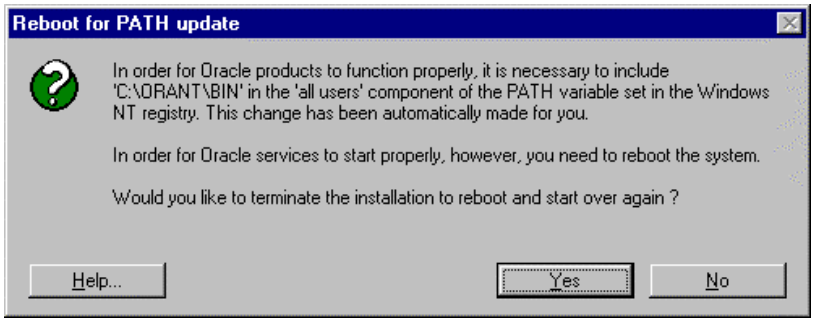

Figure 6-4. Installation Message

5. Installation is terminated. Click **OK**.

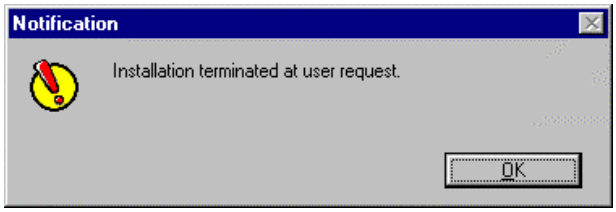

Figure 6-5. Notification

6. Run **D:\Setup.Exe** from the Start → Run Menu.

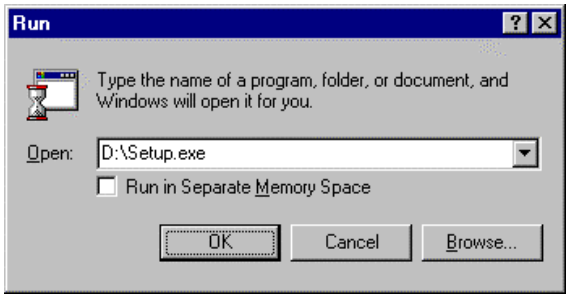

Figure 6-6. Run Application Dialog Box

7. Select the language in which you want to run Oracle.

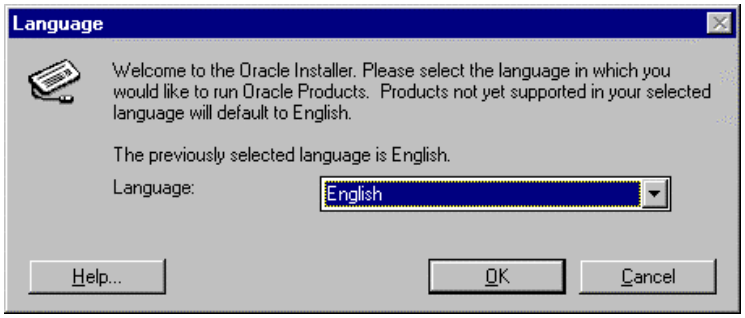

Figure 6-7. Language Selection Dialog Box

8. Leave the company name and home directory as shown if they are the same as previously entered, otherwise enter the correct information.

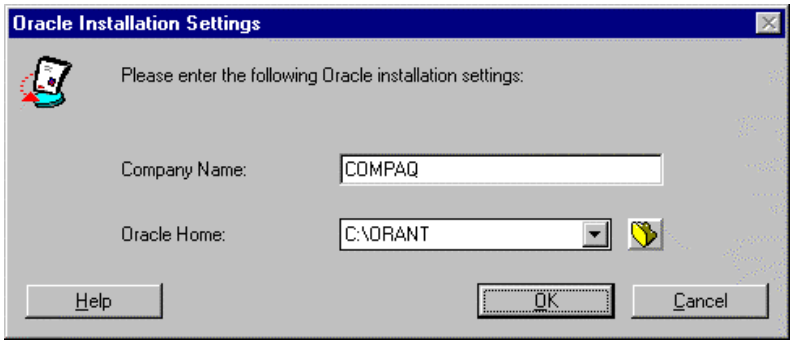

Figure 6-8. Installation Settings Dialog Box

9. Select Oracle7 Workgroup Server Products, click **OK**.

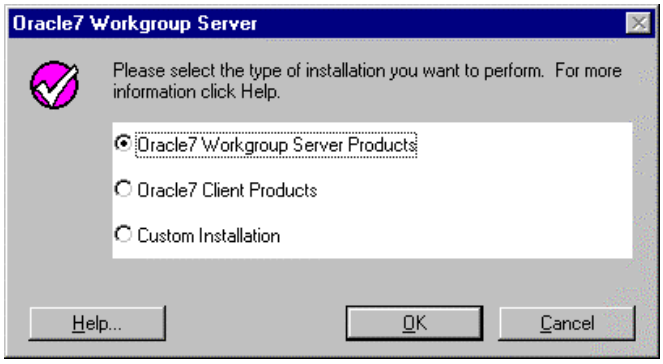

Figure 6-9. Installation Type Dialog Box

10. Select **CD-ROM**, to leave the documentation on the CD-ROM, otherwise select the **Hardrive**. When selecting **Hardrive**, make sure there is enough free disk space for the documentation in addition to the Oracle files.

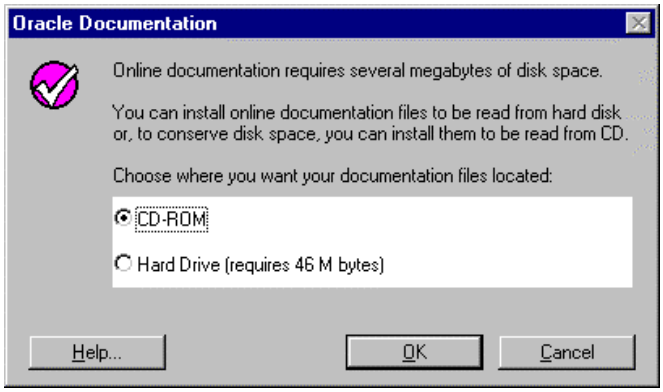

Figure 6-10. Oracle Documentation Selection Dialog Box

11. Oracle Installer now analyzes file dependencies and installs all necessary files in your home directory **C:\ORANT**. This takes some time, depending on the speed of your CD-ROM drive.

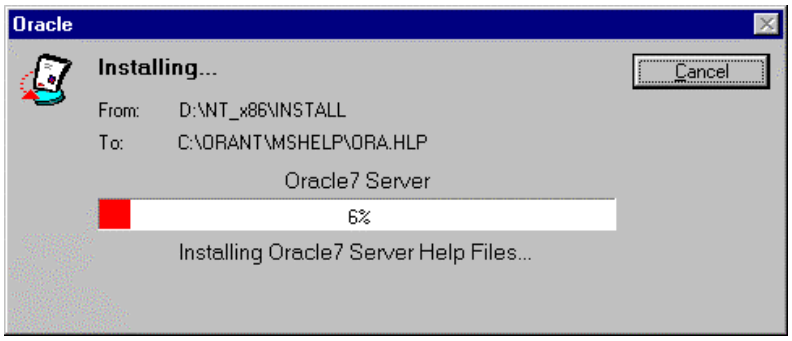

Figure 6-11. Installation Progress Dialog Box

12. Please read the message after the installation. Review the Release Notes in the Oracle CD. ODBC configuration is not discussed in this document since no ODBC connection is used in this cluster.

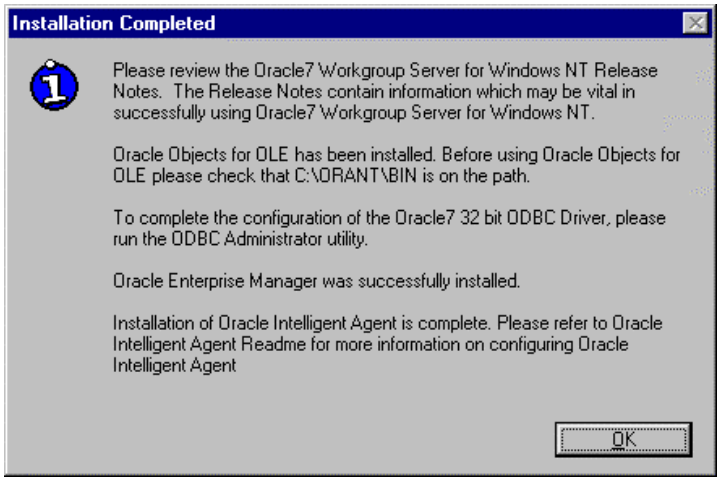

Figure 6-12. Installation Complete Message

Important: Repeat this installation on the second server.

Installing Oracle Client Products in the Workstation

1. Insert the CD into the first server, and click **Yes** when prompted if you want to install Oracle.

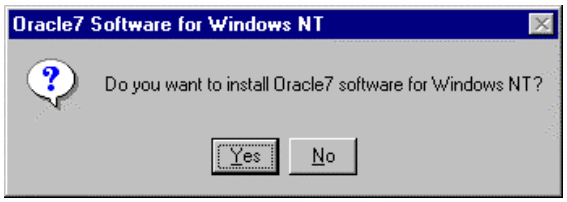

Figure 6-13. Installation Message

2. Select the language in which you want to run Oracle.

Figure 6-14. Language Selection Dialog Box

3. Enter your Company name and a **local** home directory, such as **C:\ORANT**.

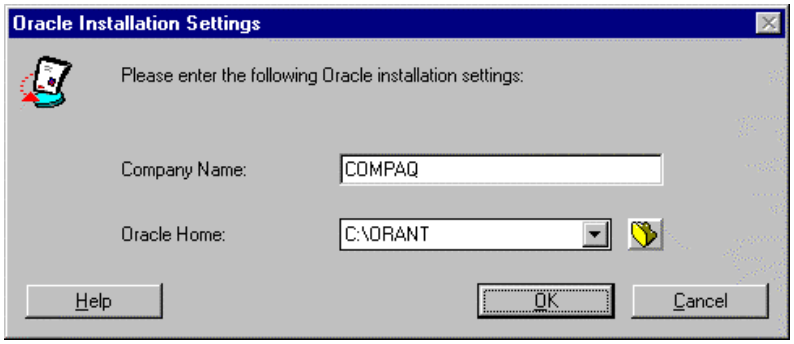

Figure 6-15. Installation Settings Dialog Box

4. Click **Yes** when prompted to restart the installation and update the path, otherwise proceed to Step 9.

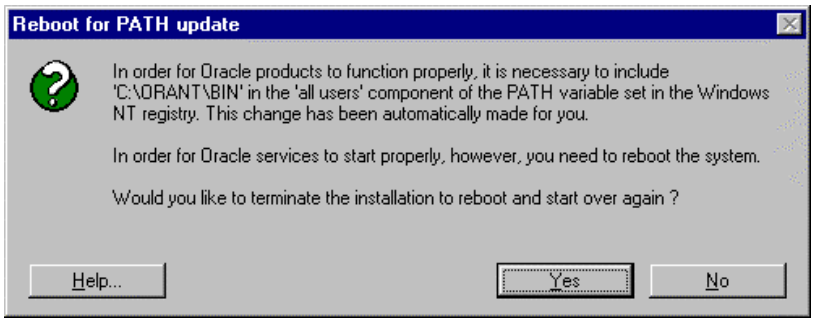

Figure 6-16. Reboot for PATH Update Dialog Box

5. Installation is terminated. Click **OK**.

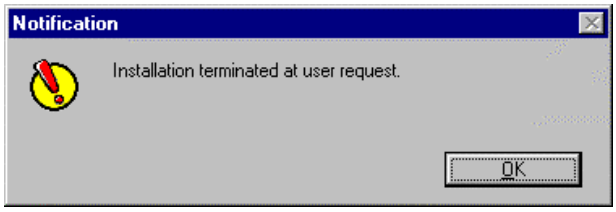

Figure 6-17. Notification Message

6. Run **D:\Setup.Exe** from the Start → Run Menu.

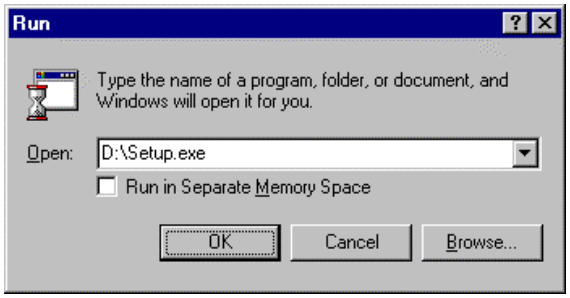

Figure 6-18. Run Program Dialog Box

7. Select the language in which you want to run Oracle.

Figure 6-19. Select Language Dialog Box

8. Leave the company name and home directory as shown if the same as previously entered, otherwise enter the correct information.

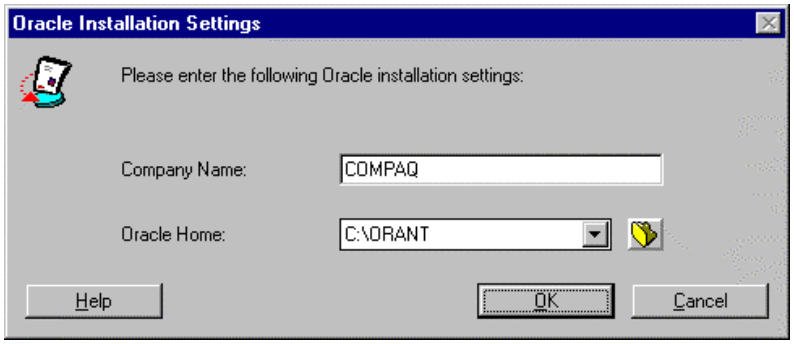

Figure 6-20. Installation Settings Dialog Box

9. Select Oracle7 Client Products, click OK.

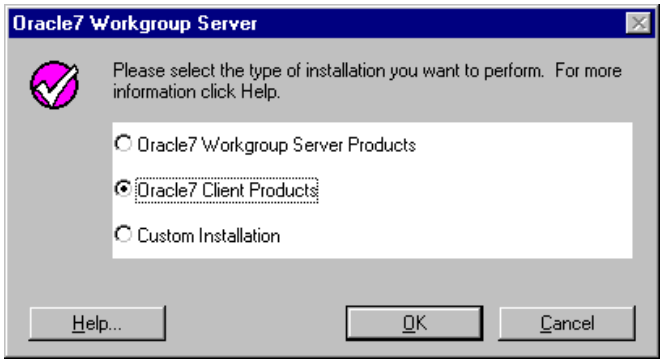

Figure 6-21. Workgroup Server Installation Type Selection Dialog Box

10. Select Database Administrator as the type of client installation. Click **OK**.

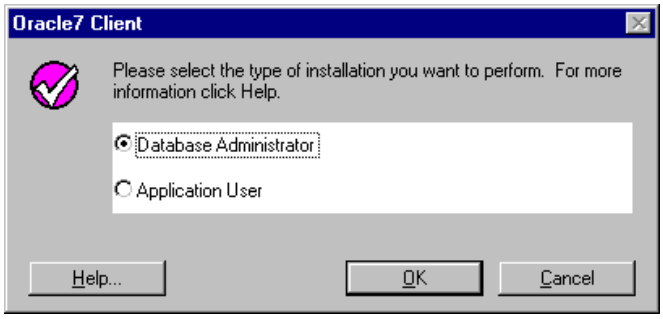

Figure 22. Client Installation Type Selection Dialog Box

11. Select **CD-ROM**, to leave the documentation on the CD-ROM, otherwise select the **Hardrive**. If selecting **Hardrive**, make sure there is enough free disk space for the documentation in addition to the Oracle files.

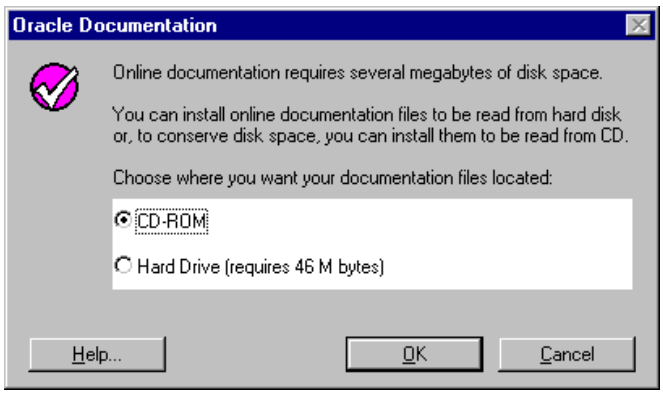

Figure 6-23. Documentation Dialog Box

12. Oracle Installer now analyzes file dependencies and installs all necessary files in your home directory **C:\ORANT**. This takes some time, depending on the speed of your CD-ROM drive.

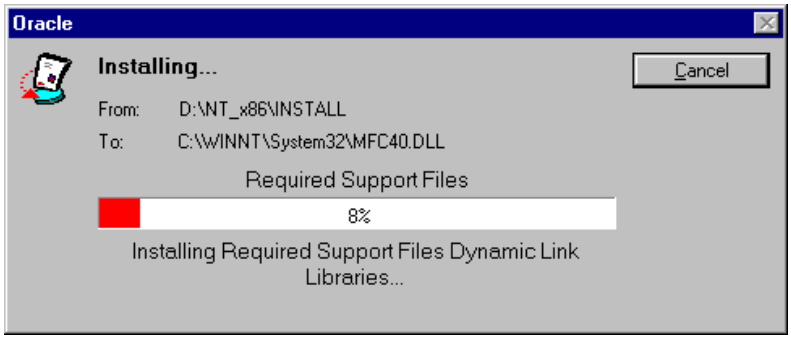

Figure 6-24. Installation Progress Dialog Box

13. Please read the message after the installation. Review the Release Notes in the Oracle CD. ODBC configuration is not discussed in this document since no ODBC connection is used in this cluster.

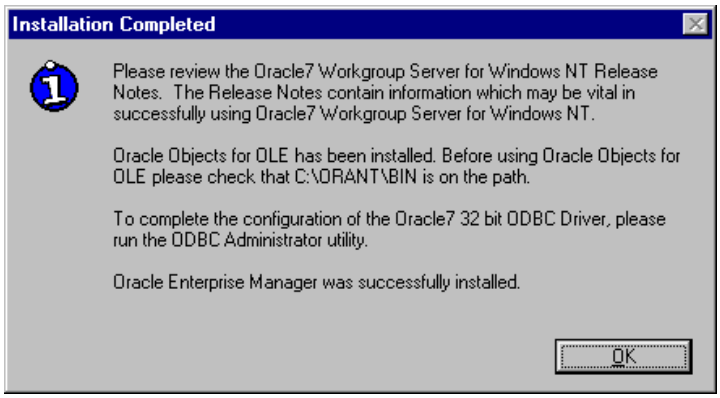

Figure 6-25. Installation Complete Dialog Box

Chapter 7
Oracle Cluster Configuration

Deleting the Starter Database

1. Run the Oracle Instance Manager from the **Start → Programs → Oracle for Windows NT** Menu.
2. The Server ID **ORCL** is used for the cluster so a new initialization INITxxxx.ORA file does not have to be created. Click **Delete** to remove the ORCL instance and any related files.

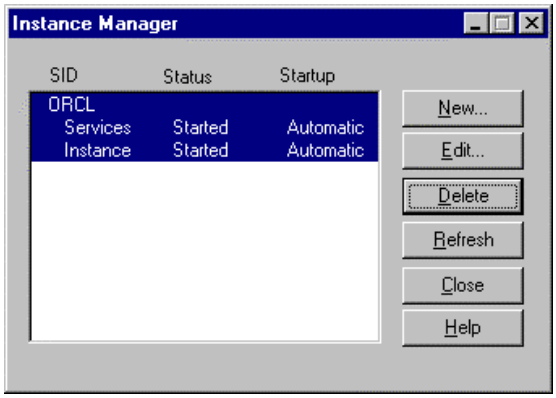

Figure 7-1. Instance Manager Dialog Box

3. A warning message displays. Click **Yes** to continue deleting the local ORCL instance.

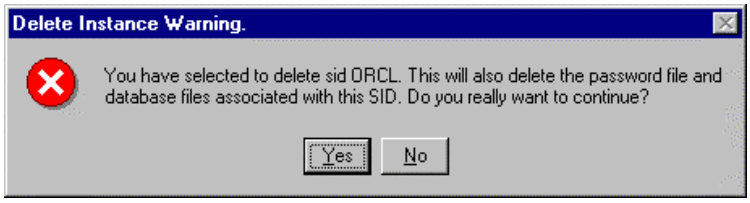

Figure 7-2. Warning Message

4. Enter the Database name **ORACLE**, and the default DBA authorization password **ORACLE**, then click **OK**.

Figure 7-3. Database Information Dialog Box

The Oracle Instance Manager should be empty now. Exit the Oracle Instance Manager.

Note: If this procedure does not delete the starter database, you must create a new Server ID and Instance Name for this cluster (for example, ORA7 for your Server ID, and ORACLE7 for your Instance Name).

Updating the HOSTS Files

To establish a SQL*Net connection in a client that supports TCP/IP, a way must be provided for the client and the servers to see each server and cluster participating in the network.

One way of doing this is to use the **C:\WINNT\SYSTEM32\DRIVERS\ETC\HOSTS** file. Add the IP addresses of the cluster and the servers to the **HOSTS** file of each machine. Add the following IP addresses in the **HOSTS** file of both servers and the client workstation:

```
110.12.120.111  or7node1
110.12.120.112  or7node2
110.12.120.120  or7cluster
110.12.120.121  or7group
```

The cluster IP address for **OR7CLUSTER** is the IP address entered when Cluster Server was installed, and is used for cluster administration. The cluster IP address for **OR7GROUP** is the IP address that is used by the client workstations to connect to the Oracle cluster.

Note: Make sure that the **HOSTS** file is saved without any file extension. You should replicate this file to both servers and the workstation.

Creating Subfolders in the Shared Drive

Make sure that both servers can access the shared storage **Disk F:**.

From the Cluster node that currently owns Disk F:, create a folder **DBORCL** with the subfolders **DATA**, **LOG**, and **TRACE**. This is the location of the data files, log files and trace files of your database.

```
DISK F:
|
|___ DBORCL
       |
       |___ DATA (data files, control files)
       |
       |___ LOG (log files)
       |
       |___ TRACE (trace files)
```

Configure SQL*Net

This step should be done on all machines running client programs, such as Server Manager SVRMGR23.EXE or SQL*Plus. For this cluster, **both servers and the workstation** are **configured**. The procedure is the same for all of them.

1. From the Start → Programs → Oracle for Windows NT Menu, run the **SQL*Net Easy Configuration** program.

2. Select View Configuration Information, then click **OK**.

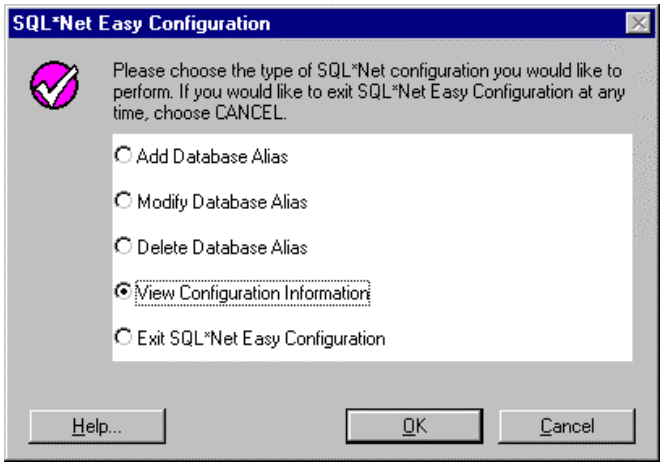

Figure 7-4. Configuration Type Dialog Box

3. Make sure that you are creating a new SQL*Net database alias for the Oracle instance **ORACLE**. **ORACLE** should not be in the list. Click **Back** to return to the Main Menu.

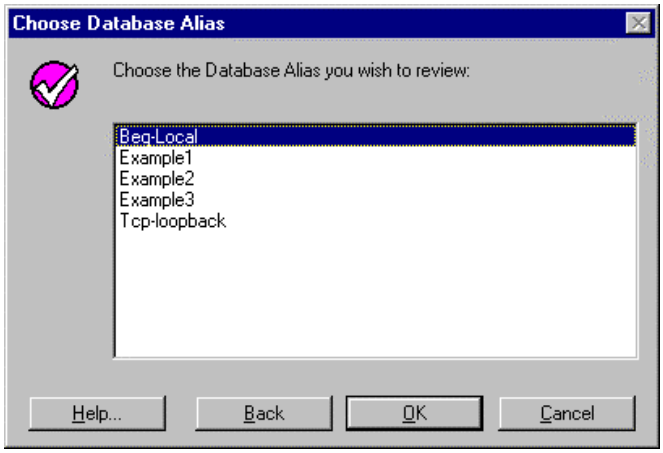

Figure 7-5. Database Alias Selection Dialog Box

4. From the main menu, select the **Add Database Alias**, then click **OK**. Enter **ORACLE** as the Database Alias. This is the name the clients use to access this database instance. Click **OK**.

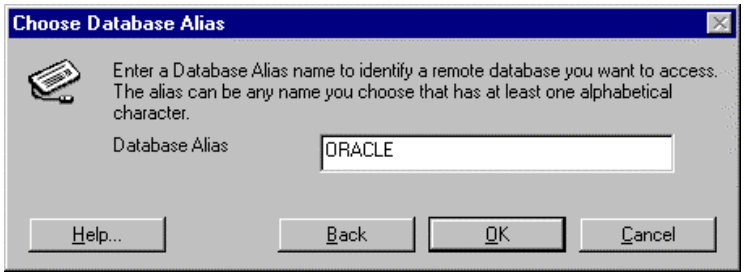

Figure 7-6. Database Alias Selection Screen

5. Select **TCP/IP** as the network protocol to be used to connect to the database, then click **OK**.

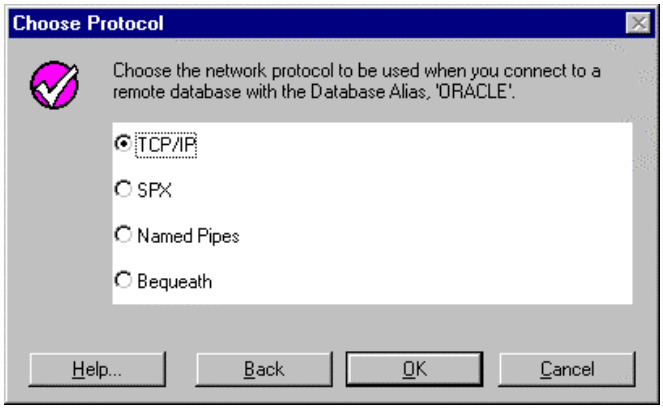

Figure 7-7. Protocol Selection Dialog Box

6. Enter the Oracle Cluster group name **OR7GROUP** and the Server ID **ORCL** for this database. Click **OK**.

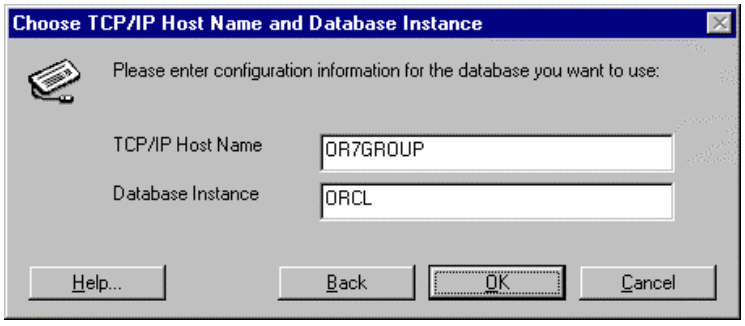

Figure 7-8. TCP/IP Host Name and Database Instance Selection Dialog Box

7. Verify that all information is correct, then click **OK**.

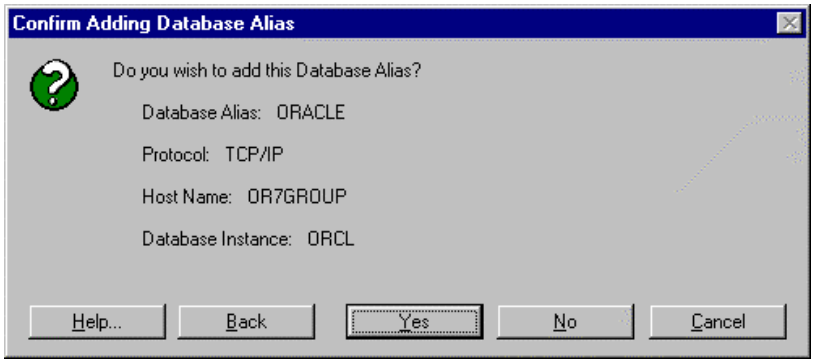

Figure 7-9. Confirmation Dialog Box

8. Exit the SQL*Net Configuration, then click **OK** to confirm.

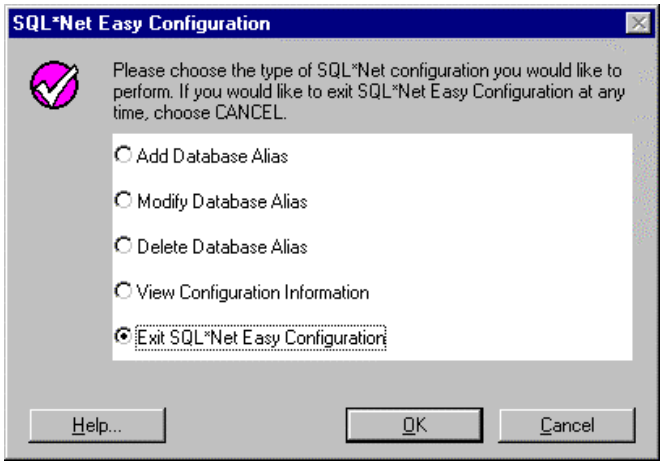

Figure 7-10. Exit Configuration Dialog Box

Important: These steps should be done on both servers and the workstation.

Modifying the Listener File in Node 1

1. Create a backup of the **C:\ORANT\NETWORK\ADMIN\LISTENER.ORA** file.
2. From the Cluster Administrator, make sure that all resources in the **ORA7GROUP** are on-line in **OR7NODE1**.
3. From the **Control Panel → Services** icon, stop the **OracleTNSListener** Services.

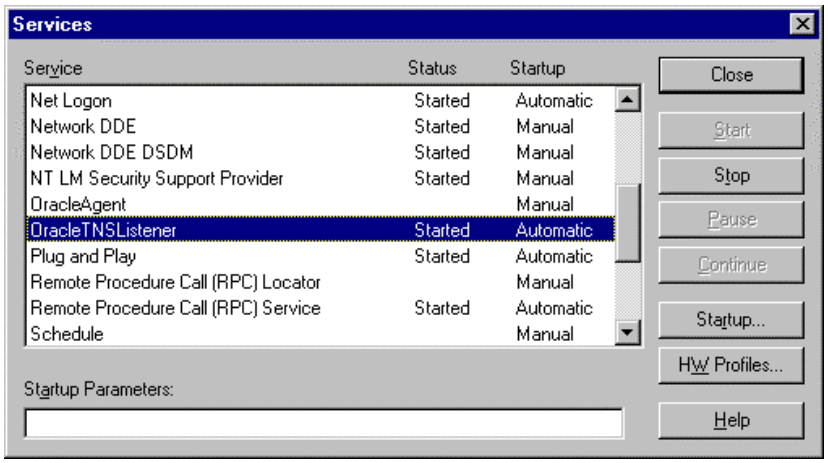

Figure 7-11. Services Dialog Box

4. Change the **Startup** option of the **OracleTNSListener** from **Automatic** to **Manual**.

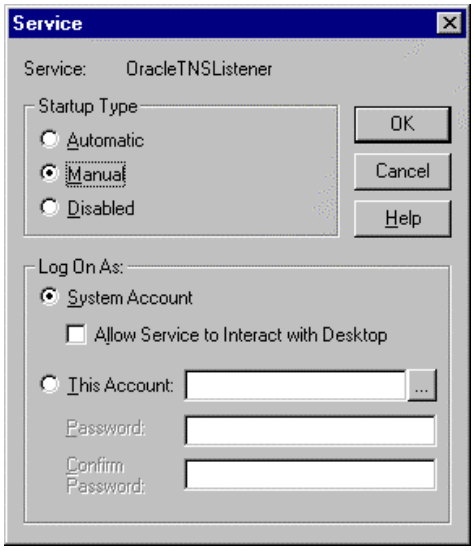

Figure 7-12. OracleTNSListener Service Dialog Box

5. Using Notepad, open the **C:\ORANT\NETWORK\ADMIN\LISTENER.ORA** file.

6. Place all references to **OR7NODE1** into **OR7GROUP**. Your updated listener should look similar to this:

```
################
# Filename......: listener.ora
# Node..........: local.world
# Date..........: 24-MAY-97 13:23:20
################
LISTENER =
  (ADDRESS_LIST =
        (ADDRESS=
          (PROTOCOL= IPC)
          (KEY= oracle.world)
        )
        (ADDRESS=
          (PROTOCOL= IPC)
          (KEY= ORCL)
```

(continued)

```
        )
        (ADDRESS=
          (COMMUNITY= NMP.world)
          (PROTOCOL= NMP)
(SERVER= OR7GROUP)
          (PIPE= ORAPIPE)
        )
        (ADDRESS=
          (COMMUNITY= TCP.world)
          (Host = OR7GROUP)
          (PROTOCOL= TCP)
          (Port= 1521)
        )
        (ADDRESS=
          (COMMUNITY= TCP.world)
          (Host = OR7GROUP)
          (PROTOCOL= TCP)
          (Port= 1526)
        )
  )
STARTUP_WAIT_TIME_LISTENER = 0
CONNECT_TIMEOUT_LISTENER = 10
TRACE_LEVEL_LISTENER = ADMIN
SID_LIST_LISTENER =
  (SID_LIST =
    (SID_DESC =
      (SID_NAME = ORCL)
    )
  )
PASSWORDS_LISTENER = (oracle)
```

7. From the **Control Panel → Services** icon, start the OracleTNSListener Service.

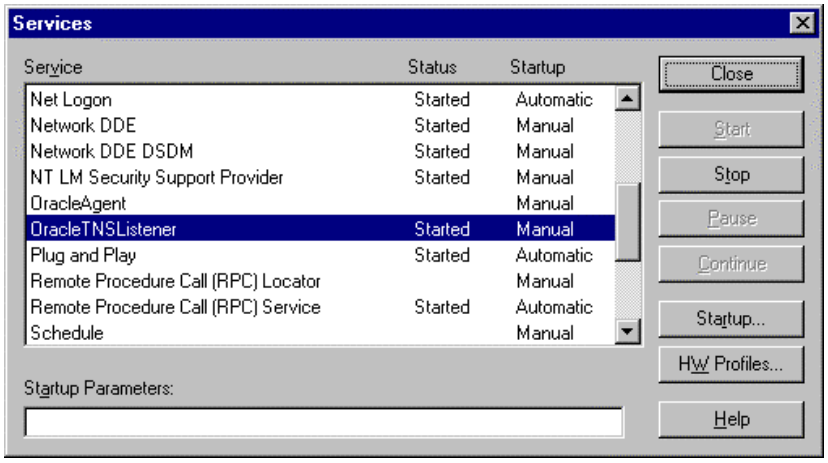

Figure 7-13. Services Dialog Box

8. From an MS-DOS command box, test the new TNS (Transparent Network Substrate) Listener by using the TNSPING.EXE command.

```
C:\>tnsping oracle

TNS Ping Utility for 32-bit Windows: Version 2.3.3.0.0 -
Production on 25-NOV-97 16:11:25

Copyright (c) Oracle Corporation 1995.  All rights
reserved.

Attempting to contact
(ADDRESS=(COMMUNITY=tcp.world)(PROTOCOL=TCP)(Host=OR7GROU
P
)(Port=1521))
OK (30 msec)

C:\>
```

Modifying the Listener File in Node 2

1. Create a backup of the C:\ORANT\NETWORK\ADMIN\LISTENER.ORA file.
2. From the Cluster Administrator, make sure that all resources in the **ORA7GROUP** is On-line in **OR7NODE2**.
3. From the **Control Panel → Services** icon, stop the **OracleTNSListener** Services.

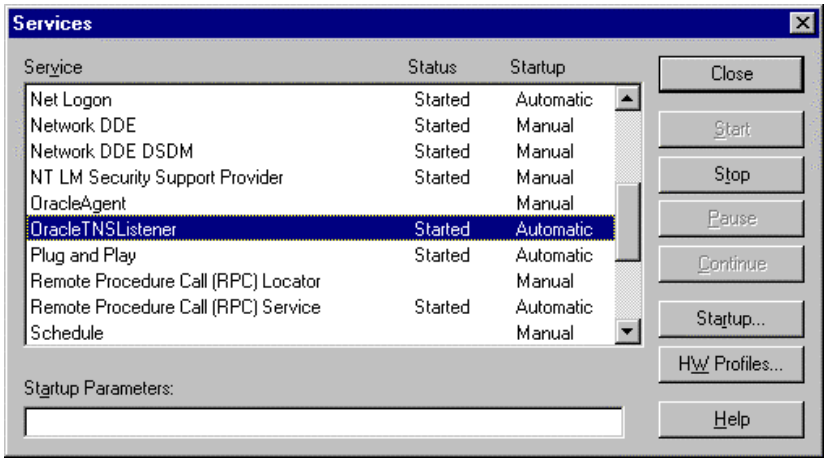

Figure 7-14. OracleTNSListener Services Dialog Box

4. Change the **Startup** option of the **OracleTNSListener** from **Automatic** to **Manual**.

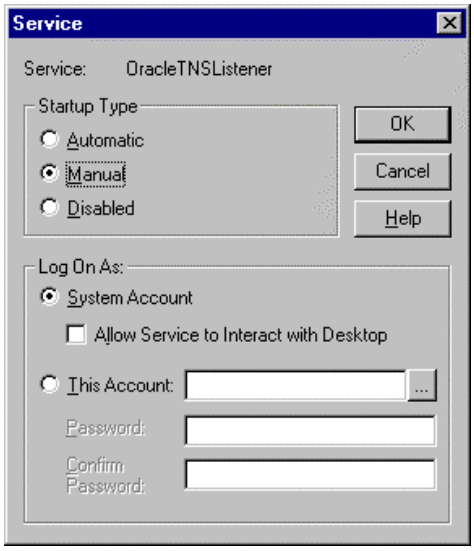

Figure 7-15. Startup Options Dialog Box

5. Using Notepad, open the **C:\ORANT\NETWORK\ADMIN\LISTENER.ORA** file.

6. Place all references to **OR7NODE1** into **OR7GROUP**. Your updated listener should look similar to:

```
#################
# Filename......: listener.ora
# Node..........: local.world
# Date..........: 24-MAY-97 13:23:20
#################
LISTENER =
  (ADDRESS_LIST =
        (ADDRESS=
          (PROTOCOL= IPC)
          (KEY= oracle.world)
        )
```

(continued)

```
        (ADDRESS=
          (PROTOCOL= IPC)
          (KEY= ORCL)
        )
        (ADDRESS=
          (COMMUNITY= NMP.world)
          (PROTOCOL= NMP)
          (SERVER= OR7GROUP)
          (PIPE= ORAPIPE)
        )
        (ADDRESS=
          (COMMUNITY= TCP.world)
          (Host = OR7GROUP)
          (PROTOCOL= TCP)
          (Port= 1521)
        )
        (ADDRESS=
          (COMMUNITY= TCP.world)
          (Host = OR7GROUP)
          (PROTOCOL= TCP)
          (Port= 1526)
        )
  )
STARTUP_WAIT_TIME_LISTENER = 0
CONNECT_TIMEOUT_LISTENER = 10
TRACE_LEVEL_LISTENER = ADMIN
SID_LIST_LISTENER =
  (SID_LIST =
    (SID_DESC =
      (SID_NAME = ORCL)
    )
  )
PASSWORDS_LISTENER = (oracle)
```

7. From the **Control Panel → Services** icon, start the OracleTNSListener Service.

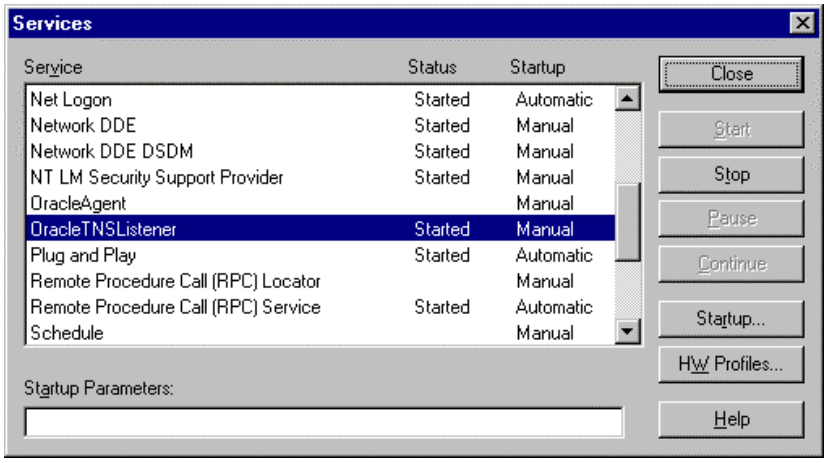

Figure 7-16. OracleTNSListener Services Dialog Box

8. From an MS-DOS command box, test the new TNS (Transparent Network Substrate) Listener by using the TNSPING.EXE command.

```
C:\>tnsping oracle

TNS Ping Utility for 32-bit Windows: Version 2.3.3.0.0 -
Production on 25-NOV-97 16:11:25

Copyright (c) Oracle Corporation 1995.  All rights
reserved.

Attempting to contact
(ADDRESS=(COMMUNITY=tcp.world)(PROTOCOL=TCP)(Host=OR7GROU
P
)(Port=1521))
OK (30 msec)

C:\>
```

Updating the Initialization File

Using the notepad, modify the initialization file **C:\ORANT\DATABASE\INITORCL.ORA**.

1. Change the location of the control files to point to your shared folder:

```
control_files = (F:\DBORCL\DATA\CTL1ORCL.ORA,
                 F:\DBORCL\DATA\CTL2ORCL.ORA)
```

2. Modify the following lines as shown:

```
background_dump_dest = F:\DBORCL\TRACE
user_dump_dest = F:\DBORCL\TRACE
```

Use **ORCL** as the four-character SID (Server ID) for the database. Use the same instance name **ORACLE** and the default listener name **LISTENER** for this cluster.

Creating the Oracle Database and Instance in Node 1

Perform all the following steps in Node 1.

1. Shutdown OR7NODE2 before creating the database in OR7NODE1. This ensures that all resources in the OR7GROUP are running in OR7NODE1, and prevents the resource group from moving to OR7NODE2 if a failure happens.

2. From the **Start → Programs → Oracle for Windows NT**, run the Oracle Instance Manager. Click **New** to create a new instance.

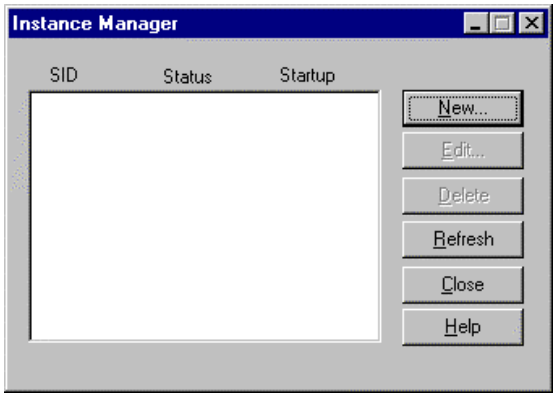

Figure 7-17. Instance Manager Dialog Box

3. Enter the following information. **Do not click OK** since this will create a database in the default location **C:\ORANT**. Click **Advanced** to change the location of the data files.

   SID = ORCL

   DBA Authorization Password = dba

   Confirm DBA Authorization Password = dba

   Max number of DBA = 5

   Startup after creation = Services, Instance

   Instance Startup Mode = Automatic

   Parameter Initialization Filename = **C:\ORANT\DATABASE\INITORCL.ORA**

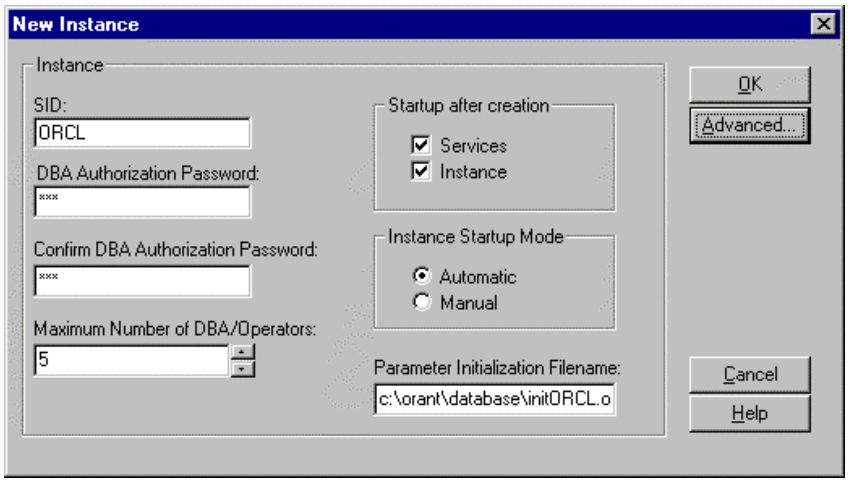

Figure 7-18. New Instance Dialog Box

4. Enter the following information, then click **OK**.

   Database = ORACLE

   Log File = “F:\DBORCL\DATA\LOG1ORCL.ORA” size 200K, “F:\DBORCL\DATA\LOG2ORCL.ORA” size 200K

   Data File = “F:\DBORCL\DATA\SYSORCL.ORA” size 10M

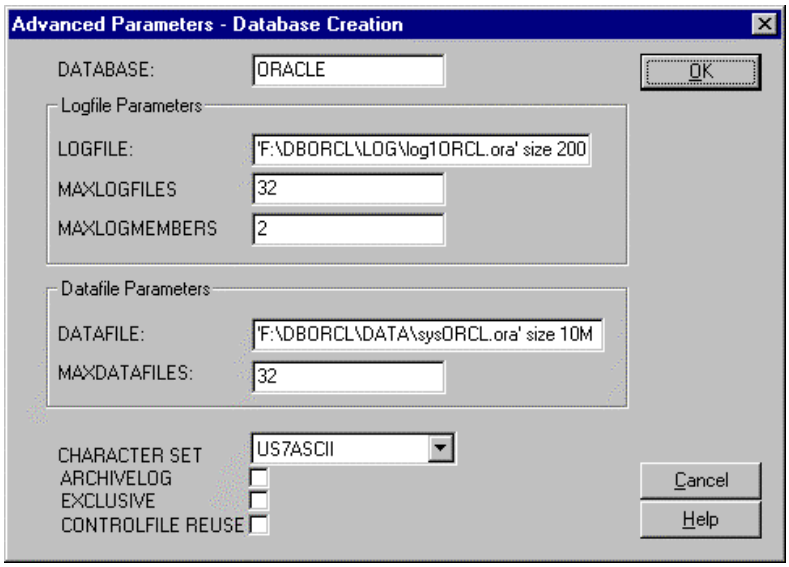

Figure 7-19. Advanced Parameters Dialog Box

5. The following message is displays, confirming that the password file is created. Click **OK**.

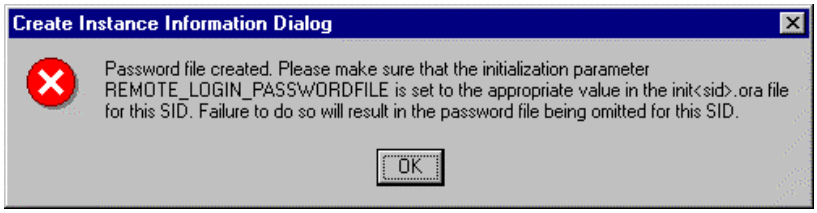

Figure 7-20. Password Creation Message

6. Another message displays, confirming that the database is created. Click **OK** to run a couple of scripts in the background.

Important: Do not restart or close Windows NT while this process is running.

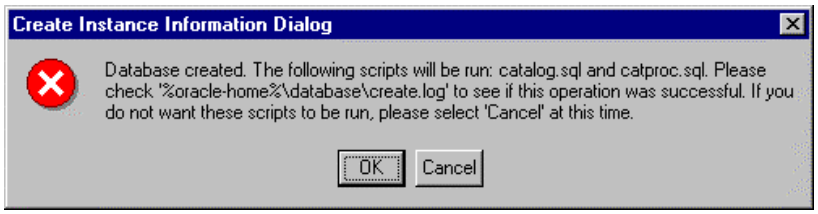

Figure 7-21. Database Creation Message

7. Monitor the SVRMGR23.EXE process from the Windows NT Task Manager. **This process may take a long time to complete. Please wait until the SVRMGR23.EXE process is completed** (i.e. until the SVRMGR23.EXE is no longer showing in the process list).

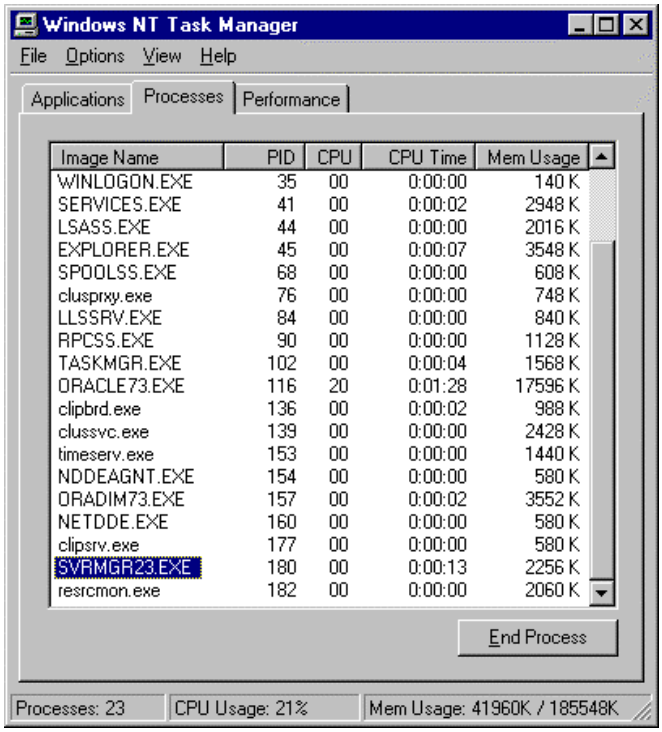

Figure 7-22. Windows NT Task Manager

8. From an MS-DOS command box, test the new database by running the Server Manager SVRMGR23.EXE.

```
C:\> svrmgr23

Oracle Server Manager Release 2.3.3.0.0 - Production

Copyright (c) Oracle Corporation 1994, 1995. All rights
reserved.

Oracle7 Workgroup Server Release 7.3.3.0.0 - Production
Release
With the distributed option
PL/SQL Release 2.3.3.0.0 - Production

SVRMGR> connect internal/dba
Connected to an idle instance.
SVRMGR> startup
ORACLE instance started.
Total System Global Area       7505584 bytes
Fixed Size                       35748 bytes
Variable Size                  7052044 bytes
Database Buffers                409600 bytes
Redo Buffers                      8192 bytes
Database mounted.
Database opened.
SVRMGR> shutdown immediate
Database closed.
Database dismounted.
ORACLE instance shut down.
SVRMGR> exit
```

9. Type **Exit** to close the Server Manager. Close the MS-DOS command box.

10. From the Control Panel → Services icon, change the **Startup** option of the following Services:

```
OracleServiceORCL       Manual
OracleStartORCL         Manual
OracleTNSListener       Manual
```

11. Shutdown OR7NODE1.

Creating the Instance in Node 2 for the existing Database

Perform the following steps on Node 2.

1. Shutdown OR7NODE1 before creating the instance in OR7NODE2 for the database in your shared drive. This ensures that all resources in the OR7GROUP are running in OR7NODE2, and prevents the resource group from moving to OR7NODE1 if a failure happens.
2. From an MS-DOS command box, enter the following command:

```
C:\>ORADIM73 -NEW -SID orcl -INTPWD dba -STARTMODE auto
-PFILE c:\orant\database\initorcl.ora
```

3. From an MS-DOS command box, test the new database by running the Server Manager SVRMGR23.EXE.

```
C:\> svrmgr23

Oracle Server Manager Release 2.3.3.0.0 - Production

Copyright (c) Oracle Corporation 1994, 1995. All rights
reserved.

Oracle7 Workgroup Server Release 7.3.3.0.0 - Production
Release
With the distributed option
PL/SQL Release 2.3.3.0.0 - Production

SVRMGR> connect internal/dba
Connected to an idle instance.
SVRMGR> startup
ORACLE instance started.
Total System Global Area       7505584 bytes
Fixed Size                       35748 bytes
Variable Size                  7052044 bytes
Database Buffers                409600 bytes
Redo Buffers                      8192 bytes
Database mounted.
Database opened.
SVRMGR> shutdown immediate
```

(continued)

```
Database closed.
Database dismounted.
ORACLE instance shut down.
SVRMGR> exit
```

4. Type **Exit** to close the Server Manager. Close the MS-DOS command box.

5. From the **Control Panel → Services** icon, change the **Startup** option of the following Services:

```
OracleServiceORCL       Manual
OracleStartORCL         Manual
OracleTNSListener       Manual
```

6. Shutdown OR7NODE2.

Chapter 8
Creating the Oracle Cluster Group Services

Start OR7NODE1 before creating the rest of the resources in the cluster group OR7GROUP.

Important: Perform all the following steps in OR7NODE1.

Creating the Cluster Resource Listener Service

Use the following information to setup your TNS Listener resource.

Table 8-1
Oracle Database (TNS Listener) IP Address Properties

Resource Name	LSNROR7
Resource Type	Generic Service
Service Name	OracleTNSListener
Dependencies	Disk F:, Network Name NMOR7

1. Enter **LSNROR7** as the resource name. Select **Generic Service** as the resource type. Click **Next**.

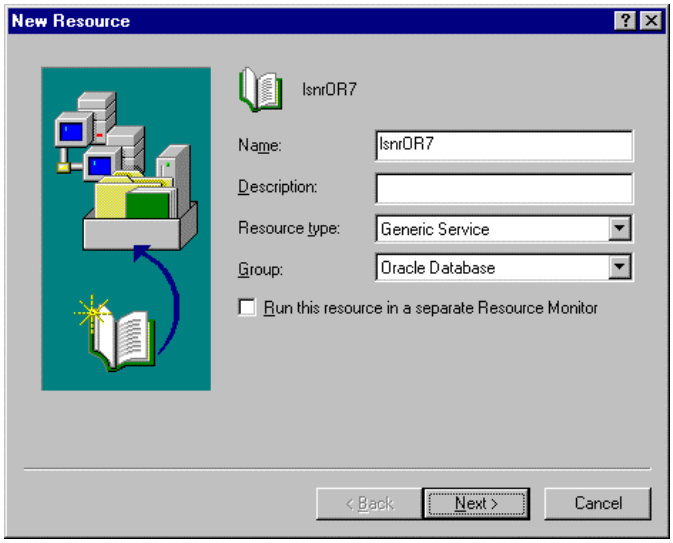

Figure 8-1. Create New Resource Dialog Box

2. Enter both nodes as possible hosts for the resource. Click **Next**.

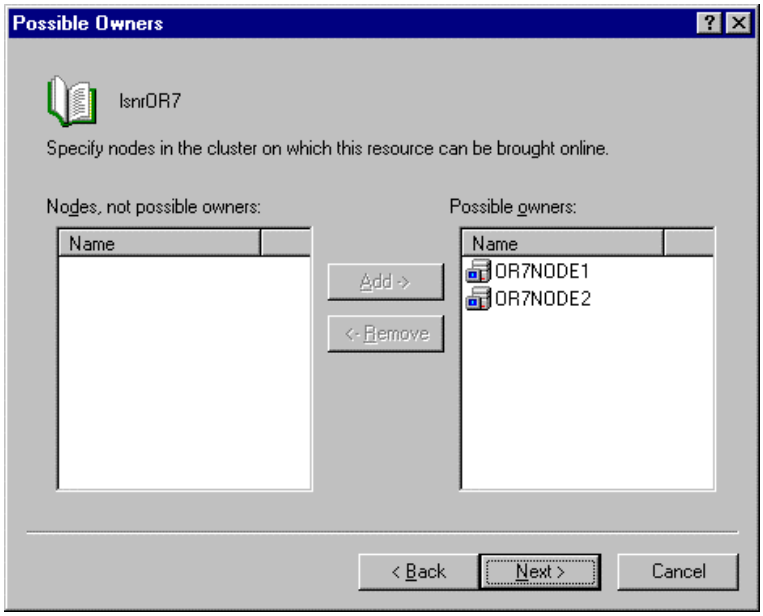

Figure 8-2. Possible Owners Dialog Box

3. Select **Disk F:** and **NMOR7** resources as dependencies. Click **Next**.

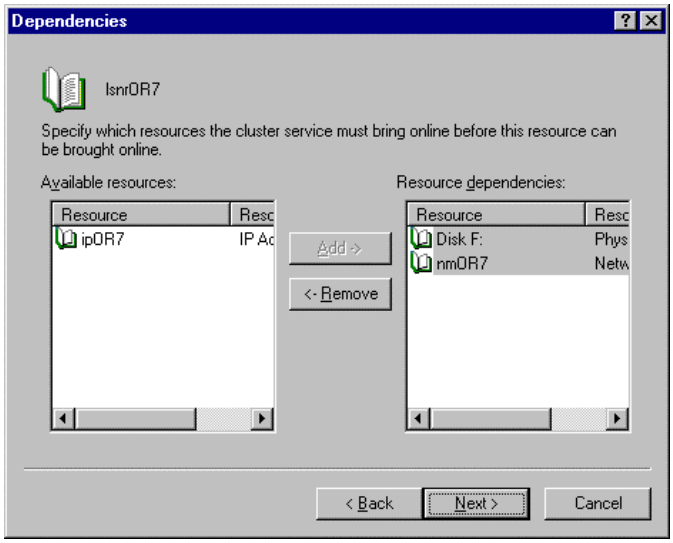

Figure 8-3. Dependencies Dialog Box

4. Enter **OracleTNSListener** as the service name to be used. Click **Next**.

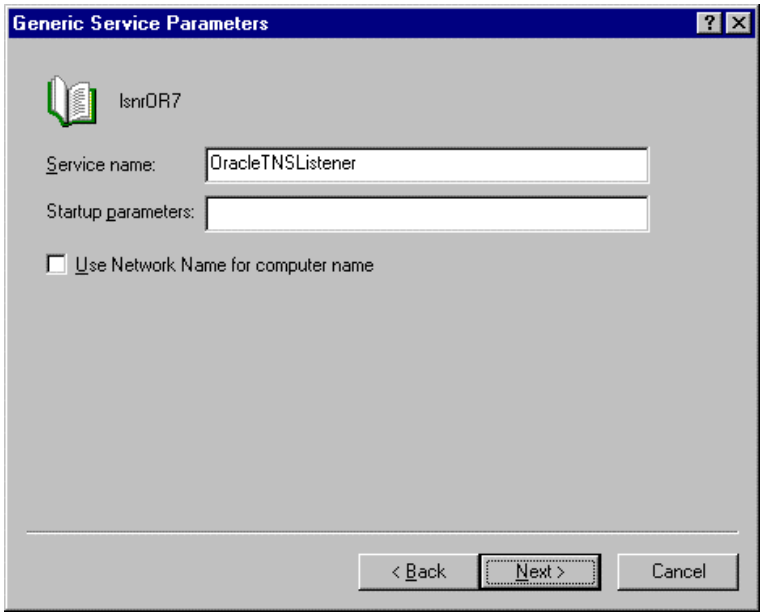

Figure 8-4. Parameters Dialog Box

5. Leave the Registry Replication blank. Click **Finish**.

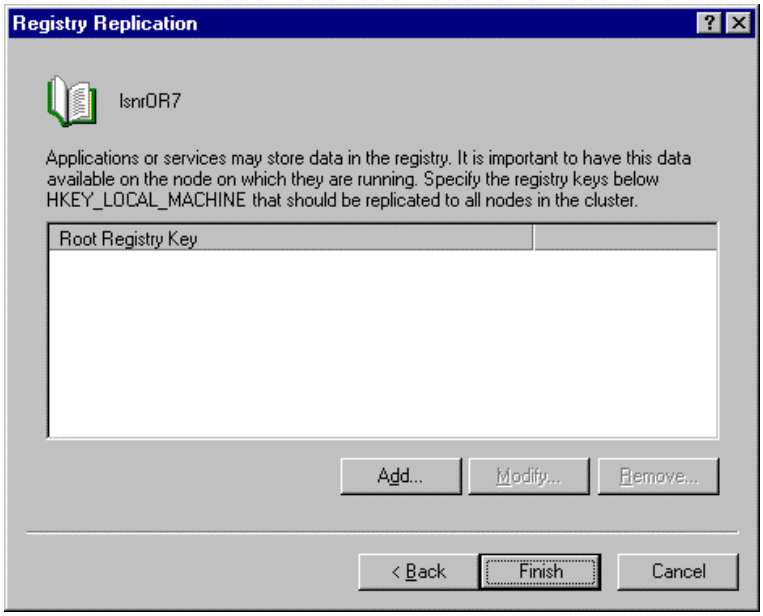

Figure 8-5. Registry Replication Dialog Box

6. Bring the Oracle Database resource group on-line. Your Cluster Administrator should look similar to the following:

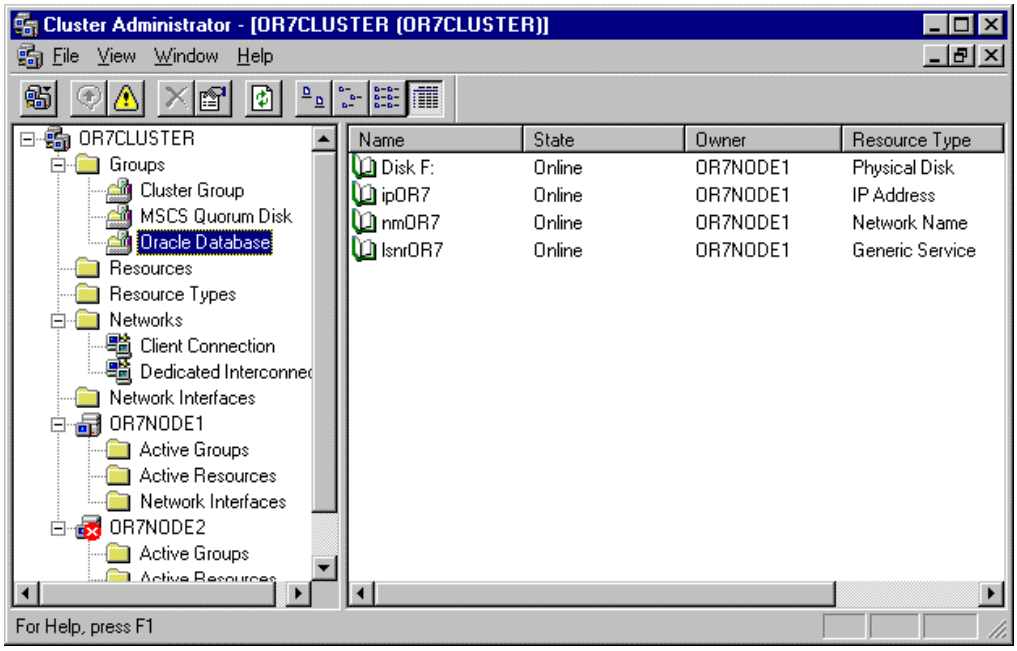

Figure 8-6. Cluster Administrator

Creating the Cluster Resource Oracle Server

Use the following information to setup your Oracle Server resource.

Table 8-2
Oracle Database (Oracle Instance) IP Address Properties

Resource Name	SRVCOR7
Resource Type	Generic Service
Service Name	OracleServiceORCL
Dependencies	LSNROR7

1. Enter **SRVCOR7** as the resource name. Select **Generic Service** as the resource type. Click **Next**.

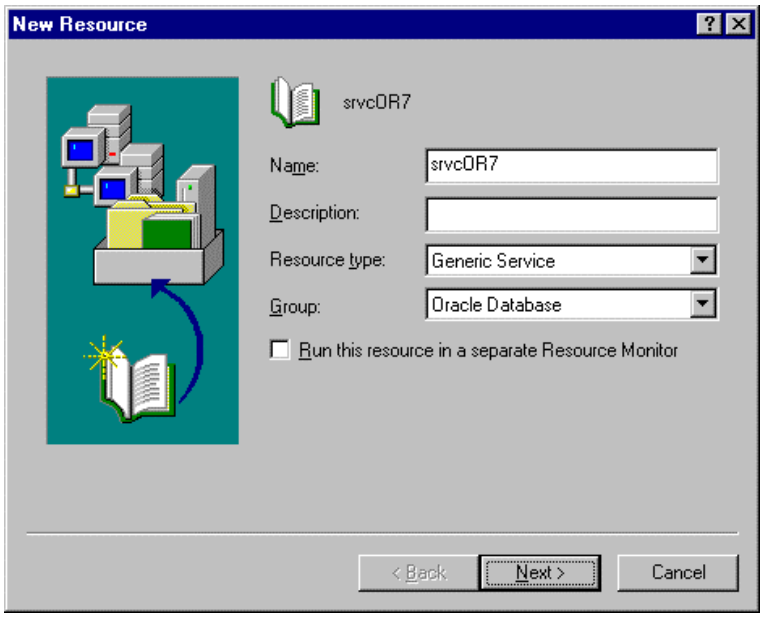

Figure 8-7. New Resource Dialog Box

2. Enter both nodes as possible hosts for the resource. Click **Next**.

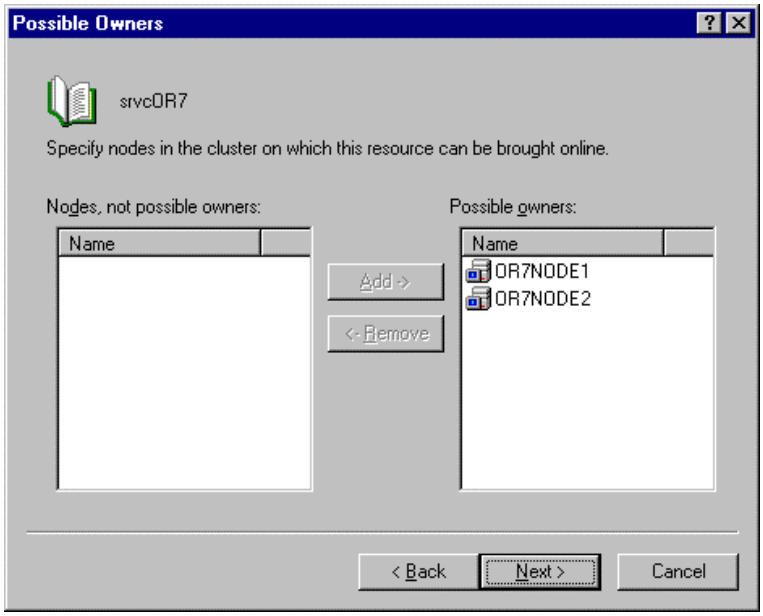

Figure 8-8. Possible Owners Dialog Box

3. Select **LSNROR7** resources as a dependency. Click **Next**.

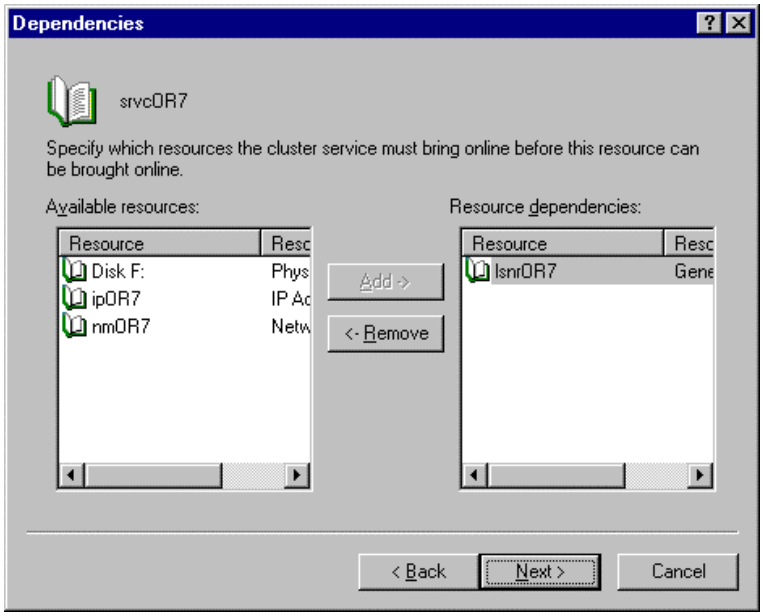

Figure 8-9. Dependencies Dialog Box

4. Enter **OracleServiceORCL** as the service name to be used. Click **Next**.

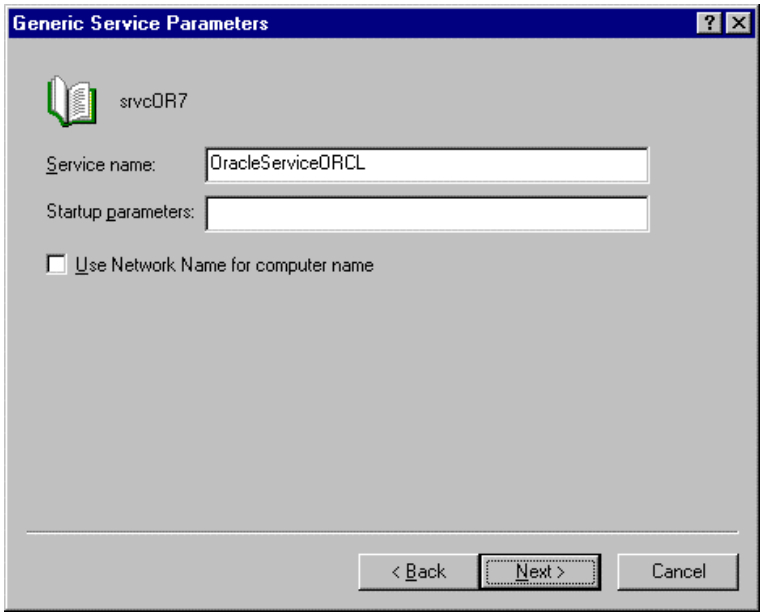

Figure 8-10. Parameters Dialog Box

5. Leave the Registry Replication blank. Click **Finish**.

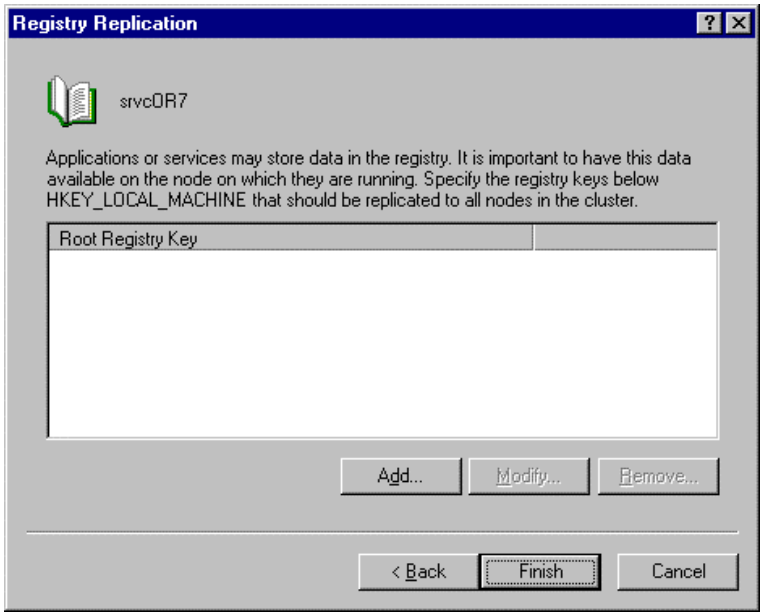

Figure 8-11. Registry Replication Dialog Box

6. Bring the Oracle Database resource group on-line. Your Cluster Administrator should look similar to the following:

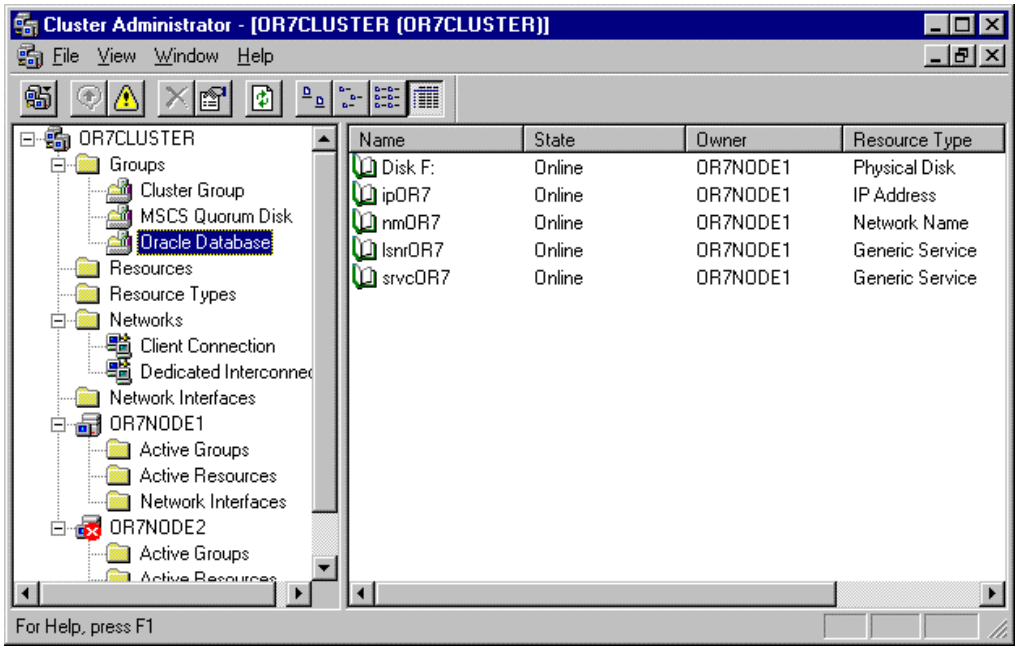

Figure 8-12. Cluster Administrator

Chapter 9
Testing the Cluster

To test the cluster:

1. Make sure both servers are started.
2. Try to move the entire **Oracle Database** resource group from OR7NODE1 to OR7NODE2.

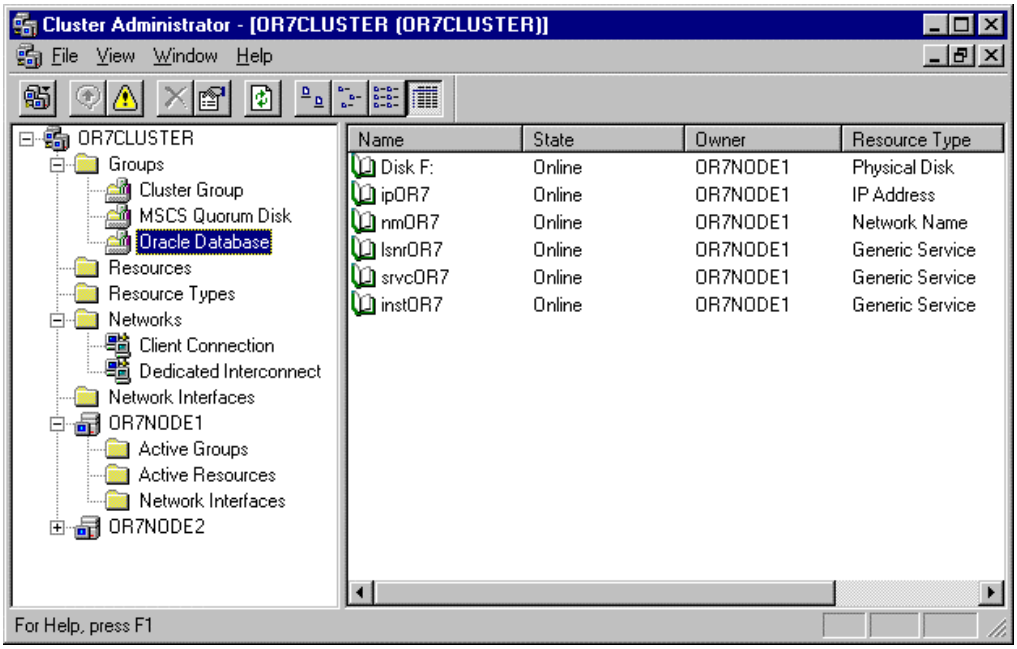

Figure 9-1. Cluster Administrator - Before Moving the Oracle Database Resource Group

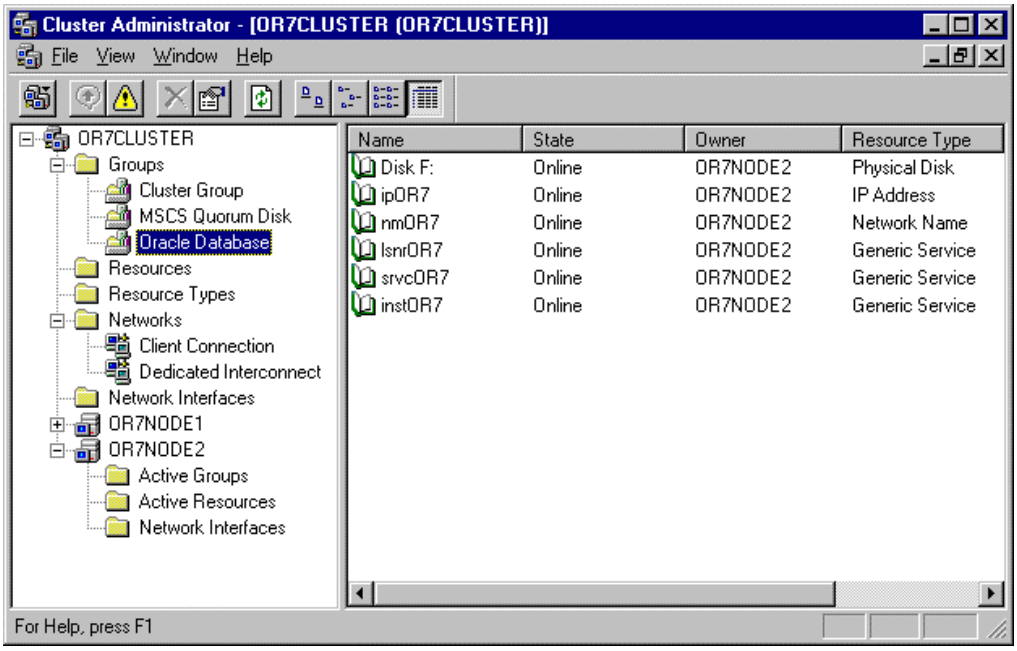

Figure 9-2. Cluster Administrator - After Moving the Oracle Database Resource Group

The owner of all resources should now be changed from OR7NODE1 to OR7NODE2.

Chapter 10
Verifying the Failover

From your client workstation, launch the SQL*Plus application by clicking on the Program icon and login to your database using the following command:

```
C:> PLUS33 SYSTEM/MANAGER@ORACLE
```

Try to connect to the same database using the same connect string, as shown below:

```
SQL*Plus: Release 3.3.2.0.2
Copyright (c) Oracle Corporation 1979, 1994.  All rights
reserved.

Connected to:
Oracle7 Server Release 7.3.2.2.0 - Production Release
With the distributed, replication and parallel query
options
PL/SQL Release 2.3.2.2.0 - Production

SQL> connect system/manager@oracle
Connected.
SQL>
```

Verify that your client can connect to the database using the above connect string. Currently, OR7NODE1 is running the database services.

Move the database group from OR7NODE1 to OR7NODE2 to simulate a failover.

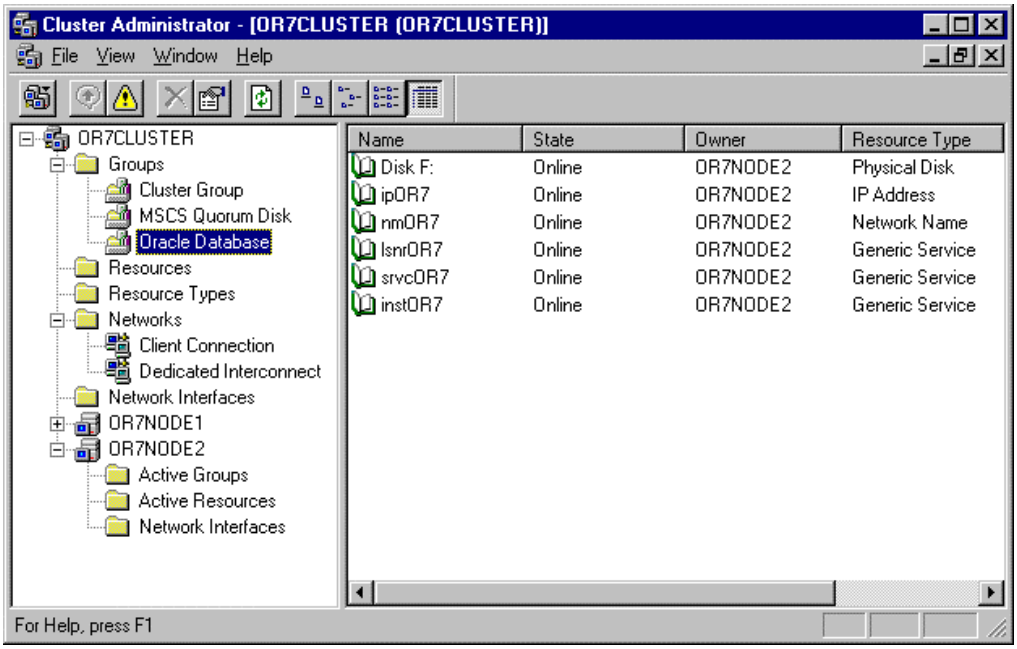

Figure 10-1. Cluster Administrator

After the failover, try to connect to the database using the same connect string as shown below:

```
SQL*Plus: Release 3.3.2.0.2
Copyright (c) Oracle Corporation 1979, 1994.  All rights
reserved.

Connected to:
Oracle7 Server Release 7.3.2.2.0 - Production Release
With the distributed, replication and parallel query
options
PL/SQL Release 2.3.2.2.0 - Production

SQL> connect system/manager@oracle
Connected.
SQL> connect system/manager@oracle
ERROR: ORA-12571: TNS:packet writer failure

Connected.
SQL>
```

Since the client is connected to the cluster, as configured in your SQL*Net, even with a failover of the database services, the client can re-connect to the database using the same database alias.

Chapter 11
Client and Server Considerations

A Re-Connect is still needed

This solution does not provide a way to remove the need to re-connect to the database. A client application that checks for a connection to the database prior to executing any requests from the database will work very well. A simple algorithm is shown below to illustrate this:

```
// FUNCTION TO CHECK IF THERE IS A CONNECTION, RE-CONNECT
IF THERE IS NONE

FUNCTION CHECK_CONNECTION(CONNECT_ATTEMPTS): BOOLEAN
BEGIN
   FOR I = 1 TO CONNECT_ATTEMPTS DO
   BEGIN
      IF (CONNECTED) THEN
         RETURN (TRUE)
      ELSE
         CONNECT TO DATABASE
      END IF
   END
   RETURN (FALSE)
END

// MAIN PROGRAM

BEGIN
   IF CHECK_CONNECTION (CONNECT_ATTEMPTS) = TRUE THEN
       ::
      READ THE DATABASE
       ::
   END IF
END
```

The algorithm shown above is a good way to request services from the database, whether the client wants to read or write to the database, and whether the client is connected to a single server or a cluster.

The use of a web browser that is connected to a web server, which is also using an Oracle database server is an example of an explicit re-connect of a client, and should be avoided.

Data Loss and Data Integrity

It is important to understand that this TechNote is not intended to solve the problems of data integrity and data loss. It is, however, intended to work with existing techniques that resolve data loss and data integrity, and at the same time, minimize the downtime of the system. It can also provide a means of scheduling maintenance to the servers without depriving clients access to the system.

Data loss and data integrity should be resolved using the techniques already presented by Oracle, such as the use of redo logs and rollback segments. You would have to implement these techniques similar to how you would implement them if you were connected to a single server. Since you are keeping all the logs, trace files, and data files in the shared drive, the surviving server uses the same set of files in the event of a failover.

Appendix
Model Hardware Configuration

Server 1

Compaq ProLiant 2500, 1P P5/200, 128-MB memory

Table A-1
Server Configuration Information

Slot Number	Description	IRQ
Embedded NIC	Netelligent 10/100 TX (Client Network)	5
Embedded Disk Ctrl	C875 Ultra Fast Wide SCSI Controller	14
Slot 1	Netelligent 10/100 TX (Interconnect Network)	9
Slot 2	Fibre Channel Host Adapter	11

Table A-2

Disk Configuration Information

Disk Controller	Controller Order	SCSI ID or Logical Drv	Disk Configuration	Raid Level	Drive Letter
Embedded Disk Ctrl	1st	0	2.1-GB	none	C
	2nd	5	CD-ROM	n/a	D
Fibre Host Ctrl	3rd	0	2 x 2.1-GB	1	E
		1	2 x 2.1-GB	1	F

Server 2

Compaq ProLiant 2500, 1P P5/200, 128-MB memory

Table A-3
Server Configuration Information

Slot Number	Description	IRQ
Embedded NIC	Netelligent 10/100 TX (Client Network)	5
Embedded Disk Ctrl	C875 Ultra Fast Wide SCSI Controller	14
Slot 1	Netelligent 10/100 TX (Interconnect Network)	9
Slot 2	Fibre Channel Host Adapter	11

Table A-4
Disk Configuration Information

Disk Controller	Controller Order	SCSI ID or Logical Drv	Disk Configuration	Raid Level	Drive Letter
Embedded Disk Ctrl	1st	0	2.1-GB	none	C
	2nd	5	CD-ROM	n/a	D
Fibre Host Ctrl	3rd	0	2 x 2.1-GB	1	E
		1	2 x 2.1-GB	1	F

Glossary

Availability The ability of a system to appear to the users as if it was operating continuously, even though failures have occurred.

Cluster A cluster is a set of independent computer systems working together as a single system. The cluster software directs client requests to one of the systems in the cluster based on resource accessibility, server load, and application requirements. When a system in the cluster fails the cluster software will recover and disperse the work from the failed system amongst the remaining systems in the cluster. When the overall load exceeds the capabilities of the systems in the cluster, additional systems may be added to the cluster. The failure of a system in the cluster will not affect the other systems and in certain cases the client applications should be unaware of the failure, except for some loss of service for a short period of time.

Given the above definition, it would be hard to distinguish a cluster from a general distributed system; therefore, clusters provide the following additional capabilities:
Every system in the cluster must have full connectivity and communications with all other systems in the cluster.
Systems in the cluster know all the other systems in the cluster; they join in order to form a cluster. The cluster systems also know when systems join or leave the cluster, via a connection manager
Systems in the cluster know what resources are available on all other members of the cluster.
There is a virtual circuit, provided by the connection manager that provides reliable delivery of messages between systems in the cluster

Connect Descriptor A specially formatted description of the destination for a network connection. Connect descriptors are constructed using a set of keywords and values mapped to service names.

Failover The time it takes a resource to failover from one system to another. For example, if a resource is to be polled at a 1-second interval, then we might decide that after two poll intervals, that a resource is dead and that it should be restarted or failed over. Also, let's assume that the local restart count for the resource is 4, then we would require 4 of these timeouts before the resource is failed over to another system. Given a 1-second interval, we would have a period of 2 seconds times the 4 retries - or a total of 8 seconds before the resource failed over to another system.

Listener An executable program that enables an Oracle7 server to accept connections from client machines over SQL*Net.

Member A system which has joined a cluster

Node *See* system.

Oracle Instance Every time a database is started, a system global area (SGA) is allocated and Oracle background processes are started. The system global area is an area of memory used for database information shared by the database users. The combination of the background processes and memory buffers is called an Oracle instance.

Quorum A voting mechanism maintained by the connection manager. A simple quorum voting mechanism is used to indicate when more than half of the available votes are present in the cluster. Note, that a member (or many members) in the cluster can have more than one vote. In addition, a resource available within the cluster can have a vote. This latter ability allows two systems cluster to function in the presence of a communication or system failure, by having the system that continues functioning and reserves the resource to obtain a quorum.

There are also change quorums and operational quorums, which are described in quorum.doc.

Quorum resource	A quorum resource is any resource that has writable persistent storage for maintaining cluster database change log records.
Reliability	The dictionary defines this as meaning dependable or trustworthy Our intention is to make the cluster, as a whole, reliable, in that clients can depend on the services provided by the cluster.
Resource	A physical or logical entity that is managed (see resource manager below). They include, but are not limited to: disks, tapes, printers, CD-ROMs, IP address, etc.
Server	The layer of software running on a cluster node which performs service registration or advertisement. Any software that uses this layer will inherit its properties.
Server application	The actual application software running on a cluster node, regardless of whether it does service registration.
Service	A data set or operation set exported by application servers to their clients. This term is distinguished from a Windows NT Service.
Service Name	A short, convenient name mapped to a network address contained in a TNS connect descriptor. Users need only know the appropriate service name to make a TNS connection.
Shared disk	A cluster organization that requires all systems of the cluster to have access to shared Disks.
Shared nothing	A cluster organization in which disks are ‘owned’ by a single system of the cluster and requests for disk access is directed to the system that ‘owns’ the disk.
Shared resource	A cluster organization in which some resources are accessible to all systems in the cluster.

SID	Oracle Server ID, this is not the same as the SID of Windows NT.
Single system image	A cluster that is 'logically' viewed by clients as being a single system. Services are provided by the cluster that is available to clients, without distinction as to which system in the cluster is actually providing the service. The cluster is also managed as a single system, with the ability to administer all systems in the cluster from a single system within the cluster.
Startup Wait Time	This parameter sets the number of seconds that the listener sleeps before it responds to the first listener control STATUS command. This feature assures that a listener with slow communication lines has adequate time to start up before responding to a STATUS request. The default of 0 means that the listener responds to the first STATUS command without any delay. Any positive numeric value is acceptable.
SQL*Net	The Oracle client/server communication software that offers transparent operation to Oracle tools or databases over any type of network protocol and operating system.
System	A complete computer system capable of operating independently (either a single processor system or an SMP system), for example, a PC.
TNS	Transparent Network Substrate (TNS) is the Oracle networking technology that provides a single application interface to all industry-standard networking protocols.

Index

A

C

D

E

F

H

I

L

M

N

O

P

R

S

T

U

V