

DIGITAL Clusters for Windows NT

Service Pack 3 Addendum

December 1998

This addendum supplies information on installing and configuring DIGITAL Clusters for Windows NT Version 1.1 with Service Pack 3.

Also included are new features, Registry Synchronization, Network Failover, and information on configuring script-based Oracle Failover Objects.

Revision/Update Information: This addendum supplements the *DIGITAL Clusters for Windows NT Administrator's Guide*, Second Edition, AA-QVUTB-TE. and *DIGITAL Clusters for Windows NT Administrator's Guide Addendum*, AV-R171B-TE.

Operating System and Version: Microsoft Windows NT Version 1.1 with Service Pack 3

Software Version: DIGITAL Clusters for Windows NT Version 1.1 with Service Pack 3

**Compaq Computer Corporation
Houston, Texas**

December 1998

The information in the addendum 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 THE MATERIAL.

This addendum contains information protected by copyright. No part of this guide may be photocopied or reproduced in any form without prior written consent from Compaq Computer Corporation.

© Compaq Computer Corporation 1998. All rights reserved. Printed in U.S.A.

Compaq and the Compaq logo registered United States Patent and Trademark Office.

DIGITAL, Prioris, and the DIGITAL logo are trademarks of Compaq Computer Corporation.

Intel is a registered trademark of Intel Corporation.

Lotus and Lotus Notes are registered trademarks of Lotus Development Corporation.

NetBIOS is a trademark of Micro Computer Systems, Inc.

Netscape and Netscape Navigator are registered trademarks of Netscape Communications Corporation.

Microsoft, MS-DOS, WIN32, Windows, Windows 95, and Windows NT are registered trademarks of Microsoft Corporation.

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.

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

Table of Contents

1 DCNT VERSION 1.1 SP3 UPGRADE & INSTALLATION 1–1

Upgrading to DIGITAL Clusters for Windows NT Version 1.1 Service Pack 3

..... 1–1

Requirements for Upgrading to DCNT Version 1.1 Service Pack 3..... 1–1

DIGITAL Clusters for Windows NT Version 1.1 Service Pack 3 Upgrade Overview

..... 1–1

Upgrading a Cluster Server to DCNT Version 1.1 Service Pack 3 1–3

Installing DIGITAL Clusters for Windows NT Version 1.1 Service Pack 3 1–8

Requirements for Installing DCNT Version 1.1 Service Pack 3 1–8

Before You Start 1–9

DIGITAL Clusters for Windows NT Version 1.1 Service Pack 3 Installation Overview

..... 1–10

Rebooting Your Server Systems 1–11

Installing DIGITAL Clusters for Windows NT Version 1.1 Service Pack 3 on the First

Server 1–12

Installing DIGITAL Clusters for Windows NT Version 1.1 SP3 Software on the

Second Server..... 1–23

Completing and Verifying the Installation..... 1–24

Problems? 1–24

Configuring the Cluster Software..... 1–25

Installing the Remote Cluster Administrator.....	1-26
Installation Requirements	1-26
Installing the Remote Cluster Administrator Tool	1-26
Upgrading the Remote Cluster Administrator to Version 1.1 Service Pack 31-34	
Installation Requirements	1-34
Upgrading to Remote Cluster Administrator Version 1.1 Service Pack 3 ...	1-34
2 USING REGISTRY SYNCHRONIZATION.....	2-1
Before You Start	2-2
Using Registry Synchronization	2-3
3 USING NETWORK FAILOVER.....	3-1
Before You Start	3-2
Using Network Failover	3-3
4 USING ORACLE FAILOVER OBJECTS.....	4-1
Before You Start	4-2
Creating Command Scripts	4-4
Sample command script for starting an Oracle failover object	4-4
Sample command script for stopping an Oracle failover object.....	4-6
Hints and Tips for Creating Command Scripts	4-8
Using Script-based Oracle Failover Object Screens	4-10
Using Parameter-based Oracle Failover Object Screens	4-14

About This Addendum

This addendum supplies updated information on installing and configuring DIGITAL Clusters for Windows NT Version 1.1 with Service Pack 3.

It also supplies installation and user information related to two new features, Registry Synchronization and Network Failover, and updated information on configuring script-based Oracle failover objects.

Audience

This addendum is for system administrators who will manage the DIGITAL Clusters software. The addendum assumes that you are familiar with the tools and methodologies needed to maintain your hardware, operating system, and network.

Organization

This addendum consists of four chapters, as follows:

- | | |
|-----------|---|
| Chapter 1 | Presents the changes to the Installation procedure for DIGITAL Clusters for Windows NT Version 1.1 with Service Pack 3. |
| Chapter 2 | Presents a new feature, Registry Synchronization. |
| Chapter 3 | Presents a new feature, Network Failover. |
| Chapter 4 | Presents the configuration changes to Oracle failover objects using a new script-based format. |

Conventions

The following conventions are used in this addendum:

Convention	Meaning
DIGITAL Clusters 1. <i>n</i>	Refers to DIGITAL Clusters for Windows NT Version 1. <i>n</i> , where <i>n</i> has a value of 0 or 1.
Chapter <i>n</i>	Chapter references throughout this addendum refer to chapters in the <i>DIGITAL Clusters for Windows NT Administrator's Guide</i> .
Bold	Bold type indicates the actual commands, words, or characters that you type in a dialog box or at the command prompt.
<i>Italic</i>	Italic type indicates a placeholder for information on parameters that you must provide. For example, if the procedure asks you to type <i>filename</i> , you must type the actual name of a file. Italic type also indicates new terms and the titles of other manuals in the DIGITAL Clusters for Windows NT package. Italic type is used for emphasis within procedures as well.
ALL UPPERCASE	All uppercase letters indicates an acronym.
Monospace	Monospaced type represents examples of screen text or entries that you might type at the command line or in initialization files.
▶	A right triangle indicates a procedure with sequential steps.

Related Information

Several other key sources of information included in the DIGITAL Clusters for Windows NT package will help you plan for and use the cluster software:

- Online release notes
- *DIGITAL Clusters for Windows NT Administrator's Guide*
- *DIGITAL Clusters for Windows NT Configuration and Installation Guide*
- *DIGITAL Clusters for Windows NT Administrator's Guide Addendum*
- *DIGITAL Clusters for Windows NT Service Pack 2 Addendum*
- Online help

1

DCNT Version 1.1 SP3 Upgrade & Installation

This chapter describes how to install to DIGITAL Clusters for Windows NT® Version 1.1 Service Pack 3. It also describes the procedures for upgrading your DIGITAL™ cluster to DIGITAL Clusters for Windows NT® Version 1.1 Service Pack 3, installing the Remote Cluster Administrator and upgrading the Remote Cluster Administrator to Version 1.1 Service Pack 3.

Upgrading to DIGITAL Clusters for Windows NT Version 1.1 Service Pack 3

To upgrade DIGITAL Clusters for Windows NT Version 1.1 Version 1.1 Service Pack 3, you must install this upgrade on each cluster server node.

Requirements for Upgrading to DCNT Version 1.1 Service Pack 3

Before this upgrade, make sure each of your cluster server nodes meets the following system requirements:

- The cluster server operating system must be Windows NT 4.0 Service Pack 3 or newer.
- DIGITAL Clusters for Windows NT Version 1.1 must be installed.
- Screen resolution must be VGA or better.

DIGITAL Clusters for Windows NT Version 1.1 Service Pack 3 Upgrade Overview

Use the following procedure to perform a rolling upgrade:

1. Insure that all users are disconnected from database resources.

Installation

2. With both servers in the cluster running, manually failover all cluster resources to one of the servers.
3. On the other server (the server without any cluster resources), exit the Cluster Administrator.
4. Upgrade to DIGITAL Clusters for Windows NT Version 1.1 Service Pack 3 and reboot the server. For a detailed procedure, see the next section.
5. Manually failover all cluster resources to the server that was just upgraded to DIGITAL Clusters for Windows NT Version 1.1 Service Pack 3.
6. Exit the Cluster Administrator
7. Now upgrade on the second server to DIGITAL Clusters for Windows NT Version 1.1 Service Pack 3. For a detailed procedure, see the next section.
8. Reboot the second server.

Upgrading a Cluster Server to DCNT Version 1.1 Service Pack 3

Use the following procedure to upgrade a cluster server to DIGITAL Clusters for Windows NT Version 1 Service Pack 3.

► **To upgrade a cluster server to DIGITAL Clusters for Windows NT Version 1.1 Service Pack 3:**

1. Run Setup.exe, which is located in the following directory:

..\Clu1-1nt.40\CluSP3\Setup.exe

2. The Welcome dialog box appears. Click Next.



Installation

3. Enter your name and the name of your company.

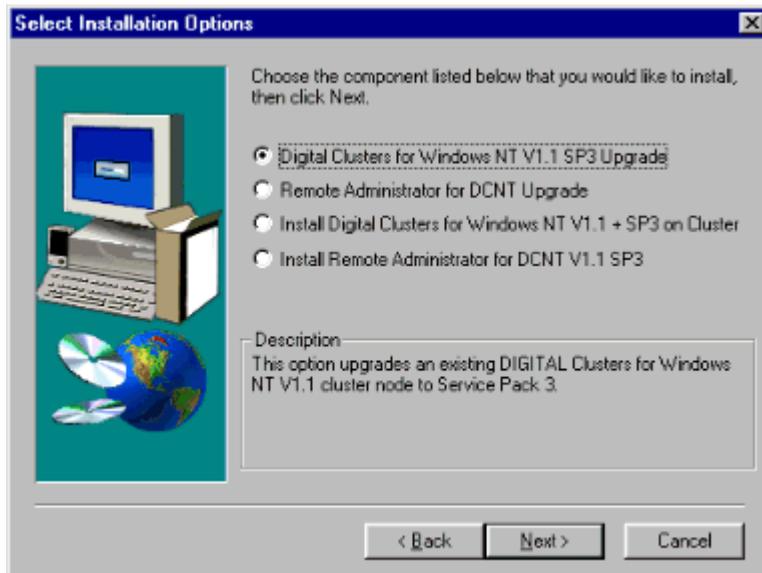
Confirm the information, then click Next to proceed to the next screen.

If needed, you can return to a previous screen by clicking Back and change the information you entered.

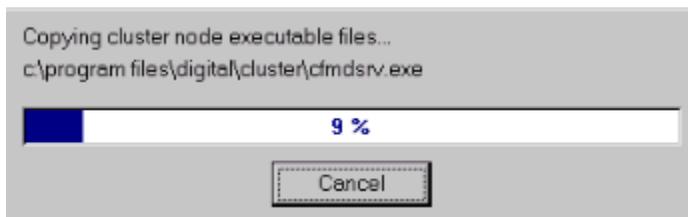


The image shows a Windows-style dialog box titled "User Information". The title bar is blue with a close button (X) on the right. The main area has a light gray background. On the left, there is a vertical panel with a teal background containing icons of a computer monitor, keyboard, mouse, and a globe. To the right of this panel, the text reads: "Type your name below. You must also type the name of the company you work for." Below this text are two white text input fields. The first field is labeled "Name:" and the second is labeled "Company:". At the bottom of the dialog box, there are three buttons: "< Back", "Next >", and "Cancel".

4. In the Select Installation Options dialog box, select the DIGITAL Clusters for Windows NT V1.1 SP3 Upgrade radio button, and then click Next.



The installation program now indicates that it is copying server files to your system.



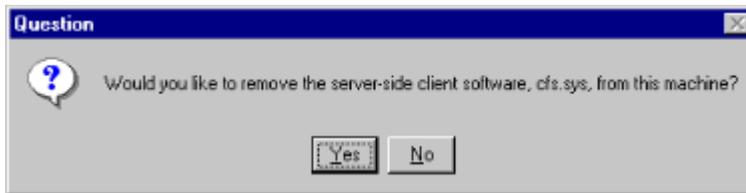
Installation

- When the server files are copied, you now have the option of removing the server-side client software from your server.

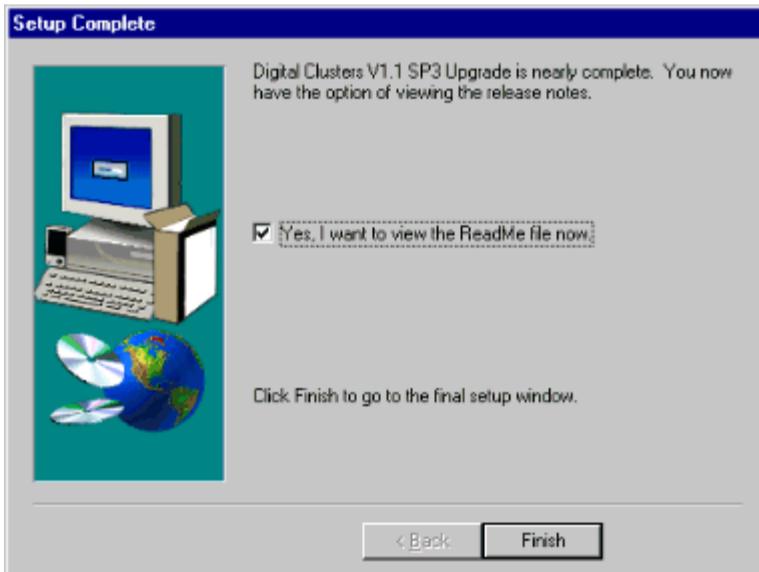
Note

If you have previously installed DIGITAL Clusters for Windows NT Version 1.1 Service Pack 2, and remove the server-side client software at that time, this dialog box will not appear.

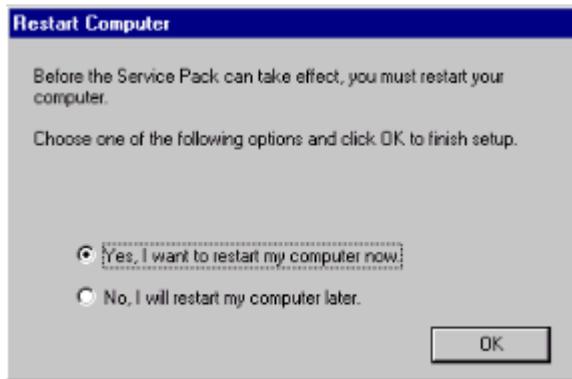
If TCP/IP is installed on your cluster servers you do not need the client software. Click Yes to remove the server-side client software.



- The Setup Complete dialog box is displayed. Compaq strongly suggests that you read the ReadMe file at this time. Click Finish to complete this installation.



7. For the Version 1.1 Service Pack 3 upgrade to take effect, you must restart your computer. Select Yes, and click the OK button.



Installing DIGITAL Clusters for Windows NT Version 1.1 Service Pack 3

This section describes how to install DIGITAL Clusters for Windows NT Version 1.1 with Service Pack 3 on your two cluster servers.

For an overview of the full cluster setup procedure, see Chapter 2 of the DIGITAL Clusters for Windows NT Configuration and Installation Guide.

Note:

This installation option is only available on the CD-ROM release of DCNT V1.1 Service Pack 3.

Requirements for Installing DCNT Version 1.1 Service Pack 3

This table lists the software and hardware requirements for installing and running DIGITAL Clusters server software:

Clusters Version 1.1 SP3	Microsoft Windows NT Server 4.0 with NT Service Pack 3 or later
Hardware	Certified DIGITAL Cluster hardware*
Hard disk space (approximate)	5 MB
Distribution media	CD-ROM

*The supported Hardware list can be found at www.windows.digital.com/clusters

Other Requirements

Your cluster servers must also meet the following system and licensing requirements:

- The two cluster servers must be running in an Ethernet or FDDI network, using one of the following PC LAN protocols:
 - NetBEUI
 - TCP/IP
 - IPX/SPX

The two servers must have at least one protocol in common.

- Both servers must belong to the same Windows NT Server domain. You can use primary and backup domain controllers (PDCs and BDCs) as cluster servers.

- Your Windows NT Server licensing for a server should cover the total number of clients that may connect to that server if a failover occurs from the other server.

Before You Start

Before installing the cluster server software, complete the following checklist:

- Ensure that your hardware is properly installed (see DCNT Configuration and Installation Guide, Chapter 3).**

Caution

If the disks on your shared bus contain data you want to preserve, install the server software on both servers *before* creating and connecting the shared bus. Otherwise, there is danger of disk corruption if the system power for both servers is turned on while the shared bus is connected and the cluster software is not yet installed.

- Create a cluster administrator account.**

When you install the server software, the Setup program prompts you for the user name, domain name, and password of an account to be used as the cluster administrator account. This account must be in the Domain Administrators group and have the advanced user rights to log on as a service.

DIGITAL recommends that you create such an account specifically to serve as the cluster administrator account and suggests using `ClusterAdmin` as the user name.

▶ **To create the cluster administrator account:**

1. Start Windows NT and log in to an account with Domain Administrator privileges.
2. Start the User Manager for Domains and create an account with the user name `ClusterAdmin` in the Domain Administrators group.
3. Set up the account rights using the Policy menu and give the account the advanced user right to log on as a service.

Installation

- **Identify your shared buses.**
During the installation, Setup scans the system hardware for bus adapters and asks you to specify which buses are shared. You should know the port and bus numbers of your shared buses before you start the installation process.
- ▶ **To identify your shared buses on a Window NT 4.0 server system:**
 1. Open the Windows NT Control Panel, SCSI Adapters applet.
 2. Expand the list under each device.
 3. Select a device, then choose Properties to see the information for that device.
 4. Look for the port and bus that match the configuration of your shared bus, with the correct adapter type and combination of disk types and SCSI IDs.
 5. Note the port and bus numbers of all your shared buses.

DIGITAL Clusters for Windows NT Version 1.1 Service Pack 3 Installation Overview

This section lists the basic installation steps. The following sections describe each step in detail. The related procedures for each of these steps are provided in the following sections.

- ▶ **To install the server software:**
 1. Turn off the power for the shared storage and turn on the power for both cluster servers.
 2. Install the server software on the first server. After the installation, leave the server turned on but do not reboot it.
 3. Install the server software on the second server). After the installation, leave the server turned on but do not reboot it.
 4. Turn on power for the shared storage, then reboot the servers *one at a time*
 5. With the power on for the entire cluster, verify that the software is installed properly and that you can see all your shared disks
 6. Use the Cluster Administrator to assign shared disks to failover groups and Windows NT Explorer to create file shares.

Rebooting Your Server Systems

After you install the cluster server software on each server, the system displays instructions to reboot your servers, *one at a time*. It is critical that you understand the proper way to reboot your servers.

Caution

Unsynchronized access to the shared bus can corrupt data on your shared disks.

► **To safely reboot your servers:**

1. Reboot one server.
2. Wait until the logon prompt appears on that server.
3. Reboot the second server.

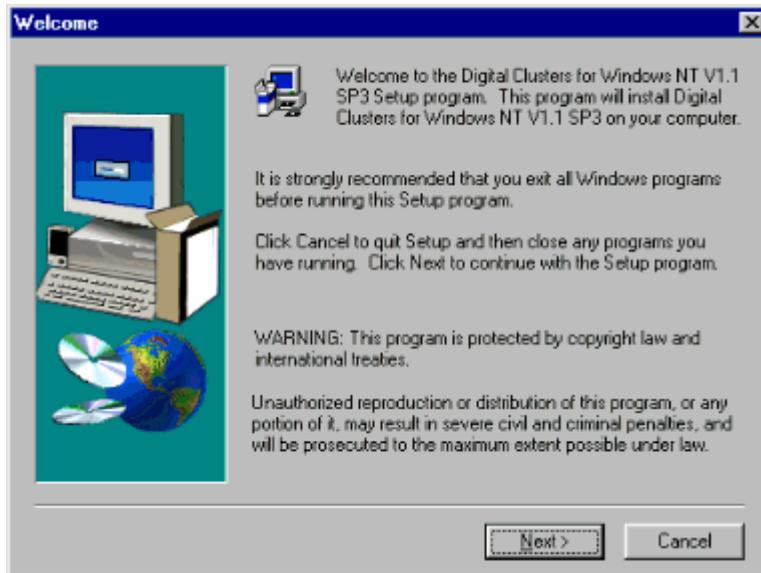
Installing DIGITAL Clusters for Windows NT Version 1.1 Service Pack 3 on the First Server

Caution

Before you start the installation, make sure the power for the shared storage is off. Unsynchronized access to the shared disks can corrupt the disks.

► **To install the DIGITAL Clusters server software on the first server:**

1. Start Windows NT and log in to the cluster administrator account. Close all open applications.
2. Run Setup.exe, which is located in the following directory:
..\Clu1-1nt.40\CluSP3\Setup.exe
3. The Welcome dialog box appears. Click Next.



4. Enter your name and the name of your company.

Confirm the information, then click Next to proceed to the next screen.

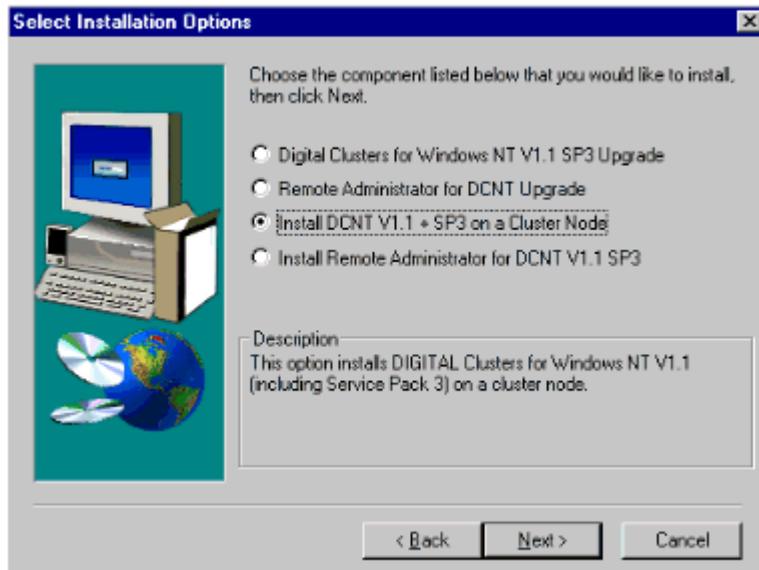
If needed, you can return to a previous screen by clicking Back and change the information you entered.



The image shows a Windows-style dialog box titled "User Information". The dialog box has a blue title bar with a close button (X) in the top right corner. On the left side, there is a vertical panel with a teal background containing icons of a computer monitor, keyboard, mouse, and a globe. To the right of this panel, the text reads: "Type your name below. You must also type the name of the company you work for." Below this text are two text input fields. The first field is labeled "Name:" and the second is labeled "Company:". At the bottom of the dialog box, there are three buttons: "< Back", "Next >", and "Cancel".

Installation

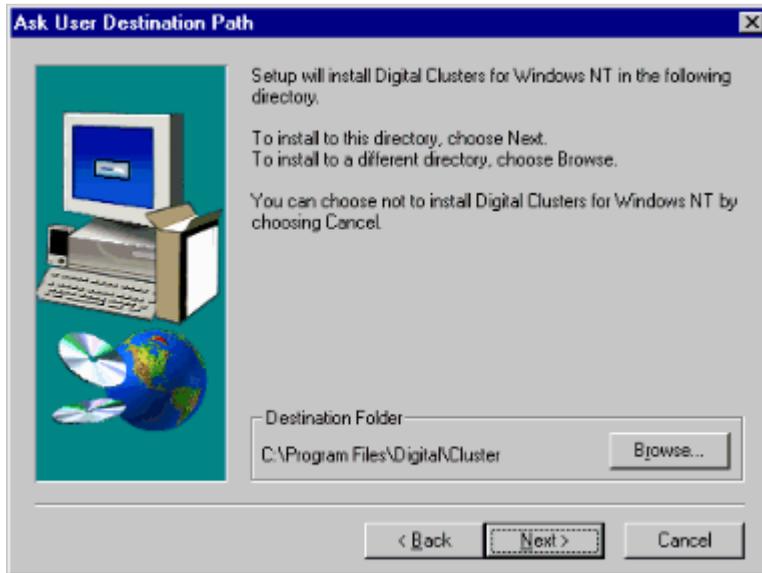
5. In the Select Installation Options dialog box, select the Install DCNT V1.1 + SP3 on a Cluster Node radio button, and then click Next.



6. Confirm the directory path for the installed files.

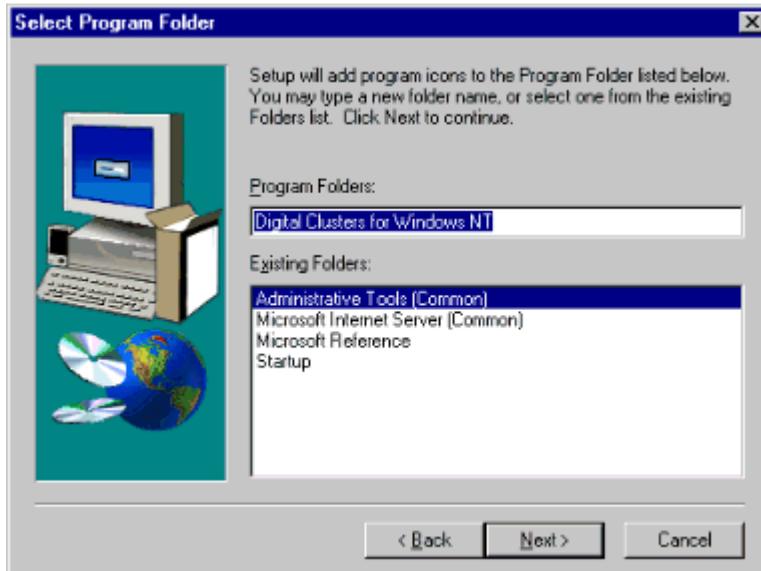
To change the directory path, click Browse and select a new directory.

If the specified destination directory does not exist, the Setup program asks you if you want to create it.

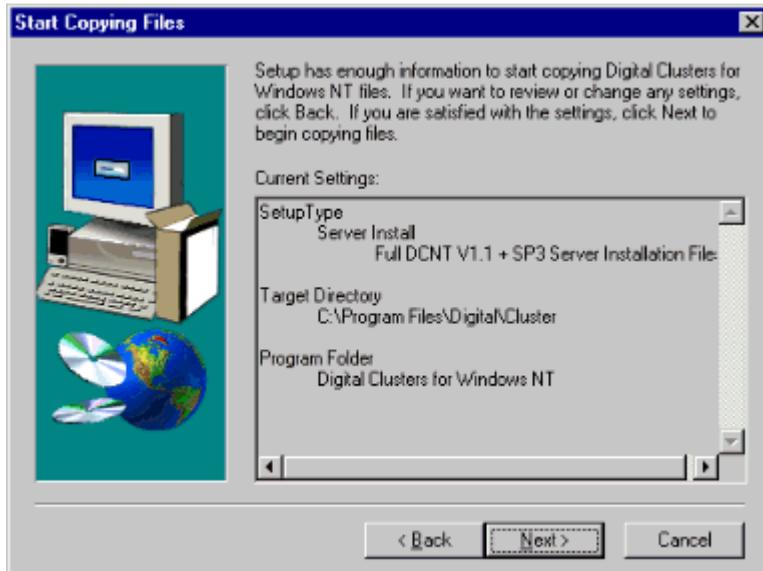


Installation

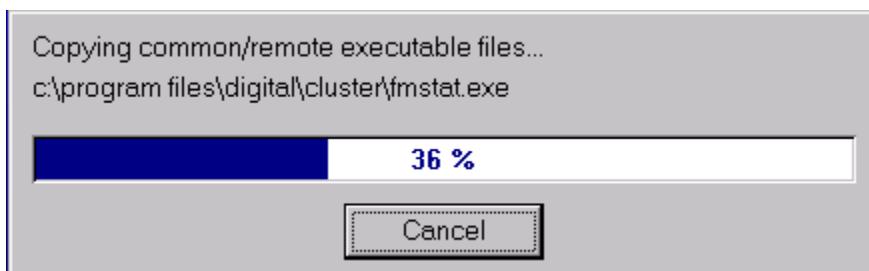
7. Select the program folder for the program icons. You can select an existing folder from the Existing Folders list or create a new folder by typing the new folder name in the Program Folders box.



- Verify the information you have supplied to the Setup program. If you want to change any of the supplied information, click Back repeatedly until you return to the window containing the information you want to change.

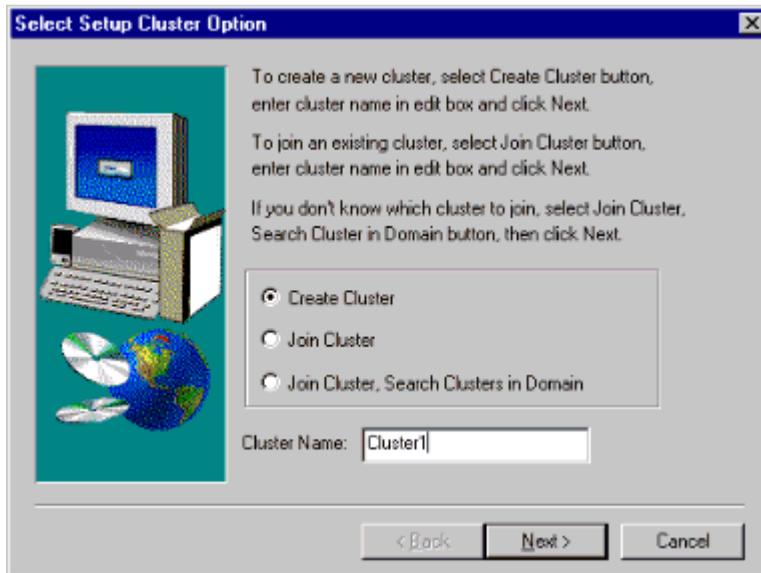


If the information displayed is correct, click Next. The Setup program copies files to the destination directory. While Setup is copying files, the program displays a status meter to mark its progress.



Installation

9. The Setup program prompts you to select whether you want to create a new cluster or join an existing cluster. Since you are installing the software on the first server of your cluster, select Create Cluster. Enter the name you want to give your cluster. Cluster names can contain up to 15 characters.



Multiple Clusters in a LAN

If you have more than one cluster in your LAN, make sure the first eight characters of cluster names are unique. Otherwise, it is possible for one cluster to see resources owned by another cluster.

Examples

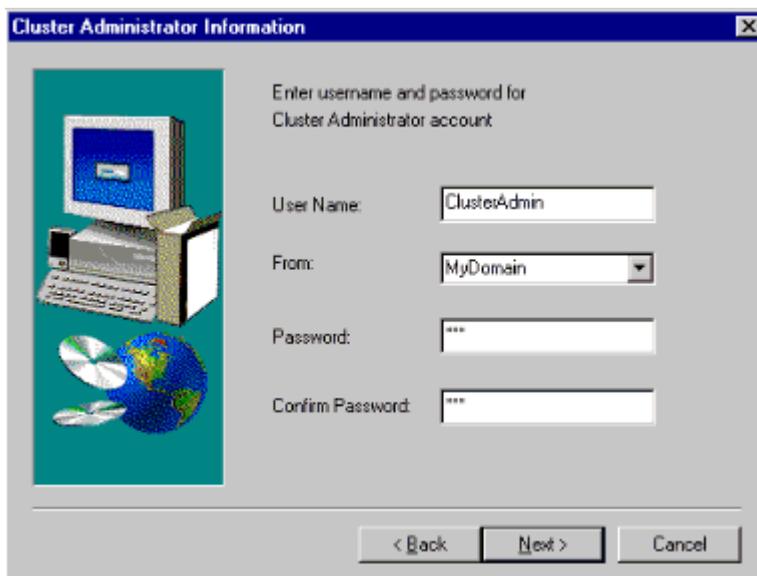
Valid: CLUSTER1 and CLUSTER2
Invalid: CLUSTER1A and CLUSTER1B

10. Enter the user name, domain name, and password of an account in the Domain Administrators group that has the advanced user right to log on as a service.

Note

DIGITAL recommends that you create such an account specifically to serve as the cluster administrator account and suggests using ClusterAdmin as the user name. See Before You Start

The system does not display your password on the screen. You must enter the password again to confirm it and guard against typing errors.



Cluster Administrator Information

Enter username and password for Cluster Administrator account

User Name: ClusterAdmin

From: MyDomain

Password: ***

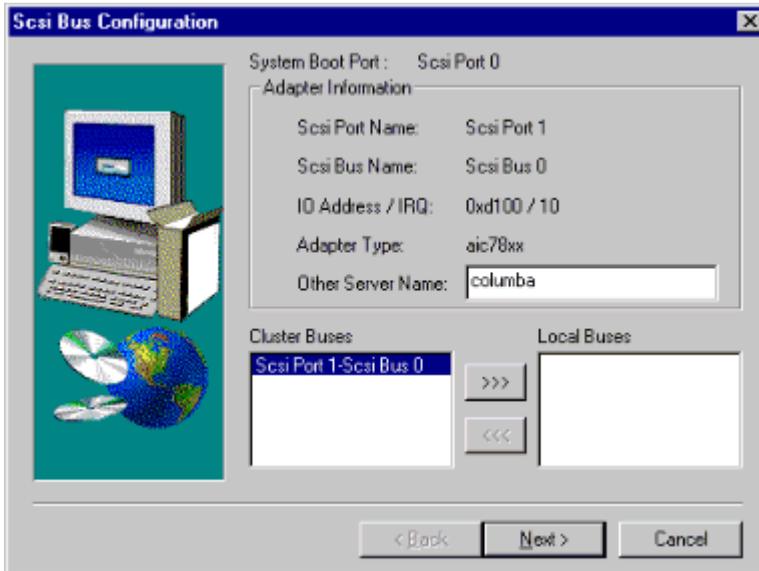
Confirm Password: ***

< Back Next > Cancel

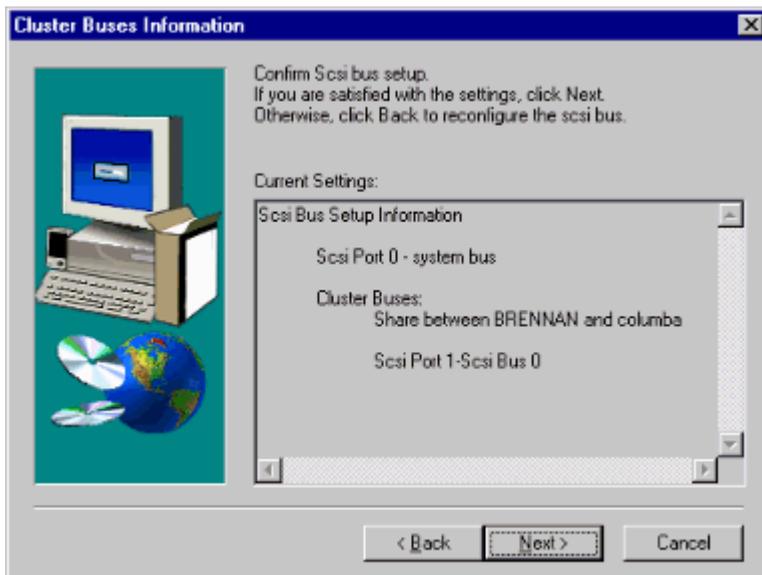
Installation

11. The Setup program scans the system hardware, looking for SCSI bus adapter cards. For each adapter port, select whether the bus is a local bus or a bus shared with another server. If it is a shared bus, enter the name of the other server.

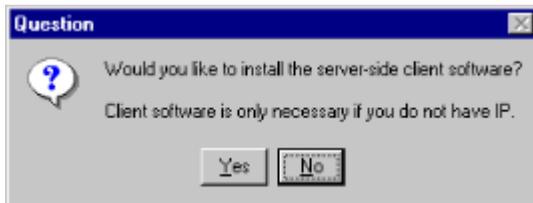
For details on identifying your shared bus, see Before You Start



12. Review and verify the bus adapter information you have supplied. To change any information, choose Back and reenter the bus information.

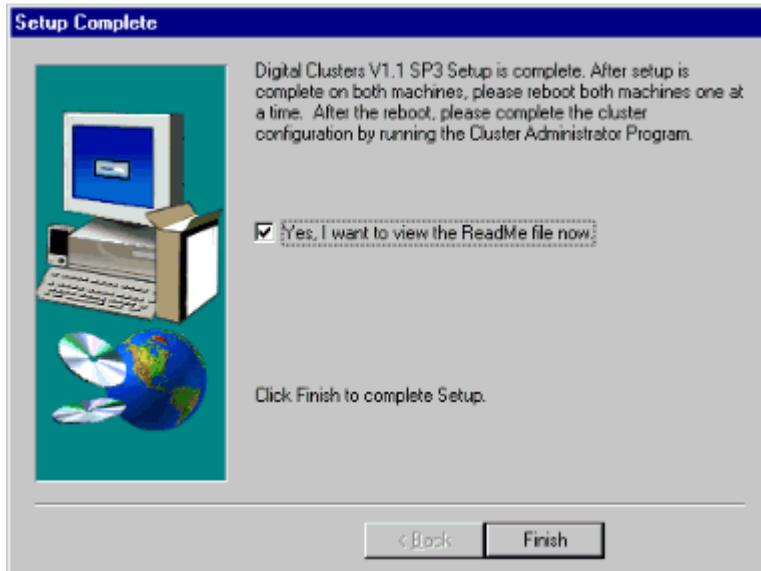


At this point you have the option of installing the server-side client software on your server. If TCP/IP is installed on your cluster servers you do not need the client software. The default selection, No, will not install the server-side client software.



Installation

13. The Setup program creates the Cluster program folder, then displays the Setup Complete dialog box. Choose Finish to complete Setup. DIGITAL strongly recommends that you read the ReadMe file at this time.



14. Leave the system power turned on, but *do not* reboot the system at this time.

Installing DIGITAL Clusters for Windows NT Version 1.1 SP3 Software on the Second Server

Caution

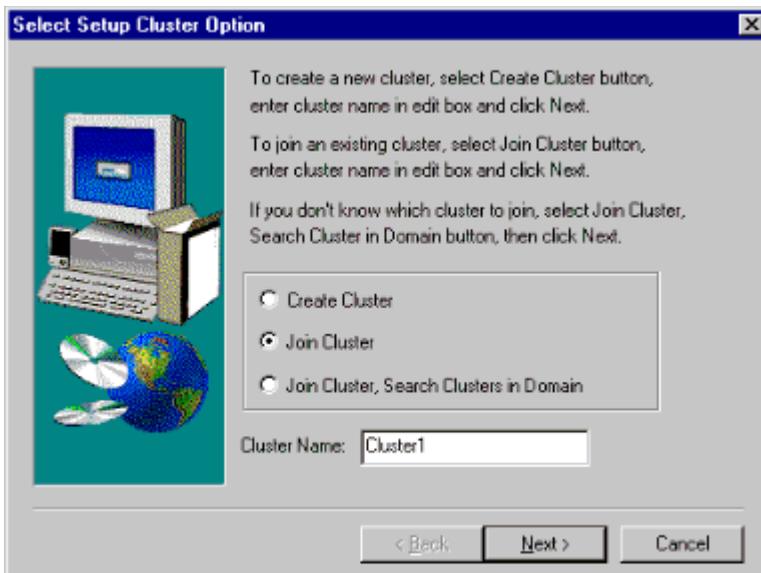
Before you start the installation, make sure the power for the shared storage is off. Unsynchronized access to the shared disks can corrupt the disks.

► To install the server software on the second server:

1. Perform steps 1 to 8 of the first server software installation
2. Select Join Cluster and enter the name of the cluster.

If you forgot the name of the cluster, select Join Cluster, Search Clusters in Domain to have Setup search for existing clusters.

Searching for existing clusters in a domain can take several minutes. DIGITAL strongly recommends that you do not select Search Cluster in Domain unless you absolutely have to.



Installation

3. Complete steps 10 to 14 in the section *Installing the Software on the First Server*.

Completing and Verifying the Installation

After installing the cluster server software on the second cluster server, use the following procedure to start the servers and verify the software installation.

▶ **To complete the installation of the server software on both servers:**

1. Turn on the power for the shared storage.
2. Reboot the servers, *one at a time*.

After you reboot both servers, verify that the software was installed properly.

▶ **To verify that the software was installed properly:**

1. Wait about 1 minute for the registries on both systems to be updated with the cluster configuration. Then run the Cluster Administrator on both servers. Select *Class View* from the drop-down list and click on *Disk*. You should see a list of your shared disks. You cannot view the shared cluster resources on either server until the cluster software is installed on both systems. At that point, you can view the shared resources from both systems.
2. If you do not see your shared disks, use the *Services* applet in the *Control Panel* to verify that the following services started on both servers:
 - Cfmnd Server
 - Cluster Failover Manager
 - Cluster Name Service

Problems?

Some problems require reinstalling the cluster software:

- **The cluster services are not running.** This indicates the cluster software is not installed properly. The most likely reason is that the software was not installed from the proper user account—that is, an account in the *Domain Administrators* group with the advanced user right to log on as a service. See *Before You Start* for information on creating the cluster administrator. To correct this problem, you must reinstall the cluster software.
- **The cluster services are running but you cannot see your shared disks.** You may have specified the bus adapter for your shared SCSI bus incorrectly. If so, you must

reinstall the software. Otherwise, see the Troubleshooting chapter of the *Administrator's Guide* for more information on possible causes of disk problems.

► **To reinstall the cluster software:**

1. Turn off the power for the shared storage.
2. Uninstall the server software on the affected server (DIGITAL Clusters for Windows NT Configuration Guide, Appendix A). Then reboot the server.
3. Log in to the server by using a properly configured cluster administrator account and reinstall the server software as described in the section Installing the Software on the Second Server. Make sure you follow the instructions exactly.

Configuring the Cluster Software

After installing the server software, you must run the Cluster Administrator to place your shared disks in failover groups and Windows NT Explorer to create file shares. See your *Administrator's Guide* for details.

Installing the Remote Cluster Administrator

DIGITAL Clusters for Windows NT Version 1.1 Service Pack 3 provides the capability of performing cluster administration tasks from a system which is not a member of the cluster. This section describes the system requirements for installing this tool and provides a detailed installation procedure.

Note:

If Remote Cluster Administrator Version 1.1 Service Pack 2 is already installed, proceed to the next section and perform an upgrade.

Note:

Do not install Remote Cluster Administrator on a cluster node. Cluster nodes running DIGITAL Clusters for Windows NT Version 1.1 SP2 or SP3 already have the ability to remotely administer clusters.

Installation Requirements

Before installing the Remote Cluster Administrator, make sure your system meets the following requirements.

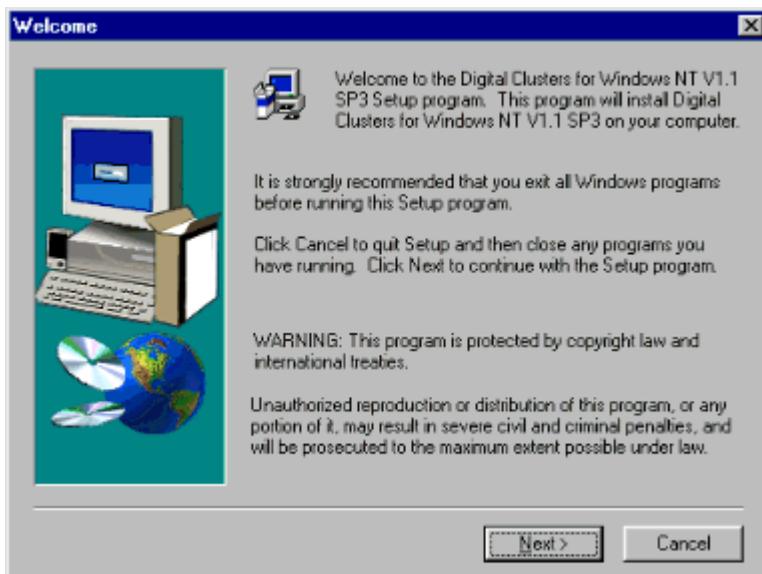
- The operating system must be any version of Windows NT.
- The user installing Remote Cluster Administrator must have administrative privileges on the system.
- Screen resolution must be VGA or better.

Use the following procedure to install the DIGITAL Clusters for Windows NT Remote Cluster Administrator.

Installing the Remote Cluster Administrator Tool

► **To install the Remote Cluster Administrator:**

1. Run Setup.exe, which is located in the following directory:
 `..\Clu1-1nt.40\CluSP3\Setup.exe`
2. The Welcome dialog box appears. Click Next.



Installation

3. Enter your name and the name of your company.

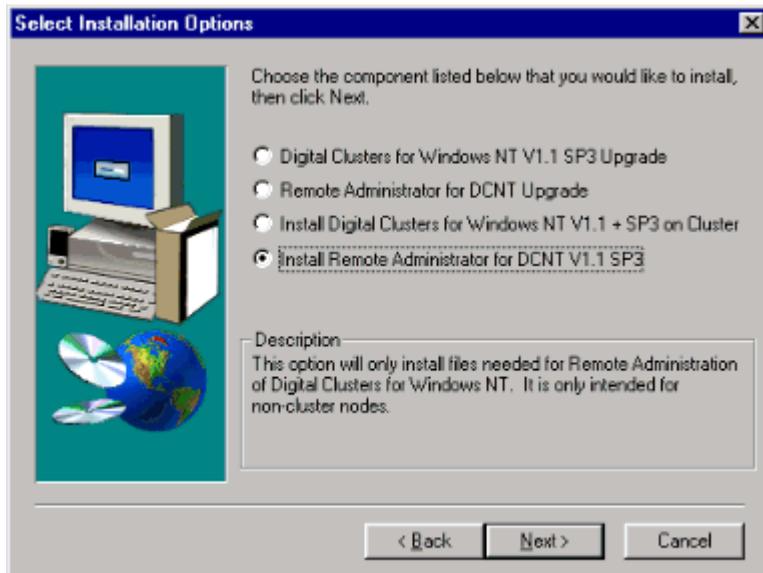
Confirm the information, then click Next to proceed to the next screen.

If needed, you can return to a previous screen by clicking Back and change the information you entered.



The image shows a Windows-style dialog box titled "User Information". The dialog has a blue title bar with a close button (X) in the top right corner. On the left side, there is a vertical panel with a teal background containing icons of a computer monitor, keyboard, mouse, and a globe. To the right of this panel, the text reads: "Type your name below. You must also type the name of the company you work for." Below this text are two text input fields. The first field is labeled "Name:" and the second is labeled "Company:". At the bottom of the dialog, there are three buttons: "< Back", "Next >", and "Cancel".

4. In the Select Installation Options dialog box, select the Install Remote Administrator for DCNT 1.1 SP3 radio button, and then click Next.

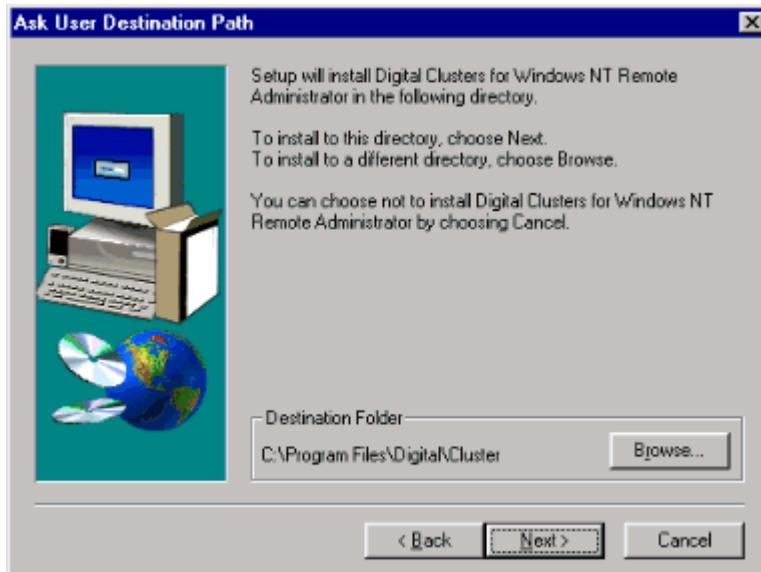


Installation

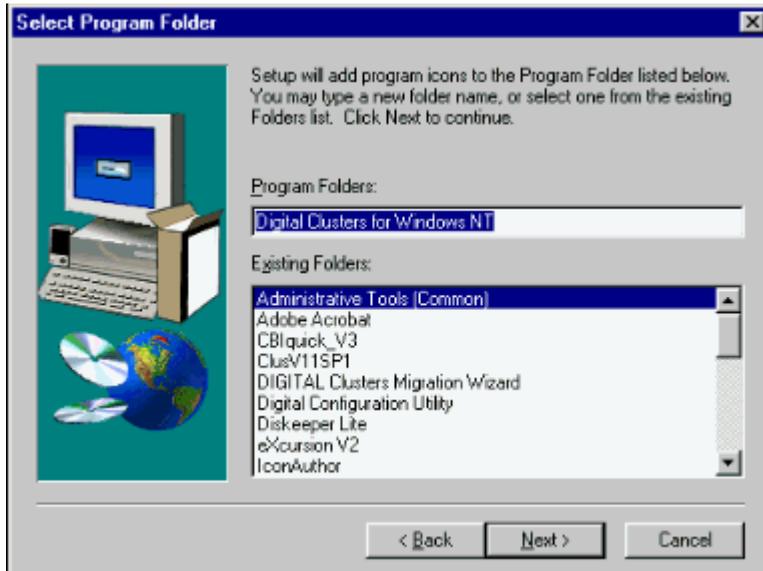
5. Confirm the directory path for the installed files.

To change the directory path, click Browse and select a new directory.

If the specified destination directory does not exist, the Setup program asks you if you want to create it.

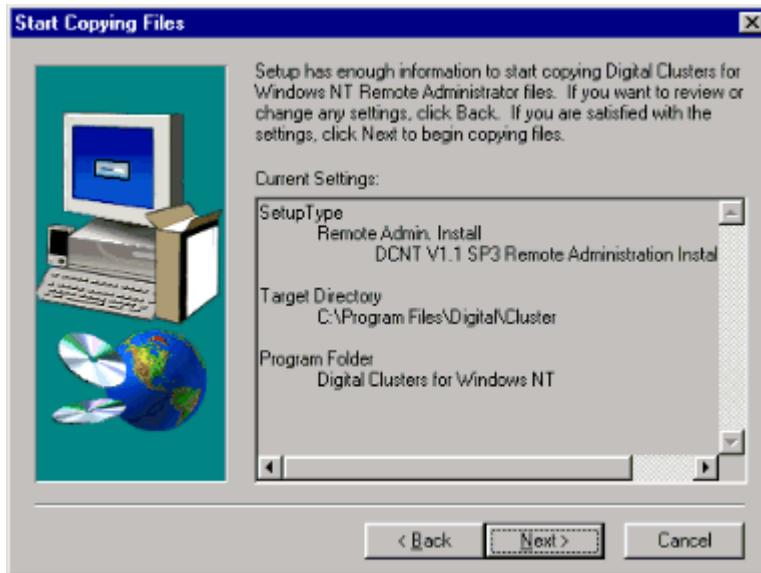


6. Select the program folder for the program icons. You can select an existing folder from the Existing Folders list or create a new folder by typing the new folder name in the Program Folders box.

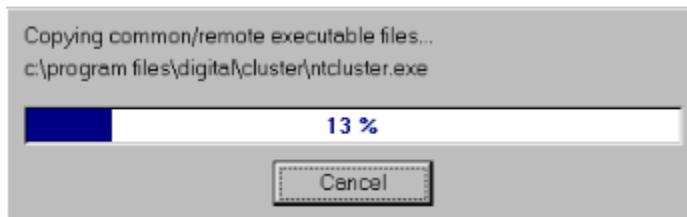


7. Verify the information you have supplied to the Setup program. If you want to change any of the supplied information, click Back repeatedly until you return to the window containing the information you want to change.

Installation



If the information displayed is correct, click Next. The Setup program copies files to the destination directory. While Setup is copying files, the program displays a status meter to mark its progress.



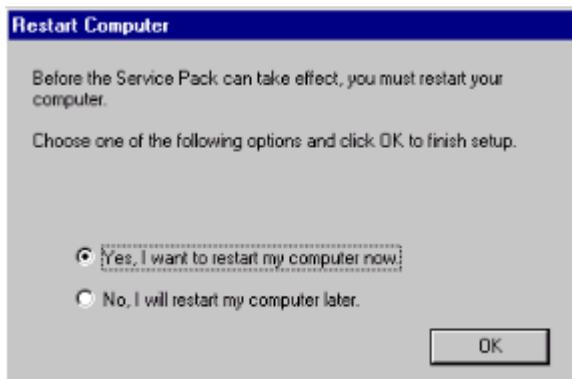
- The Setup program creates the Cluster program folder, then displays the Setup Complete dialog box. Click Finish to complete Setup. Compaq strongly recommends that you read the ReadMe file at this time.

This completes the Remote Cluster Administration installation.



Note

The Setup program may prompt you to reboot if shared or locked files were installed.



Upgrading the Remote Cluster Administrator to Version 1.1 Service Pack 3

This section describes the system requirements for upgrading to Remote Cluster Administrator Version 1.1 Service Pack 3 and provides a detailed upgrade procedure.

Installation Requirements

Before upgrading the Remote Cluster Administrator to Version 1.1 Service Pack 3, make sure your system meets the following requirements.

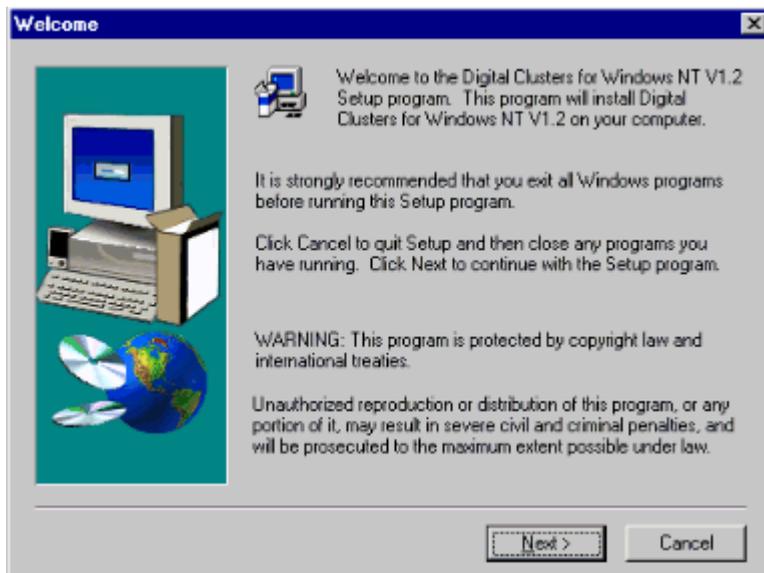
- The user installing Remote Cluster Administrator must have administrative privileges on the system.
- The system must already have Remote Cluster Administrator Version 1.1 Service Pack 2 installed.

Use the following procedure to upgrade to DIGITAL Clusters for Windows NT Remote Cluster Administrator Version 1.1 Service Pack 3.

Upgrading to Remote Cluster Administrator Version 1.1 Service Pack 3

► **To upgrade to Remote Cluster Administrator Version 1.1 Service Pack 3:**

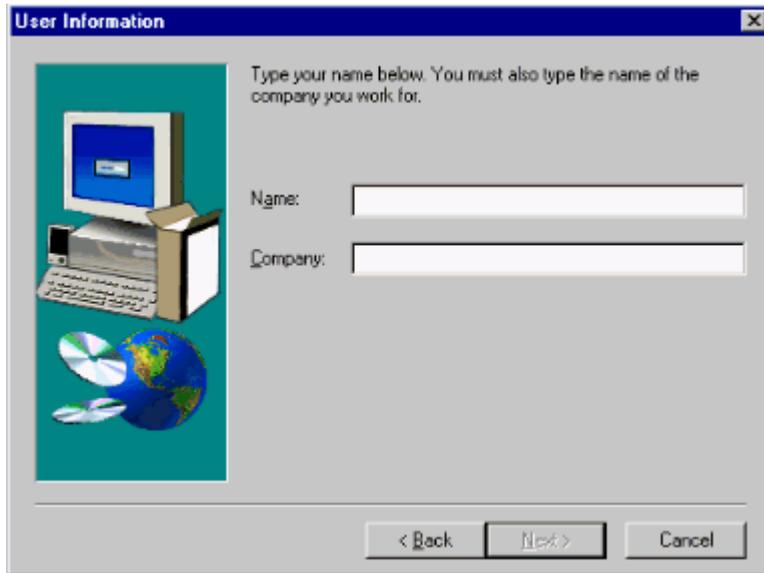
1. Run Setup.exe, which is located in the following directory:
..\Clu1-1nt.40\CluSP3\Setup.exe
2. The Welcome dialog box appears. Click Next.



3. Enter your name and the name of your company.

Confirm the information, then click Next to proceed to the next screen.

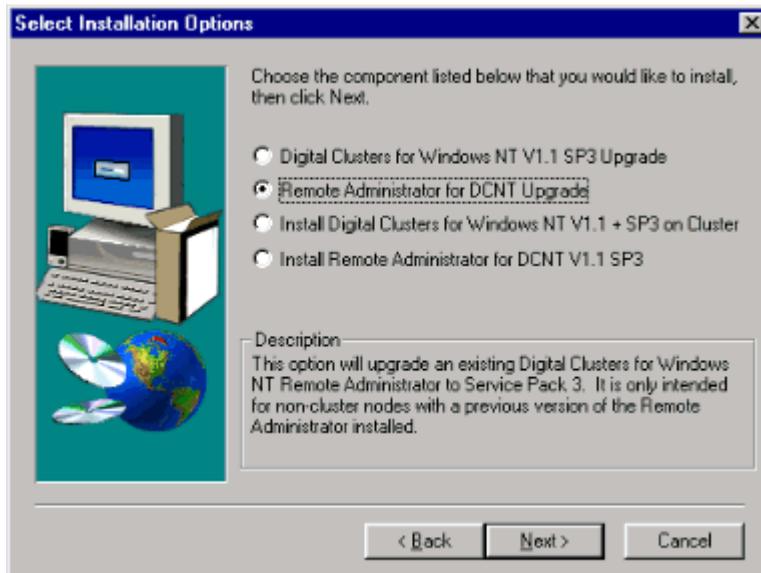
If needed, you can return to a previous screen by clicking Back and change the information you entered.



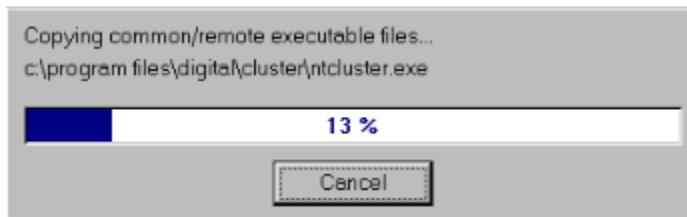
The image shows a Windows-style dialog box titled "User Information". The dialog box has a blue title bar with a close button (X) in the top right corner. The main area is light gray. On the left side, there is a vertical rectangular area with a teal background containing icons of a computer monitor, keyboard, mouse, and a globe. To the right of this area, the text reads: "Type your name below. You must also type the name of the company you work for." Below this text are two white text input fields. The first field is labeled "Name:" and the second is labeled "Company:". At the bottom of the dialog box, there are three buttons: "< Back", "Next >", and "Cancel".

Installation

4. In the Select Installation Options dialog box, select the Remote Administrator for DCNT Upgrade radio button, and then click Next.



The Setup program now copies files to the destination directory. While Setup is copying files, the program displays a status meter to mark its progress.



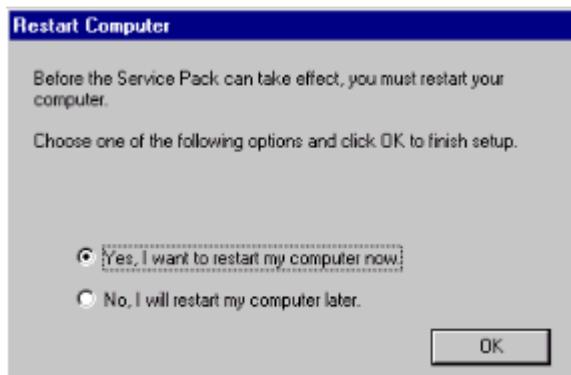
- The Setup program then displays the Copying Complete dialog box. Click Finish to complete Setup. Compaq strongly recommends that you read the ReadMe file at this time.

This completes the Remote Cluster Administration upgrade procedure.



Note

The Setup program may prompt you to reboot if shared or locked files were installed.



Using Registry Synchronization

The Cluster Administrator screens and menus include a new type of failover object, registry synchronization failover object. This new failover object allows application- or service-specific registry information to move with that application or service upon failover. This information can be any state information or data that is saved in the registry. Like most other failover objects, a registry synchronization failover object can be created, modified, deleted and added to, or removed from, a failover group.

A registry synchronization failover object defines a registry key under the `HKEY_LOCAL_MACHINE` key whose hive is copied to the other node upon failover. A *hive* is a key, its values, all of its sub-keys, and their values (see the “Before You Start” section for hive restrictions). Any changes made to the object’s defined hive while it is online are propagated to the other node at the time of failover.

Although hive information is only shared with the other node at failover time, the hive information is stored in a “snapshot” file on the cluster log disk immediately upon changes to that hive. As a result, the log disk has become increasingly important to the proper operation of the cluster (see “Before You Start” for more information on the cluster log disk).

Before You Start

To use registry synchronization, there are no special hardware requirements and the operating system requirements are the same as those for DIGITAL Clusters for Windows NT (DCNT). DCNT Service Pack 3 must be installed on both cluster nodes.

Follow the instructions for installing DIGITAL Clusters 1.1 Service Pack 3 in the “Installation” section of this Administrator’s Guide Addendum.

When using registry synchronization, consider the following guidelines:

- In defining a registry synchronization object’s registry hive key, there are two restrictions:
 1. It cannot be a first generation descendant (child) of HKEY_LOCAL_MACHINE.
 2. Its hive cannot in any way overlap the hive of another registry synchronization object in the cluster.
- Since almost any key can be defined for registry synchronization, in an attempt to make registry synchronization objects as flexible as possible, great care must be taken in defining the hive whose information will be duplicated on the other node upon failover. For instance, defining the “SOFTWARE\Microsoft” key for a registry synchronization object would essentially make the two machines identical from the standpoint of Windows NT and render your cluster unusable. As a rule of thumb, do not define a hive that includes machine-specific information.
- Pay strict attention to the placement of objects in a group. A registry synchronization object should always be placed in the group before an application or service that uses the registry synchronization object’s hive information.
- Remember that registry synchronization is not dynamic. That is, changes to an object’s hive do not immediately propagate to the other cluster node. Changes are propagated only at the time of failover.
- Registry synchronization does not keep the registry keys defined by a registry synchronization object from becoming out-of-sync. The out-of-sync condition can occur if changes are made to the registry hive on the node that does not have that object on line. The registry hives remain in sync on failover as long as changes are only made by the node that owns that resource.
- Previously, the Cluster Log Disk was only used by DCNT to store cluster changes if one of the cluster nodes was down. Now, with the addition of the registry synchronization failover object, having the Log Disk on line is mandatory for the proper operation of a registry synchronization object regardless of whether the other

node is up or down. If the Cluster Log Disk is off line no registry synchronization can occur.

- You can set a flag to make the group that owns a registry synchronization object go off line if that object fails because of Log Disk failure or any other failure. If this flag is not set, you will not have any immediate indication that synchronization did not occur. Set this flag by checking the **Take the group OFFLINE....** check box in the Create or Modify Registry Synchronization Failover Object dialog boxes.

Using Registry Synchronization

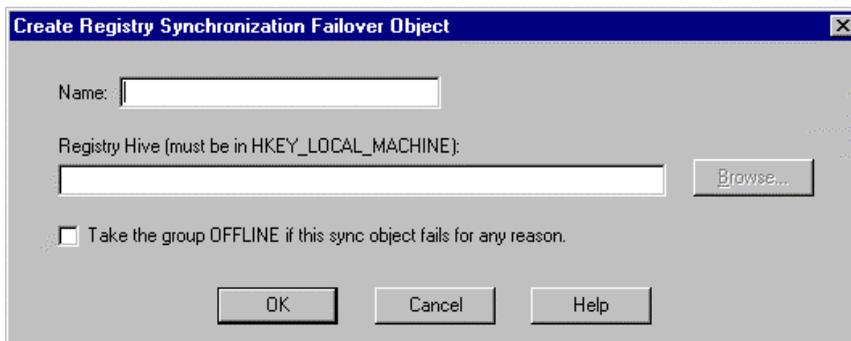
Use the Cluster Administrator program to create, modify, or delete a registry synchronization failover object.

Note

Right click support is available in Cluster Administrator.

► To create a registry synchronization failover object:

1. From the Create menu, select Registry Synchronization Failover Object or from Class View right-click on Sync Failover Objects and select Create Sync from the drop-down menu.
2. The Create Registry Synchronization Failover Object dialog box is displayed.



3. Enter the name of the failover object.
4. Enter the registry hive.

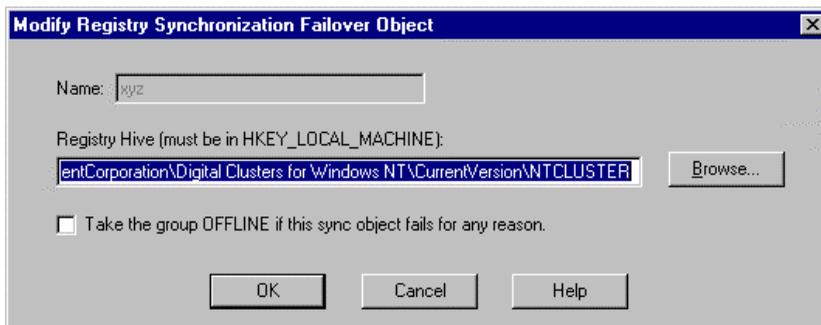
Using Registry Synchronization

5. Note the **Take the group OFFLINE if this sync object fails for any reason** check box. You may wish to check this box. If you don't select this feature you may lose important data. For example, if the log disk is off line (the registry uses the log disk), the hive information would not be saved. There would be no indication that the information is not saved and therefore it would not be failed over. If something does not complete successfully, the system takes the group off line, thereby indicating that a problem exists.

6. Click OK.

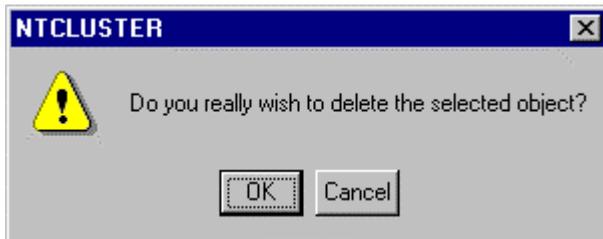
► **To modify a registry synchronization failover object:**

1. In Class View (or Cluster View if the object is in a group), click the registry synchronization failover object you want to modify.
2. From the Modify menu, select Registry Sync Failover Object or choose Modify Sync from the right-click drop-down menu.
3. The Modify Registry Synchronization Failover Object dialog box is displayed. Note that the selected registry synchronization failover object is displayed in the Name box. This entry cannot be changed during this procedure.
4. Make registry hive changes as needed, observing the restrictions in the “Before You Start” section.
5. Click OK.



► **To delete a registry synchronization failover object:**

1. In Class View (or Cluster View if the object is in a group), click the object you want to delete.
2. From the Delete menu, select Registry Synchronization Failover Object or choose Delete Sync from the right-click drop-down menu.
3. A confirmation dialog box is displayed. Click OK to delete object; click Cancel to return without deleting the object.



Using Registry Synchronization

Using Network Failover

A new feature of DCNT 1.1 Service Pack 3, Network Failover, solves the problem of a cluster being oblivious to network failure. This feature allows the cluster to detect and respond to network failure by failing over a group containing an IP object that is bound to an affected network interface card (NIC).

Prior to this, if a network interface card (NIC) on an active cluster node were to fail or lose connectivity with the network, the active cluster node would be isolated and unable to perform network transactions. Although the other cluster node might still be active on the network and could assume cluster operations, the cluster would be unaware that network connectivity had been lost and thus do nothing. Now, network failover can be set to cause the group associated with this NIC to failover and the process can continue.

This feature detects whenever network connectivity is lost, whether through failure of the NIC card itself or by network failure. In the event of connectivity loss, whether direct or as a side effect of a failed NIC, DCNT performs the same action; it fails all groups containing an IP object bound to or dependent upon the isolated NIC over to the other cluster node.

The system detects lack of connectivity by systematically sending a signal to all other systems on the network (pinging) and waiting for a response. If return replies are sent, no action is taken. If all “ping” attempts fail for a fixed (user-settable) time period, it is assumed that connectivity has been lost.

When connectivity is lost, a message is sent through the private network to the partner node, advising it that connectivity is lost. The partner node then attempts, through a similar ping mechanism, to determine if it can see the outside network. If it can, it invokes a failover, taking over responsibility for the group dependent upon the isolated NIC on the cluster partner.

In addition, if the failback feature is selected for this group, the first node continues pinging other members of the network until it receives a response. Once a sufficient interval of consistent replies has occurred (this interval is user-settable), the first node requests failback of the group from the partner node.

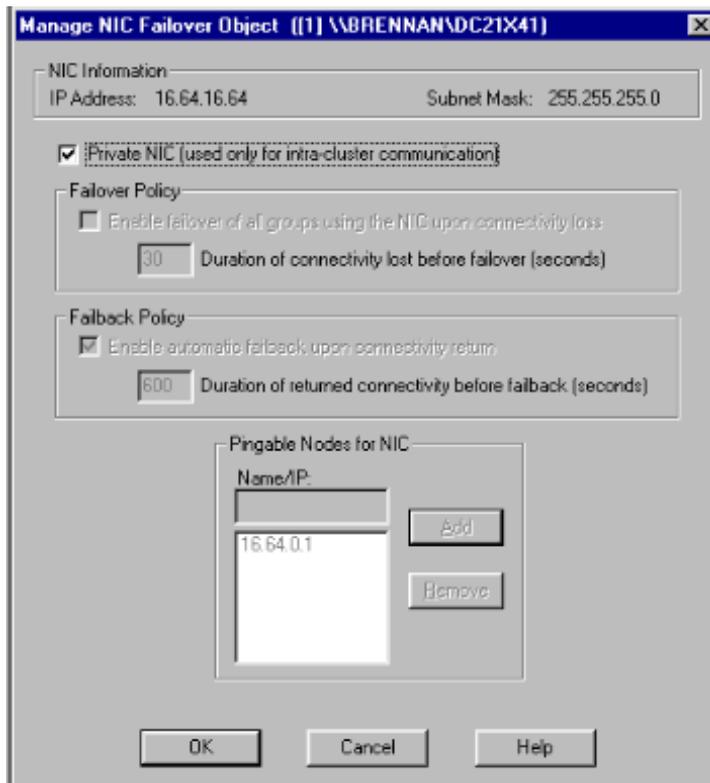
Before You Start

Before setting up network failover objects, take the following preliminary steps:.

1. Understand how your network is configured.
2. Determine which nodes you want to ping. Choose these systems carefully. Use the following guidelines:
 - The system should be nearby on the network, so that wider network interruptions do not cause spurious failovers.
 - Ensure that the system is on line.
 - Preferably, the system should be a lesser connected system so that valuable network bandwidth into the system is not used.
 - The system does not need to be an actual system; any network device that responds to an ICMP ping packet can be used.
 - Do not use the other cluster node as a pingable device. Redundancy of the private network between cluster nodes would make this useless. To determine if the cluster is still on the network, the pingable system(s) must be on the enterprise network.
3. Under System view, check each NIC on both systems. Ensure that the NIC for the private network is designated as private. See the Manage NIC Failover Object screen instructions in the next section to learn how to do this.

Using Network Failover

- To set up a private network NIC for each cluster node:
1. In System view, select the NIC you want to modify.
 2. From the Modify menu, select NIC Failover Object. (or on selected NIC double-click or click the right mouse button).
 3. The Manage NIC Failover Object screen is displayed.

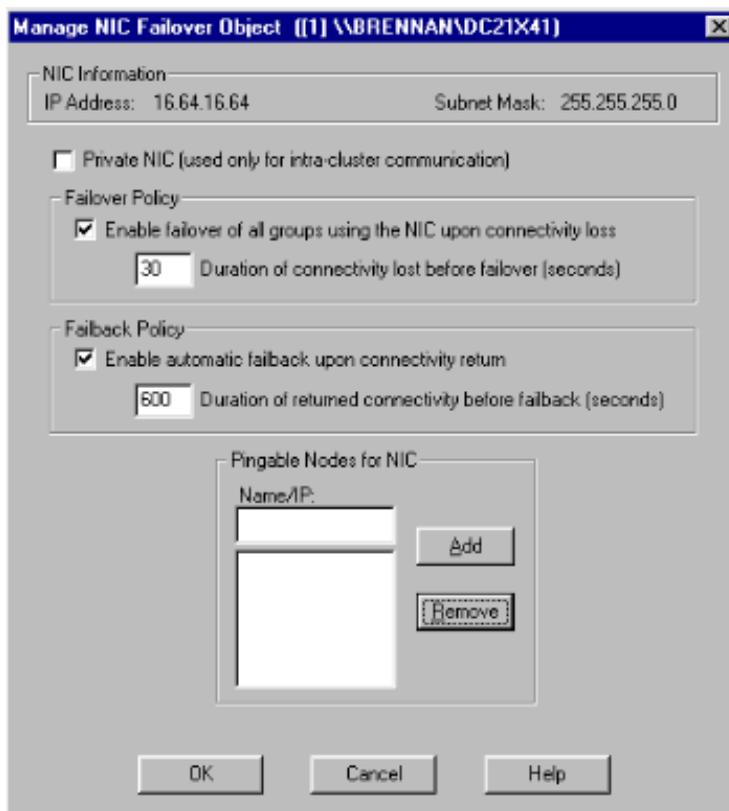


4. Check **Private NIC (used only for intra-cluster communication)** check box to designate that this is a private network NIC.
5. At the Manage NIC Failover Object screen, click OK to save settings; Cancel to exit without saving.

Using Network Failover

► **To manage a NIC failover object:**

1. In System view, select the NIC you want to modify.
2. From the Modify menu, select NIC Failover Object. (or on selected NIC double-click or click the right mouse button).
3. The Manage NIC Failover Object screen is displayed.



4. Ensure that **Private NIC (used only for intra-cluster communication)** check box is **not selected**.
5. Check **Enable failover of all groups using the NIC upon connectivity loss** check box (Note: If the private NIC is indicated, there is no failover or failback policy.) If this is an enterprise NIC, indicate that you want to set a policy and set a time by checking this box and entering the time you want. Default is 30 seconds. If you set this interval of time too short, the smallest of network interruptions may cause the

failover routine to start. During the failover period, data can be lost. Thirty seconds is the recommended duration so that a small fault cannot set off the failover routine.

6. Check **Enable automatic failback upon connectivity return** check box if you want to create a failback policy. Note that the **Duration of returned connectivity before failback (seconds)** contains a recommended default setting of 600. This 600 seconds gives the network adequate time to stabilize after coming back up before the failback begins. If you want to failback the system manually at a more convenient time, do not set a failback policy.
7. Use the **Pingable Nodes for NIC** text boxes to set the nodes you want the cluster to ping to determine if the cluster nodes are able to reach the network. You may type either the name of the node or its IP address.

Note

If you choose to type the name of the node, be aware that at failover time the network must check the DNS server to convert the name to its IP address. This may cause problems if the network is unstable. For this reason, it is preferable to use the IP address of the node.

8. Click the Add button to add the node to the pingable node list.
9. The **Check Connectivity?** dialog box is shown. This routine checks for network connectivity by pinging the network node. Either a misspelling of a name, a mistyped IP address or an offline node is indicated by the dialog box message, 'Node xxx did not respond to a ping from this NIC. Would you like to add this node to the list anyway?'
10. Click Yes to add to the list; No to cancel.

Note

The maximum allowable number of nodes on the pingable list is 16. There is no significant network overhead related to the number of nodes as only one ping to the network is sent every second. The cluster node rotates through the pingable node list. In the case of a 10-node list each node would only be pinged every 10 seconds.

11. To remove a node from the list, click on the node you want to delete. Click Remove.
12. At the Manage NIC Failover Object screen, click OK to save settings; Cancel to exit without saving.

Using Network Failover

Note

Additional NIC information is now available on the Create IP Failover Object screen.

It is advisable to check any pre-existing IP failover objects to ensure they are as you want them.

► **To create an IP failover object:**

1. To create an object, choose IP Failover Object from the Create menu, or from the Class View right-click on IP Failover Objects and select Create IP from the drop-down menu.
2. The Create IP Failover Object dialog box is displayed.
3. Create: In the IP Address text box, enter the IP address to use as a cluster IP address. The address also serves as the name of this IP object. Select this address when adding this IP object to a failover group.
4. In the Subnet Mask field:

For DIGITAL Clusters for Windows NT Version 1.1 Service Pack 3 the Subnet Mask field is automatically set to the Subnet Mask value of the selected network adapter card.

5. Optional step: In the NetBIOS Name field, enter the name (15 characters or less).
6. Select the network adapters on each cluster server which can enable the IP address.

A list box for each server displays the list of available network adapters. For DIGITAL Clusters for Windows NT Version 1.1 Service Pack 2 all network adapters, both public and private, are listed for both machines. In DCNT 1.1 Service Pack 3, the left list box displays all the public NIC cards on the machine you are currently connected to. The right list box displays all the NIC cards with the same Subnet Mask as the card currently selected in the left list box. If a network adapter is dedicated to a private network between the two cluster servers, you must specify a different network adapter to support failover.

Note

If you use DCNT V1.1 SP3 to modify an IP Object that was created with DCNT V1.1 SP2 or earlier, where (Any) was selected for a machine's adapter, the selected network adapter card for that machine will no longer be (Any). It will instead be the network adapter card that is actually being used.

The left list box displays all the public NIC cards on the machine you are currently connected to. The right list box displays all the NIC cards with the same Subnet Mask as the card currently selected in the left list box.

The *Selected NIC Information* field displays the information for the selected network adapter card.

7. Choose OK to save the modifications to the settings, or choose Cancel to cancel the settings.

Using Network Failover

Using Oracle Failover Objects

Previous DIGITAL Cluster for Windows NT (DCNT) versions use a set of Oracle database related parameters for the creation of an Oracle failover object. Those parameters control the starting and stopping of the Oracle database. This is achieved by the “real-time” creation of `.bat` and `.sql` files to start/stop the Oracle object. All existing Oracle failover objects on your system use this parameter-based format.

A new feature of DCNT 1.1 Service Pack 3, the implementation of a script-based system asks the user directly for the commands to start/stop the Oracle object. This direct control over the starting or stopping of the Oracle object allows the usage of the Multiple Oracle Homes feature introduced in Oracle 8.0.4.

This new feature also provides an automatic conversion mechanism to convert the old parameter-based format to the new script-based format.

Note

For databases created in Oracle8 and above, the script-based format must be used. The parameter-based format may not work because the executable name and listener name can change for each Oracle revision.

Before You Start

Before installing, upgrading or making changes to Oracle objects, please note the following:

Warning:

Do not modify, create, or convert Oracle failover objects while a rolling upgrade to DCNT V1.1 Service Pack 3 is in progress. Both cluster nodes must be upgraded to DCNTV1.1 Service Pack 3 before performing any of the aforementioned actions.

Modifying, creating or converting Oracle failover objects in the middle of a rolling upgrade can disrupt service availability of the Oracle databases.

- During the rolling upgrade, all existing Oracle failover objects will still function normally.
- For Oracle7 or earlier databases, it is not necessary to convert to the command script format.
- Because of the name change to the Oracle Server Manager's executable and listeners, only Oracle7 or earlier databases are supported by the parameter-based format.

The command script format is required for Oracle8 or later, or for advanced features such as Multiple Oracle Homes (introduced in Oracle 8.0.4).

- If you want to upgrade/migrate an Oracle7 database to Oracle8 and take advantage of the advanced features such as Multiple Oracle Homes, you must first convert the associated Oracle failover object to the command script format. Then follow the Oracle documentation to upgrade/migrate the Oracle database.
- To upgrade or migrate an Oracle7 database to Oracle8 you have to first remove the Oracle failover object from the failover group. Since Oracle does not support live upgrade or live migration, the Oracle database would be inaccessible during the upgrade/migration. Once the database conversion is complete you can put the object back into the original failover group.
- To convert a parameter-based Oracle failover object to a script-based format, On the cluster administrator screen in the Modify menu, choose Oracle Failover Object. Then click on the Convert to Command Scripts button. The system will automatically create the command scripts based on the existing parameters.

Note:

In DCNT V1.1 Service Pack 3 or later, an Oracle failover object can be of parameter-based or script-based format.

A script-based object can then be viewed as having two aspects -- one of parameters and the other of command scripts. For a newly created script-based object, the fields of its parameter aspect are empty. For a script-based object that has been converted from the parameter-based format, the command scripts are used and the parameter-aspect fields are kept intact, but ignored on the node running DCNT V1.1 Service Pack 3 or later. The parameter aspect of a script-based Oracle failover object will be picked up and used on a node running an earlier DCNT release. This increases the object availability in during a rolling upgrade.

Creating Command Scripts

To create a script-based Oracle failover object, you must supply command scripts. Below are samples of .cmd and .sql files and instructions on writing them.

Note

Sample .cmd and .sql files are provided in the Sample\Oracle directory of the Service Pack CD or the download.. These files are provided as templates that you can copy and modify so that you don't have to create these files from scratch.

► **To create command scripts for starting and stopping Oracle failover objects:**

Shown below are four example files, two each for starting and stopping Oracle failover objects. Both the start and stop command files use the following environment.

Cluster Nodes:	"NODE-0" and "NODE-1".
Drive Z:	The shared drive (identical drive letter assignment on both nodes), for files being shared in the cluster.
Name of the Oracle failover object:	Ora7 (Note: The associated Windows NT service for this instance is started automatically on both nodes.)
Oracle Server Manager:	SVRMGR23.EXE
Oracle installation:	C:\OraNT on NODE-0; D:\Ora73 on NODE-1
Name of the Oracle instance:	ORCL
Listener for the Oracle instance:	OracleTNSListener
Name of the Oracle database:	Oracle
Database parameter file:	Z:\OraNT\Database\initORCL.ora

Sample command script for starting an Oracle failover object

Shown below are two example files one is a .cmd file and one is a .sql file, Both use the same environment. The .cmd file launches the Oracle Server Manager to start the Oracle database and to start its listener (if used). The .sql file connects to the database and starts it.

Z:\OraNT\Dcnt\Ora7DcntStartOra.cmd

```
@echo*** Running "%0" on "%ComputerName%" for Oracle failover
object "Ora7"... [note1]
```

```

prompt $d $t $g$g [note2]
set ORACLE_SID=ORCL [note3]
:beginNode0
if not "%ComputerName%" == "NODE-0" goto endNode0 [note4]
    net start OracleTNSListener [note5]
        "C:\OraNT\bin\SVRMGR23.EXE"
    @"Z:\OraNT\Dcnt\Ora7Dcnt startOra.sql" [note6]
        goto theEnd
:endNode0
:beginNode1
if not "%ComputerName%" == "NODE-1" goto endNode1 [note4]
    net start OracleTNSListener [note5]
        "D:\Ora73\bin\SVRMGR23.EXE"
    @"Z:\OraNT\Dcnt\Ora7Dcnt startOra.sql" [note6]
        goto theEnd
:endNode1
:nodeUnknown
    @echo.
    @echo *** Error -- node "%ComputerName%" not a cluster
member!!
    @goto theVeryEnd
:theEnd
@echo.
@echo *** Done with "%0" on "%ComputerName%"
:theVeryEnd
prompt [note7]

```

Notes:

1. Use ECHO to capture the .cmd being executed and the computer name in the logfile.
2. This PROMPT statement gives a date/time stamp for each subsequent command line executed. It will also echo out the current drive and directory as the default prompt on Windows NT.
3. Set ORACLE_SID properly before launching Oracle Server Manager.
4. Compare the name of the current computer against that of the member node in the cluster – in CAPITALS. Note that %ComputerName% returns the name in upper case.

Using Oracle Failover Objects

5. Start the corresponding listener (only if a listener is in use). There may be other Oracle services that need to be started. Check your Oracle documentation for starting up an Oracle instance.
6. Launch the Oracle Server Manager with the startup .sql file on the shared drive. Note that the Oracle installation in the example is different on NODE-0 and NODE-1.
7. Use PROMPT as the last command, so that the date/time stamp of completion of the .cmd can be captured.

```
Z:\OraNT\Dcnt\Ora7DcntStartOra.sql
set ECHO ON[note1]
-- connect internal/*****@ORCL[note2]
set ECHO OFF[note3]
connect internal/oracle@ORCL
set ECHO ON[note4]
startup exclusive openOracle
pfile=Z:\OraNT\Database\initORCL.ora[note5]
exit
```

Notes:

1. Set ECHO ON at the very start, so that the statement being executed will be echoed out to the standard output. It's for trouble shooting if needed.
2. This remark indicates the next thing that is to be executed – connecting to the database. This line is to echo to the screen the user name (“internal” in the example) and the Oracle instance (ORCL in the example) that are in use.
3. Set ECHO OFF, so that the password (oracle in the example) is hidden from the standard output when it actually connects.
4. Set ECHO back ON, so that the following statements are echoed to standard output.
5. This “startup” statement opens the database “Oracle”, using a parameter file on one of the shared drives (Z:).

Sample command script for stopping an Oracle failover object

Two example files, one .cmd file and one .sql file, are shown below, using the same environment as the example above. The .cmd file is used to launch the Oracle Server Manager to shut down the Oracle database and to stop its listener (if used). The .sql file is to connect to the database and shut it down.

```

Z:\OraNT\Dcnt\Ora7DcntStopOra.cmd
@echo *** Running "%0" on "%ComputerName%" for Oracle failover
object "Ora7"...
prompt $d $t $g$g
set ORACLE_SID=ORCL
:beginNode0
if not "%ComputerName%" == "NODE-0" goto endNode0
    "C:\OraNT\bin\SVRMGR23.EXE"
@ "Z:\OraNT\Dcnt\Ora7Dcnt StopOra.sql"
    net stop OracleTNSListener[note1]
    goto theEnd
:endNode0
:beginNode1
if not "%ComputerName%" == "NODE-1" goto endNode1
    "D:\Ora73\bin\SVRMGR23.EXE"
@ "Z:\OraNT\Dcnt\Ora7Dcnt StopOra.sql"
    net stop OracleTNSListener
    goto theEnd
:endNode1
:nodeUnknown
    @echo.
    @echo *** Error -- node "%ComputerName%" not a cluster
member!!
    @goto theVeryEnd
:theEnd
@echo.
@echo *** Done with "%0" on "%ComputerName%"
:theVeryEnd
prompt

```

Notes:

1. Stop the corresponding listener (if a listener is in use). There may be other Oracle services that need to be stopped. Check your Oracle documentation for stopping an Oracle instance.

Using Oracle Failover Objects

```
Z:\OraNT\Dcnt\Ora7DcntStopOra.sql
set ECHO ON
-- connect internal/*****@ORCL
    set ECHO OFF
    connect internal/oracle@ORCL
    set ECHO ON

shutdown abort Oracle pfile=Z:\OraNT\Database\initORCL.ora [note1]
exit
```

Notes:

This statement shuts down the database with the option -- abort, so that failover or failback of the Oracle object will not be pending on an open connection. The other two options (**normal and immediate**) may cause problems in such a situation.

Hints and Tips for Creating Command Scripts

- The command script for "Command Script to Start Oracle" is used to start the Oracle database and to start any associated Windows NT services for the Oracle instance. The command script for "Command Script to Stop Oracle" is used to shutdown the database and to stop any Windows NT services for the instance.
- Use the proper ECHO ON/OFF instruction to allow proper logging and tracing for trouble shooting. For example, using the "echo off" instruction for the .cmd files is not a good idea since it leaves no traces if a failure occurs. Use SET ECHO ON as the first statement in the .sql file to leave traces. However, use SET ECHO OFF before connecting to the database so that a password or any sensitive information does not appear in the log or trace.
- It's highly recommended that all the .cmd and/or .sql files or any other files used in the command script are placed on the same shared drive that all the Oracle database and parameter files reside on. All the paths given in any such file are node-specific. For example, if one assigns the command scripts on C:\OraDCNT, then the command script files should be on C:\OraDCNT for both cluster nodes.
- Test the scripts before putting into production use.
- General tips for creating a .cmd file
 - * Use double quotes (see Example) if there is a space in the path of the files referenced in the command scripts.
 - * Use capital letters for computer names to enable you to compare with the Computer Name environment variable.

Using Oracle Failover Objects

- Sample `.cmd` and `.sql` files are provided in the `Sample\Oracle` directory of the Service Pack CD or the download. These files are provided as templates that you can copy and modify.

Using Script-based Oracle Failover Object Screens

The following procedures allow you to:

- create a new Oracle failover object in the script-based format
- modify an existing Oracle failover object
- convert a previously existing parameter-based Oracle failover object to the command script-based format.
- delete an existing failover object.
- activate the changes to an Oracle failover object

► **To create a script-based Oracle failover object:**

Caution

Do a full upgrade of both cluster nodes to DCNT version 1.1 SP3 before creating script-based Oracle failover objects. Failure to do so may cause disruption in the availability of the Oracle databases.

Use this screen to create an Oracle failover object. This screen uses the script-based format.

1. Create Oracle `.cmd` start and stop files (usually `.cmd` or `.bat` files). The start command script starts the Oracle database and its associated Windows NT services (such as listener). The stop command script shuts down the database and stops its associated Windows NT services. After creating these scripts, see samples of command scripts in Before You Start section and save them in a location and note the path address. This path must be the same with respect to both cluster nodes. It is recommended that these files be placed on one of the shared drives.
2. From the Create menu, select Oracle Failover Object or choose Oracle Failover Object from the right-click drop-down menu. Then select Script-Based Oracle Failover Object.
3. The following screen is displayed.

The screenshot shows a dialog box titled "Create Oracle Failover Object". It contains the following fields and buttons:

- Oracle Failover Object Name: [Text Input]
- Oracle Pipe Name (Optional): [Text Input]
- Oracle Start Command Script: [Text Input] with a "Browse..." button to its right.
- Oracle Stop Command Script: [Text Input] with a "View" button to its right.
- Command Output Logfile (Optional): [Text Input]
- Buttons at the bottom: "OK", "Cancel", and "Help".

1. Enter Oracle failover object name.
2. Enter Oracle pipe name (this optional entry, is only needed when *Named Pipes* is used in the Oracle instance).

Note

If you are running Cluster Administrator locally, you can use the Browse button to help find and write the path to your start and stop command files.

3. Enter the Oracle start command script name path for the start command script you created in step 1. This path must be the same for both cluster nodes. It is recommended that these files be placed on one of the shared drives.
4. Enter the Oracle stop command script name path for the stop command script you created in step 1. This path must be the same for both cluster nodes. It is recommended that these files be placed on one of the shared drives.
5. Enter Command Output Logfile (this is optional). However it is advisable to create a Command Output Logfile to capture all command output. This facilitates the troubleshooting for the cause of a failure to start or stop an object. Use a text editor such as Notepad to view the output logfile.
6. Use the View button to launch notepad.exe for viewing and editing the command script files or Command Output Logfile. Note that if Notepad is unable to find the file, it may be a result of the path destination drive not being controlled by or accessible to the node (may be controlled by the other node).

Using Oracle Failover Objects

Note

If you are running Cluster Administrator remotely, the Browse and View buttons are disabled.

7. Click OK to accept changes. Remember the new Oracle object must be added to a failover group before it is activated.
8. Choose Cancel to exit without saving changes.

► **To modify a script-based Oracle failover object:**

Caution

Do a full upgrade of both cluster nodes to DCNT version 1.1 SP3 before modifying script-based Oracle failover objects. Failure to do so may cause disruption in the availability of the Oracle databases.

Note

If you have previously converted a failover object to the new command script-based format or created a new script-based Oracle object from scratch, use the following procedure. If however, the failover object was created in the old, parameter-based format, DCNT detects it and presents you with the old Modify screen with a new button allowing you to convert to the command script format (See next section, “To convert an existing parameter-based Oracle object to the script-based format”).

1. On the Modify menu, select Oracle Failover Object or choose Modify Oracle Failover Object from the right-click drop-down menu.
2. The following Modify Oracle Failover Object screen is displayed.
3. Modify any data you want and click OK to save. See “To create a script-based Oracle failover object” section for details on each field and “To activate changes to an Oracle failover object” section for details on having your changes take effect.

The screenshot shows a dialog box titled "Modify Script-Based Oracle Failover Object". It contains the following fields and buttons:

- Oracle Failover Object Name: [Text Input]
- Oracle Pipe Name (Optional): [Text Input]
- Oracle Start Command Script: [Text Input] with a "Browse..." button to its right.
- Oracle Stop Command Script: [Text Input] with a "View" button to its right.
- Command Output Logfile (Optional): [Text Input]
- Buttons at the bottom: OK, Cancel, Help.

Using Parameter-based Oracle Failover Object Screens

- ▶ **To create a parameter-based Oracle object:**
 1. On the Create menu, from the Class or Cluster view, select the Oracle Failover Object menu item. Then select the Parameter-Based Oracle Failover Object sub-menu item.
 2. The following screen is displayed.

Create Oracle Failover Object [X]

Note: Parameter-based objects will not work with Oracle 8 or above. In this case, create a script-based Oracle failover object.

Oracle Failover Object Name:

Comment:

Oracle Username:

Oracle Password:

Oracle User Type: ▼

Oracle sid:

Oracle Pipe Name:

Oracle sid Parameter File:

Oracle Database:

Oracle Listener Name:

See DCNT Administrator's Guide for details on each field.

► **To modify a parameter-based Oracle object:**

1. From Class or Cluster View, select the Oracle failover object you want to modify.
2. On the Modify menu, select Oracle Failover Object.
3. The following screen is displayed if the object is in the parameter-based format.

Modify Parameter-Based Oracle Failover Object [X]

Note: Parameter-based objects will not work with Oracle 8 or above. In this case, create a script-based Oracle failover object.

Oracle Failover Object Name:

Comment:

Oracle Username:

Oracle Password:

Oracle User Type:

Oracle sid:

Oracle Pipe Name:

Oracle sid Parameter File:

Oracle Database:

Oracle Listener Name:

See DCNT Administrator's Guide for details on each field.

Using Oracle Failover Objects

► To convert an existing parameter-based Oracle object to the script-based format:

Use this procedure to convert an existing Oracle7 failover object to an Oracle8 failover object or to convert an existing Oracle7 failover object to a command script-based Oracle7 failover object.

1. From Class or Cluster View, select the Oracle failover object you want to convert.
2. On the Modify menu, select Oracle Failover Object or choose Modify Oracle Failover Object from the right-click drop-down menu.
3. The following screen is displayed if the object is in the parameter-based format.

Modify Parameter-Based Oracle Failover Object [X]

Note: Parameter-based objects will not work with Oracle 8 or above. In this case, create a script-based Oracle failover object.

Oracle Failover Object Name:

Comment:

Oracle Username:

Oracle Password:

Oracle User Type:

Oracle sid:

Oracle Pipe Name:

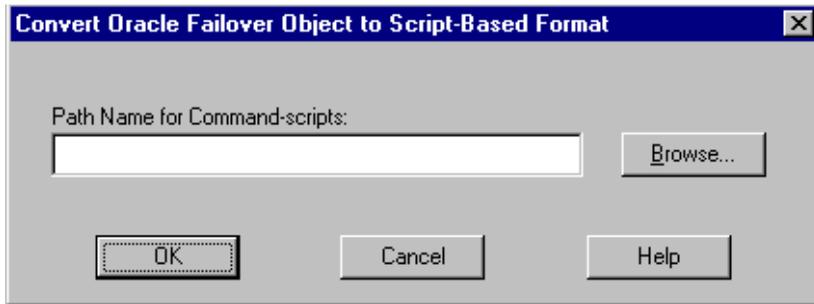
Oracle sid Parameter File:

Oracle Database:

Oracle Listener Name:

If desired, see DCNT Administrator's Guide for details on each field.

4. Click the Convert to Command Scripts button if you wish to convert to the command script-based format. The following screen is presented.



5. Enter path name for the command script.
6. Use the Browse button to select a new command script path. If necessary, use it to change the final path to ensure that both cluster nodes have access to the command script files

Note

If you are running Cluster Administrator remotely, the Browse button is disabled.

7. If you have clicked the OK button the following screen is displayed

Using Oracle Failover Objects

Convert Oracle Failover Object to Script-Based Format

Oracle Failover Object Name:

Oracle Pipe Name (Optional):

Oracle Start Command Script:

Oracle Stop Command Script:

Command Output Logfile (Optional):

Old Object Information

Oracle Username:	internal
Oracle User Type:	ORA_NORMAL
Oracle sid:	Ora7
Oracle Pipe Name:	
Oracle sid Parameter File:	Y:\OraDCNT\Ora7_init.ora
Oracle Database:	Oracle
Oracle Listener Name:	Listener
Comment:	

Note

Note the old object information at the bottom of the screen which is shown if this is a conversion on a previously existing Oracle7 failover object.

1. Oracle Failover Object Name is the same as in the parameter-based format. Change to the name is not allowed..
2. Change Oracle pipe name if necessary (this is an optional entry, it is only needed when Name Pipes is used in the Oracle instance).

Note

If you are not running Cluster Administrator remotely, you can use the Browse button to help find and write the path to your start and stop command files.

3. Change the Oracle start command script name path for the start command script if you want to use a different script than the one created by the previous conversion process.

This path must be the same for both cluster nodes. It is recommended that these files be placed on one of the shared drives.

4. Change the Oracle stop command script name path for the stop command script if you want to use a different script than the one created in the previous conversion process. This path must be the same for both cluster nodes. It is recommended that these files be placed on one of the shared drives.
5. Enter Command Output Logfile (this is optional). However it is advisable to create a Command Output Logfile to capture all command output. This facilitates the troubleshooting for the cause of a failure to start or stop an object. Use a text editor such as Notepad to view the output logfile.
6. Use the View button to launch notepad.exe for viewing and editing the command script files or Command Output Logfile. Note that if Notepad is unable to find the file, it may be a result of the path destination drive not being controlled by the node (online on the other node).

Note

If you are running Cluster Administrator remotely, the Browse and View buttons are disabled.

7. Click OK to accept changes. Remember it is necessary to failover/failback or remove objects from the failover group to activate changes. See “To activate changes to Oracle failover object” section for details.
8. Choose Cancel to exit without saving changes. The Oracle failover object will stay in the parameter-based format.

► **To delete an Oracle failover object:**

1. From the Class or Cluster View, select the Oracle failover object you want to delete.
2. From the Delete menu, choose Oracle Failover Object or choose Delete Oracle Failover Object from the right-click drop-down menu, or click the Delete button on the toolbar. The Oracle failover object is deleted from the cluster database and is unavailable for adding to a group.

► **To activate changes to an Oracle failover object:**

1. To activate any of the changes to an Oracle failover object, you must either failover or failback the Oracle failover object or remove the failover object from a group and then put it back into the group. If you choose to use the remove/replace procedure, be aware you may want to retain the failover order you had previously established.