

# DIGITAL Clusters for Windows NT

---

## Administrator's Guide

Part Number: AA-QVUTB-TE

**February 1997**

This guide provides a conceptual overview of the DIGITAL Clusters for Windows NT™ product; gives procedures for using Cluster Administrator to perform routine cluster administration tasks; presents procedures for configuring applications to take advantage of clustering; and offers troubleshooting procedures.

**Revision/Update Information:** This is a revised guide.

**Operating System and Version:** Microsoft® Windows NT 4.0  
with Service Pack 2  
Microsoft Windows NT 3.51  
with Service Pack 5

**Software Version:** DIGITAL Clusters for Windows NT  
Version 1.1  
DIGITAL Clusters for Windows NT  
Version 1.0 with Service Pack 2

**Digital Equipment Corporation  
Maynard, Massachusetts**

---

**February 1997**

Digital Equipment Corporation makes no representations that the use of its products in the manner described in this publication will not infringe on existing or future patent rights, nor do the descriptions contained in this publication imply the granting of licenses to make, use, or sell equipment or software in accordance with the description.

Possession, use, or copying of the software described in this publication is authorized only pursuant to a valid written license from DIGITAL or an authorized sublicensor.

© Digital Equipment Corporation 1997. All rights reserved.

The following are trademarks of Digital Equipment Corporation: Alpha AXP, AlphaGeneration, AlphaServer, AlphaStation, DIGITAL, OpenVMS, Prioris, ServerWORKS, StorageWorks, and the DIGITAL logo.

The following are third-party trademarks:

Adaptec is a trademark of Adaptec Inc.

Banyan is a registered trademark and Street Talk is a trademark of Banyan Systems, Inc.

CLARiON is a registered trademark of Data General Corporation.

Intel is a registered trademark of Intel Corporation.

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

Macintosh is a registered trademark of Apple Computer, Inc.

NetBIOS is a trademark of Micro Computer Systems, Inc.

Netscape and Netscape Navigator are trademarks of Netscape Communications Corporation.

Novell is a registered trademark of Novell, Inc.

NT is a trademark of Northern Telecom Limited.

Oracle, SQL\*DBA, SQL\*Net, and SQL\*Plus are registered trademarks and Oracle7, Oracle7 Workgroup Server, Oracle Enterprise Manager, Oracle Names, Oracle Network Manager, Oracle Server, and Oracle Server Manager are trademarks of Oracle Corporation.

OS/2 is a registered trademark of International Business Machines Corporation.

SQL Server is a trademark of Sybase, Inc.

Sun is a registered trademark of Sun Microsystems, Inc.

UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Ltd.

Windows NT is a trademark, and Microsoft, MS-DOS, Windows, and Windows 95 are registered trademarks of Microsoft Corporation.

All other trademarks and registered trademarks are the property of their respective holders.

---

# Table of Contents

## 1 Introduction

Overview .....	1-1
What Is DIGITAL Clusters for Windows NT?.....	1-3
Example DIGITAL Clusters for Windows NT Configuration .....	1-4
Benefits of DIGITAL Clusters for Windows NT .....	1-5
High Availability .....	1-5
Scaleability .....	1-6
Industry Standards and Commodity Hardware.....	1-6
How DIGITAL Clusters Fits in a PC LAN Environment.....	1-7
Supported Server Architectures .....	1-8
Supported Clients.....	1-8
Physical Connections.....	1-9
Management of Resources and Services.....	1-9

## 2 Understanding Cluster Concepts

Cluster Concepts and Terminology .....	2-1
Failover .....	2-1
Spare Capacity .....	2-1
Failover Objects.....	2-2
Failover Groups .....	2-2
Failover Policy.....	2-2
Failback.....	2-3
Example of Database Failover.....	2-3

### 3 Configuring Database Software for Failover

Microsoft SQL Server Installation and Configuration Requirements.....	3-1
New Features and Enhancements in DIGITAL Clusters 1.1 .....	3-1
Prerequisite Information .....	3-2
Software Requirements .....	3-2
Designating Primary and Failover Servers .....	3-2
SQL Server Database Failover .....	3-2
Access to Shared Disks .....	3-3
Configuration Requirements .....	3-3
Configuration and Run-Time Recommendations.....	3-3
Restrictions.....	3-4
Configuring SQL Server for Failover .....	3-4
Configuring Clients to Access SQL Server Databases .....	3-6
Verifying SQL Server Failover.....	3-9
Adding a Shared Disk to a Failover Group for SQL Server .....	3-11
Unenrolling an SQL Server Database from Failover Support .....	3-11
Expanding or Shrinking an SQL Server Database.....	3-12
Moving an SQL Server Database and Disk to a Different Failover Group .....	3-13
Moving SQL Server Databases to a New Cluster.....	3-13
Resetting the Suspect Status of an Unenrolled SQL Server Database .....	3-14
Oracle7 Server Installation and Configuration Requirements .....	3-14
New Features in DIGITAL Clusters 1.1 .....	3-15
Prerequisite Information .....	3-15
Software Requirements .....	3-15
Oracle Server Database Failover .....	3-15
Configuration Requirements .....	3-16
Creating and Configuring an Oracle Instance for Failover.....	3-16
Example listener.ora File.....	3-21
Configuring Clients to Access Oracle Server Databases .....	3-22
Configuring the tnsnames.ora File .....	3-23
Client Connections During A Failover .....	3-25
Initiating Manual Failover of an Oracle Instance .....	3-25

### 4 Configuring Lotus Notes for Failover

About Lotus Notes.....	4-1
Domino 4.5 Installation and Configuration ( <i>VI.1</i> ).....	4-2
Clustering Methods .....	4-2
IP Failover for a Single Domino Server .....	4-2
IP Failover and Partitioned Domino Servers.....	4-3
Before You Start .....	4-4
Check Windows NT Settings .....	4-5

Create Failover Groups .....	4-5
Installation .....	4-5
Software Requirements .....	4-5
Setting Up Partitioned Domino Servers on the First Cluster Server.....	4-6
Install the First Domino Server and Notes Client on the First Cluster Server.....	4-6
Register Additional Domino Servers on the First Cluster Server.....	4-11
Install a Second Domino Server on the First Cluster Server.....	4-14
Install the Domino Servers as Windows NT Services.....	4-17
Setting Up Partitioned Domino Servers on the Second Cluster Server .....	4-18
Creating Script Files and Completing the Failover Groups .....	4-19
Creating Failover Scripts.....	4-19
Completing the Failover Groups .....	4-20
Adding Failover Support for the Domino Web Server.....	4-20
Lotus Notes 4.11 Installation and Configuration ( <i>V1.0 SP1</i> ) .....	4-21
Software Requirements.....	4-21
Clustering Methods.....	4-21
Setting Up Databases on Both Servers.....	4-22
Failover Method: A Script and Notes Directory Pointer Files.....	4-23
Installation .....	4-23
Before You Start.....	4-23
Install Lotus Notes on the First Cluster Server .....	4-24
Registering the Second Lotus Notes Server.....	4-26
Installing Lotus Notes on the Second Cluster Server.....	4-28
Registering Users and Installing Notes Workstation Clients .....	4-30
Creating the Directory Pointer Files .....	4-30
Sample Script Files .....	4-32
Lotus Notes 4.11 Failover Issues.....	4-34

## 5 Configuring Web Servers for Failover

Overview .....	5-1
Before You Start.....	5-2
Microsoft IIS Installation and Configuration .....	5-3
Software Requirements.....	5-3
Installation .....	5-3
Configuration.....	5-4
Netscape Enterprise Server Installation and Configuration.....	5-5
Software Requirements.....	5-5
Installation .....	5-6
Configuration.....	5-7

## 6 Getting Started with Cluster Administrator

Cluster Administrator Overview .....	6-1
Starting Cluster Administrator .....	6-2
Quitting Cluster Administrator .....	6-2
Displaying the Cluster Topology .....	6-2
Displaying the System View .....	6-3
Displaying the Cluster View .....	6-4
Displaying the Class View .....	6-5

## 7 Configuring Your Cluster

Configuration Steps .....	7-1
---------------------------	-----

## 8 Managing a Cluster

Managing a SCSI Adapter Configuration .....	8-2
Managing Disk Aliases .....	8-5
Managing an Event Log .....	8-6
Managing the Log Disk .....	8-8
Managing Manual Failover .....	8-9
Managing SQL Server Databases .....	8-11

## 9 Working with Failover Objects and Groups

Working with an Oracle Failover Object .....	9-1
Creating an Oracle Failover Object .....	9-2
Modifying an Oracle Failover Object .....	9-4
Deleting an Oracle Failover Object .....	9-6
Working with a Script Failover Object .....	9-6
Creating a Script Failover Object .....	9-7
Script Failover Object Command Restrictions .....	9-8
Modifying a Script Failover Object .....	9-9
Deleting a Script Failover Object .....	9-9
Working with SQL Server Failover Objects and Groups .....	9-10
Creating an SQL Server Failover Object .....	9-10
Modifying an SQL Server Failover Object .....	9-10
Deleting an SQL Server Failover Object .....	9-11
Working with an IP Failover Object .....	9-12
Creating an IP Failover Object .....	9-13
Modifying an IP Failover Object .....	9-14
Deleting an IP Failover Object .....	9-16
Working with a Failover Group .....	9-17

Creating a Failover Group .....	9–17
Modifying a Failover Group .....	9–21
Deleting a Failover Group .....	9–24

## 10 Application Considerations

Application Handling During a Failover .....	10–2
Failover of Client Connections .....	10–2
Example of Client Connection Failover .....	10–2
Database Application Failover .....	10–4
Database Client Application Failover .....	10–5
SQL Client Application Considerations .....	10–5
Additional Client Application Considerations .....	10–5
Open Files, Named Pipes, and IP Socket Connections .....	10–5
Failover Times .....	10–6
What the User Sees During a Failover .....	10–6
Supported Clients .....	10–6
Clients Not Using the Cluster Alias .....	10–6

## 11 Troubleshooting

Configuration Problems .....	11–1
Where are my disks? I can't create a failover group. ....	11–1
My group won't come on line. ....	11–3
Failover Problems .....	11–4
Failover doesn't work .....	11–4
Client Problems .....	11–5
My client doesn't see any clusters. ....	11–5
My client doesn't see the cluster it needs. ....	11–5
My client can't access cluster resources. ....	11–6
My client hangs after failover. ....	11–6
Database Problems .....	11–7
My database isn't available .....	11–7
My database failover group won't come on line. ....	11–8
My Oracle7 Server won't fail over to the other server system. ....	11–9
My Oracle7 Server is running but the client can't access it. ....	11–9

## A Troubleshooting Tools and Resources

Software Utilities .....	A-1
regedt32 .....	A-1
Disk Administrator.....	A-1
Services Applet .....	A-2
NET SHARE Command .....	A-2
NET VIEW Command.....	A-2
NETMON Network Monitor Utility .....	A-3
CLUIVP Utility.....	A-4
CLUXFER Utility.....	A-4
Cluster Monitor Utility.....	A-5
Cluster Monitor Utility Display .....	A-7
Registry.....	A-8
DIGITAL Clusters Registry Keys.....	A-8
DIGITAL Clusters Failover Management Database (CFMD) Key.....	A-8
DIGITAL Clusters Port Driver (CluPort) Key .....	A-8
DIGITAL Clusters Disk Driver (CluDisk) Key .....	A-9
DIGITAL Clusters File System (CFS) Key.....	A-9
DIGITAL Clusters Log Watch Key .....	A-9
DIGITAL Clusters Failover Manager Key.....	A-9
DIGITAL Clusters Name Service Key.....	A-10
SCSI Device Map Key .....	A-10
DIGITAL Clusters Registry Key Tuning Parameters .....	A-10
CfmdTrace.....	A-10
CisTrace.....	A-11
ConnectionTimeout .....	A-12
FailoverEvaluateDelay .....	A-13
ReconnectWait .....	A-14
LogLevel .....	A-15
ClusterName.....	A-16
DisableFailoverNetDelay .....	A-16
DiskArbitrationInterval .....	A-17
DiskErrorThreshold.....	A-18
DiskErrorSeparation.....	A-19
FmTraceOutput .....	A-19
FmTrace .....	A-20
FmLogLevel.....	A-21
FmTraceVerbosity.....	A-22
DIGITAL Clusters Client Tuning Parameters .....	A-23
Windows 95 Clients.....	A-23
Adjusting the Time for Trying Network Connections.....	A-23
Windows NT Clients.....	A-24

Adjusting the Time for Trying Network Connections .....	A-24
TranslationResponseTimeout Parameter .....	A-24
SolicitNameServerTimeout .....	A-25
Windows for Workgroups Clients .....	A-25
Adjusting the Time for Trying Network Connections .....	A-25
Trace Log .....	A-26
Event Log .....	A-27
Blue Screen Messages .....	A-27

## **B Registry Snapshots**

DIGITAL Clusters Failover Management Database (CFMD).....	B-1
DIGITAL Clusters Port Driver (CluPort).....	B-7
DIGITAL Clusters Disk Driver (CluDisk) .....	B-8
DIGITAL Clusters File System (CFS) .....	B-9
DIGITAL Clusters Failover Manager.....	B-10
DIGITAL Clusters Name Server .....	B-11
Log Watch .....	B-12
SCSI Device Map .....	B-13

## **C DIGITAL Clusters Failover Manager Trace Log Example**

## **D DIGITAL Clusters Event Logs**

System Event Log.....	D-1
Application Event Log .....	D-2

## **Glossary**

## **Index**

## Figures

Typical DIGITAL Clusters for Windows NT Configuration .....	1-4
Database Failover Example—Normal .....	2-4
Database Failover Example—Failure .....	2-5
IP Failover Model for a Single Domino Server .....	4-3
IP Failover Model for Partitioned Domino Servers .....	4-4
Sample Lotus Notes 4.11 Failover Configuration .....	4-22
Pointer Files: Database Access Before a Failover .....	4-31
Pointer Files: Database Access After a Failover .....	4-31
NT Cluster Configuration Before Client Connection Failover .....	10-3
NT Cluster Configuration After Client Connection Failover .....	10-4

---

# About This Guide

This guide provides a conceptual overview of the DIGITAL Clusters for Windows NT™ product; gives procedures for using Cluster Administrator to perform routine cluster administration tasks; presents procedures for configuring applications to take advantage of clustering; and offers troubleshooting procedures.

## Audience

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

## Organization

This guide consists of eleven chapters and four appendices, as follows:

- |           |  |
|-----------|--|
| Chapter 1 | Provides a conceptual overview of the DIGITAL Clusters for Windows NT product and describes key features and benefits.   |
| Chapter 2 | Discusses basic cluster concepts and terminology. It also presents a database failover example.  |
| Chapter 3 | Presents the database software installation and configuration steps you must complete before using Cluster Administrator to configure for failover of Microsoft SQL Server™ and Oracle7™ Server. |
| Chapter 4 | Provides information on installing and setting up Lotus Notes in a DIGITAL Clusters for Windows NT environment.  |
| Chapter 5 | Describes how to install and configure two web servers for IP failover: Microsoft Internet Information Server (IIS) and Netscape Enterprise Server.  |

- Chapter 6 Introduces Cluster Administrator. It gives step-by-step instructions on how to start and quit Cluster Administrator and how to display the cluster topology.
- Chapter 7 Outlines the steps you need to perform to complete your initial cluster configuration.
- Chapter 8 Presents step-by-step procedures on how to manage a cluster using Cluster Administrator. Topics covered include managing an adapter configuration, managing disk aliases, managing the event log, managing the log disk, managing manual failover, and managing SQL Server databases.
- Chapter 9 Describes how to create, modify, and delete Oracle failover objects, script failover objects, SQL failover objects, IP failover objects, and failover groups using Cluster Administrator.
- Chapter 10 Discusses client application failover considerations. Specific examples of client applications are included.
- Chapter 11 Provides troubleshooting procedures for commonly encountered problems in the DIGITAL Clusters for Windows NT environment.
- Appendix A Describes some of the software tools and system resources that you can use to diagnose problems with your cluster.
- Appendix B Contains snapshots of those parts of the Windows NT Registry that pertain to DIGITAL Clusters for Windows NT.
- Appendix C Contains an example of a typical DIGITAL Clusters Failover Manager trace log.
- Appendix D Contains examples of typical DIGITAL Clusters system and application event logs.

## Conventions

The following conventions are used in this guide:

Convention	Meaning
DIGITAL Clusters 1.1	Refers to DIGITAL Clusters for Windows NT Version 1.1.
<b>1.1</b>	Indicates a feature available in DIGITAL Clusters for Windows NT Version 1.1.
<b>Bold</b>	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 Configuration and Installation Guide*
- Online help

This chapter provides a conceptual overview of the DIGITAL Clusters for Windows NT product and describes key features and benefits.

## Overview

The explosive growth of Windows NT as an enterprise-level operating system has generated a demand for new high-availability tools and system management features. These tools, based on UNIX® and other operating systems, are available in well-established computing environments. DIGITAL delivers advanced Windows NT solutions today.

Clustering technology is well understood by today's UNIX and OpenVMS™ system administrators. Multiple systems are grouped together to appear as a single system to users and to other processes. The advantages of clustering are:

- High availability of system resources. The work load of a failed system is assumed by its counterpart system to ensure continuous, uninterrupted services to end users and applications.
- Improved scalability. New resources can be added incrementally to the cluster, which results in a cost-effective growth path to high performance.
- Reduced system management costs. Clustering makes it easier to manage multiple systems, reduces the costs associated with replicated data, and allows specialized peripherals to be shared by more users.

DIGITAL Clusters for Windows NT introduces the benefits of clustering technology to today's PC client/server local area network (LAN) environments.

## Overview

DIGITAL Clusters 1.0 focused on minimizing the down time caused by software, network and system failures, and:

- Offered a low-cost, high-availability solution for PC client/server LANs.
- Was based on industry-standard hardware components and industry-standard software.
- Supported Intel® and Alpha processor-based architectures. (See the *Configuration and Installation Guide* for details on supported hardware configurations.)
- Supported failover of the NTFS file system, Microsoft SQL Server, and Oracle7 Workgroup Server.
- Supported generic application failover for additional server-based applications.
- Allowed clients to access the cluster as if it were a single system through a common cluster name.

DIGITAL Clusters 1.1 and 1.0 with Clusters Service Pack 2 (SP2) build on DIGITAL Clusters 1.0 by adding the following new features and enhancements. The features are available in both DIGITAL Clusters 1.1 and 1.0 with SP2 unless noted otherwise:

- Support for Windows NT 4.0 with NT SP2 (DIGITAL Clusters 1.1 only).
- Internet Protocol (IP) address failover support for applications typically accessed by an IP address (DIGITAL Clusters 1.1 only). See Chapters 3, 4, 5, and 9, and Chapter 1 of the *Configuration and Installation Guide* for details.
- Enhanced Microsoft SQL Server 6.5 support, including:
  - Support for multiple, independent SQL Server databases
  - Ability to run SQL Server simultaneously on both cluster servers
  - Automated configuration of SQL Server databases
  - IP failover support for SQL Server clients using the TCP/IP network protocol (DIGITAL Clusters 1.1 only)

See Chapter 3 for details.

- Enhanced Oracle7 Server support, including:
  - Failover support for Oracle7 Server 7.3
  - IP failover support for Oracle7 Server clients using the TCP/IP network protocol (DIGITAL Clusters 1.1 only)

See Chapter 3 for details.

- Support for failover of Web servers (DIGITAL Clusters 1.1 only). See Chapter 5 for details.

- Support for failover of the following versions of Lotus Notes:
  - Lotus Notes® 4.11
  - Lotus Notes 4.5 (DIGITAL Clusters 1.1 only)See Chapter 4 for details.
- Faster access to file shares after a failover.
- Faster access to file shares on systems outside the cluster.
- Support for large numbers of cluster file shares.
- Support for hidden file shares on a cluster disk.
- Cluster Administrator enhancements and additions.
- Rolling upgrade support from Windows NT 3.51 based clusters (DIGITAL Clusters 1.0 and 1.0 with SP2) to Windows NT 4.0 based clusters, while maintaining existing cluster configurations. See Chapter 7 of the *Configuration and Installation Guide* for details.

## What Is DIGITAL Clusters for Windows NT?

DIGITAL Clusters for Windows NT is a general-purpose, high-availability, scalable solution for today's PC client/server LANs. It couples two Windows NT Servers in an enterprise LAN, via a shared SCSI bus, to create a single-system environment, or *cluster*. End-user clients have access to all the cluster resources, such as shared disks, file shares, and database applications, without having to know the names of the individual servers in the cluster. In the event that one server system fails, the second cluster server will immediately assume its work load, reconnect clients, and migrate shared storage and file shares.

The software design of DIGITAL Clusters is extensible, flexible, and hardware independent. The product delivers critical, high-availability features to Windows NT Servers. Due to its extensible architecture, advanced clustering capabilities can be added easily over time, built on the same core product functionality. Software control of clustering gives users flexibility in integrating clustering with their existing application environments.

Enterprises have come to rely on high levels of availability from information systems. They are unwilling and unable to tolerate down time as they reengineer their businesses and expand their global operations. Performance statistics show that the best current option—Windows NT Server on a state-of-the-art *symmetric multiprocessing (SMP)* processor—is still likely to have about 90 hours down time per year. With the improvements that have been made to hardware, these failures are not the major reason for system down time. Instead, software failures, software maintenance, and planned upgrades are currently the major source.

DIGITAL Clusters for Windows NT, with proper power management, supply engineering, and redundant communications, can reduce the average down time to less than 12 hours per

# Overview

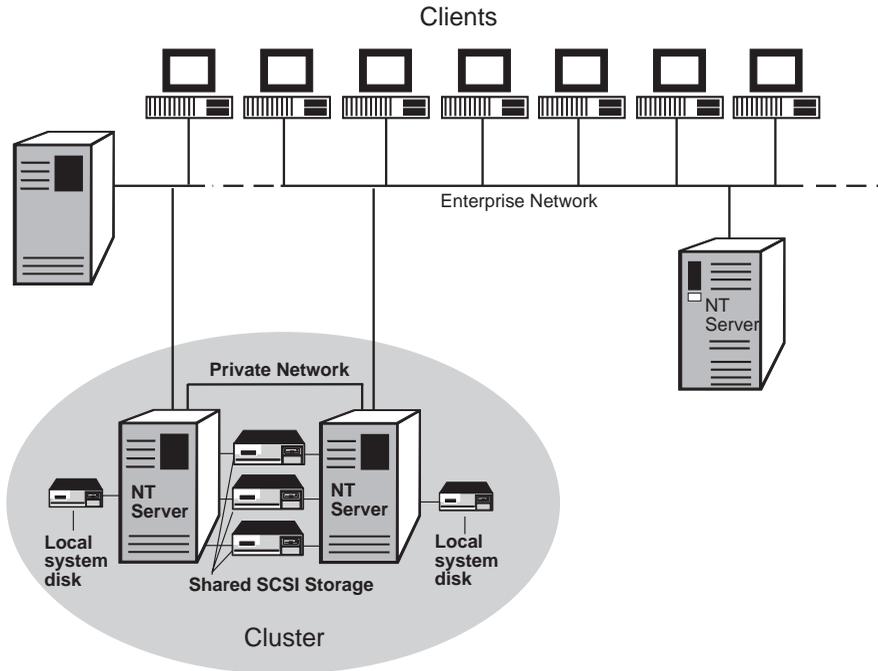
year, and in some cases, less than 1 hour per year. It can keep end users productive on business-critical applications and databases while still allowing for growth and flexibility.

## Example DIGITAL Clusters for Windows NT Configuration

Let's look at how a cluster fits into today's PC LAN environment. In a typical PC LAN configuration, a variety of client desktop systems can access several Windows NT Servers over the network.

Now we introduce the concept of a cluster. As illustrated in the following figure, a pair of Windows NT Servers, each running the cluster software, is connected via a shared SCSI bus. Multiple shared storage devices can be connected to the shared SCSI bus. The clustered servers communicate over the network. In addition, we include a second Ethernet connection that runs between the two servers.

## Typical DIGITAL Clusters for Windows NT Configuration



ZK-8758A-FH5

DIGITAL strongly recommends a second, dedicated Ethernet connection between the paired clustered servers for the following reasons:

- It eliminates another single point of failure in the cluster by ensuring that the two servers will always be able to communicate with one another.
- It avoids any unnecessary failovers that may occur in the event of a network partition where one server is under the misconception that the other server has failed, and therefore causes a failover.

The clients view the cluster as a single system without regard to cluster server names or which disk is managed by which server. The cluster software directs the clients to the correct disk or file share. The resources and services of the cluster are available to the client system as if they were local.

## Benefits of DIGITAL Clusters for Windows NT

DIGITAL Clusters for Windows NT offers system-level high availability at very low cost using industry-standard, commodity components. A cluster is addressed by clients as if it were a single server. Similarly, the cluster configuration is managed as if it were a single server. Clustering provides high levels of availability through redundant CPUs, storage, and data paths.

### High Availability

High availability is made possible through failover capabilities. Simply stated, failover quickly redirects interrupted services and resources to clients using a backup path. For failover of named pipes, network file shares, and other applications that use NetBIOS emulated over NetBEUI, DIGITAL Clusters uses a common cluster name, or *cluster alias*, which makes the location of the cluster resources transparent to the end user. DIGITAL Clusters 1.1 also introduces support for IP failover of socket-based applications. With proper configuration, your IP socket-based applications (including Microsoft SQL Server 6.5, Oracle7 Server 7.3, Lotus Notes 4.5, Microsoft Internet Information Server (IIS), and Netscape Enterprise Server 2.01) will fail over to the alternate cluster server in the event of a primary server failure. In each case, the client does not need to know how the cluster is configured or how the work load is divided among servers. The benefit to users is that they can focus on tasks rather than technology.

Enterprises can leverage their application software investments because DIGITAL Clusters for Windows NT works with both packaged and in-house applications. Depending on the application, either the user sees a message requesting a retry or the application continues uninterrupted. In-house applications also can take advantage of cluster failover capabilities by using the generic application failover feature.

## Benefits of DIGITAL Clusters for Windows NT

The level of availability provided by DIGITAL Clusters for Windows NT compares well with fault-tolerant systems. DIGITAL Clusters is more cost-effective because it does not require a complete mirrored backup of the primary system. Although the “hot standby” in fault-tolerant systems furnishes nonstop availability, it does so at the cost of a backup system that does not add any computing capacity. DIGITAL Clusters, in contrast, allows users to partition work loads and use both servers. You have the benefit of greater capacity, and you maximize your investment in current resources. DIGITAL Clusters reduces your investment in new resources because it does not depend on custom hardware or proprietary interconnects.

### Scaleability

DIGITAL Clusters for Windows NT is a scaleable solution. You can add capacity to a cluster in several dimensions. I/O and storage resources and application services can be added incrementally to efficiently and cost-effectively meet the dynamic needs of an enterprise.

The scaleability of DIGITAL Clusters rests with the partitioned data model of its software architecture. This model delivers numerous benefits, including:

- Greater flexibility. Dividing the work load into smaller components provides greater flexibility. DIGITAL Clusters supports partitioning of the work load down to the disk level.
- More efficient failover. Dividing the work load makes for more efficient failover. Unlike other failover strategies, small workload components can be failed over individually, rather than having to fail over an entire server’s work load. System resources that are unaffected by a failure avoid being migrated unnecessarily, improving both availability and performance.

The underlying design of DIGITAL Clusters enables scaleability. In contrast, both the mirrored backup and the high degree of synchronization of fault-tolerant systems undermine scaleability.

### Industry Standards and Commodity Hardware

DIGITAL Clusters for Windows NT provides investment protection. Because it is designed for industry-standard hardware, software, interconnects, and protocols, it offers numerous advantages over other solutions. By supporting both Intel and Alpha processor-based architectures, enterprises are able to leverage their current hardware investments and feel secure in their ongoing choices.

DIGITAL Clusters supports a variety of off-the-shelf hardware components such as RAID subsystems, SCSI-2 disks, and SCSI adapters. These off-the-shelf choices reduce cost in the short term and add to investment protection in the long term.

## How DIGITAL Clusters Fits in a PC LAN Environment

Similarly, supported clients for automatic failover include the most widely used desktops running Windows NT, Windows for Workgroups, and Windows 95®. Other clients, such as Macintosh®, MS-DOS®, and OS/2, also can be used with manual reconnects.

Another way DIGITAL Clusters furnishes investment protection is through support of industry-standard networking protocols between the clustered servers and clients. DIGITAL Clusters supports all the Windows NT supported network protocols: TCP/IP, NetBEUI, and IPX/SPX.

DIGITAL Clusters is fully compatible with Windows NT Server as follows:

<b>DIGITAL Clusters for Windows NT</b>	<b>Microsoft Windows NT Server</b>
Version 1.1	Version 4.0 with NT SP2
Version 1.0 with Clusters SP2	Version 3.51 with NT SP5

Unlike other solutions, DIGITAL Clusters is not a port of older technology to a new operating system; it was designed from the ground up for the Windows NT client/server environment. An *SNMP (Simple Network Management Protocol)* agent is included enabling industry-standard SNMP management tools, such as ServerWORKS™, to work with DIGITAL Clusters for Windows NT.

## How DIGITAL Clusters Fits in a PC LAN Environment

In a typical client/server LAN environment, a single-server system provides file, print, and application services to a group of desktop clients. In a cluster client/server configuration, the notion of a single server serving clients is extended to include multiple server systems. The collection of servers, or cluster, is viewed by clients as a single server. This is accomplished via cluster software, which performs the management, integration, and synchronization of the servers in the cluster, or *cluster members*. Work assigned to the cluster is partitioned across the two servers with, for example, file services furnished by one server, and database services by the other.

In a clustered environment, end-user clients have access to the combined resources of the entire cluster. Like the single-server environment, a cluster offers a single-management environment. Clients view resources and services in the cluster as if they were local. A major advantage of clustering in a LAN environment is the ability to add system components incrementally to build in component redundancy for higher availability.

As customers deploy client/server solutions in their enterprise, they are concerned about system reliability and a cost-effective growth path for the future—attributes that are critical to supporting their user community and running their business. DIGITAL's cluster technology on Windows NT is well suited to address these concerns by enhancing the availability, scalability, and management of data and key services within a client/server LAN environment.

## How DIGITAL Clusters Fits in a PC LAN Environment

### Supported Server Architectures

DIGITAL Clusters supports both Intel and Alpha processor-based Windows NT Server systems as cluster members. Any single cluster must have the same processor architecture: either both Intel processors or both Alpha processors. The two cluster members do not need to be identically configured, and one or both may be SMP systems. Each clustered server must have its own local system disk that may be used to store data or run applications.

### Supported Clients

DIGITAL Clusters for Windows NT is a LAN server-based cluster solution. End-user clients are not members of the cluster. All Windows clients—Windows NT, Windows 95, and Windows for Workgroups—with LAN connections to the clustered servers are fully supported. With the cluster client software installed, these clients can access the cluster as a single system via the cluster alias. They do not need to know the individual names of the servers to which they are connected. The cluster name service software directs the clients to the correct disk or file share.

Other clients—such as Macintosh, MS-DOS, and OS/2—can access the cluster and benefit from the cluster resources and services as well. However, these other clients must know the names of the clustered servers, and must manually reconnect in the event of a failover. Manual reconnection is also required for clients that have wide area network (WAN) connections from the clustered servers.

If a client only needs to use applications accessed through an IP address, you do not need to install the client software. For example, if the client only uses the Microsoft Internet Information Server (IIS), or only connects to Microsoft SQL Server through an IP address, you can use the cluster software's IP address failover capability. IP address failover supports *all* clients with TCP/IP connections to the clustered servers and functions in a WAN.

### Physical Connections

The two clustered servers are connected by up to three physical connections:

- Shared SCSI bus (or buses). All clustered data must reside on the disks or subsystem RAID storage connected to the shared SCSI bus. The number of shared SCSI buses that a cluster can support is limited only by the number of available slots in the Windows NT Server systems that are cluster members. One or more shared SCSI buses are a required component of DIGITAL Clusters for Windows NT.
- Enterprise network connection. This is the primary network that connects end-user clients to the clustered Windows NT Servers. Any Windows NT supported LAN, such as Ethernet or FDDI, can be used.
- Private network connection. This secondary network connection between the two clustered servers is highly recommended to eliminate a single point of communications

failure in the cluster, and to ensure that the two servers will be able to communicate with each other in the event of an outage of the enterprise LAN. The second network also allows manual failover of the cluster in the case of a network partition. Low-cost Ethernet, such as thinwire or twisted pair, is more than sufficient to accommodate the minimal message-passing traffic of the cluster software.

### Management of Resources and Services

DIGITAL Clusters allows workload partitioning to the disk level by assigning disks on the shared SCSI bus to one server at a time for management and control. Moving disk resources from one server to another is simply a matter of reconfiguration via the graphical cluster administration tool, Cluster Administrator. (See Chapters 2, 6, and 9 for details.) There is no need to turn off the power or recable hardware components.

DIGITAL Clusters supports the automatic detection of new storage devices on the SCSI bus. If you add an additional disk to your storage tower, or turn on the power on a storage device that previously was turned off, the cluster software can detect this and can automatically incorporate the new device into the cluster's resource list. To initiate this action, you run Disk Administrator after adding the new storage device. Then, using Cluster Administrator, you add the device to a failover group to bring it online. (See Chapters 2, 6, and 9 for details.)

DIGITAL Clusters supports the use of standard Windows NT management tools, such as Windows NT Explorer and the `net use` command, to manage cluster resources. The intuitive Cluster Administrator allows easy configuration and reconfiguration of the cluster. The cluster state and configuration are remotely accessible by any SNMP browser, through the standard SNMP agent and cluster *MIB* (*management information base*) included with the cluster software.

---

# Understanding Cluster Concepts

This chapter discusses basic cluster concepts and terminology. It also presents a database failover example.

## Cluster Concepts and Terminology

This section discusses concepts that you must understand to configure your cluster for failover.

### Failover

To ensure the highest level of system availability while maintaining data and system integrity, a cluster must be able to provide services and access to resources in the event of multiple failures. Such failures include software, server, storage, and LAN-related failures.

High availability is achieved in a cluster by using active backup subsystems. These backup subsystems perform routine functions and are themselves primary servers for a given set of cluster services and resources. In the case of a failure, the service or resource is relocated to an alternate path, the other cluster server. This transparent relocation of cluster services and resources is referred to as *failover*.

### Spare Capacity

A given cluster member must have sufficient processing power and memory to handle a worst-case scenario of the other cluster server being nonfunctional. Although the DIGITAL Clusters software allows for resource balancing between both cluster servers, when one server fails, the other server must have sufficient resources both to continue to handle its existing services, and to handle the services of the other server. It is reasonable to expect some increase in service delay during this period, but services must be able to run.

## Cluster Concepts and Terminology

### Failover Objects

A *failover object* is any cluster service or resource for which you want to ensure availability in the event of a system failure. Examples of failover objects include disks, groups of databases served by applications such as Microsoft SQL Server and Oracle7 Workgroup Server, and any application that can be launched and shut down in a script. Failover objects can be collected into *failover groups*, discussed in the next section.

### Failover Groups

To flexibly define and manage failover objects, the system administrator can create logical groups of cluster services and resources that are referred to as *failover groups*. Resources in a failover group will move together; either they all are on line or they all are off line. Resources in different failover groups move independently. Resources in a failover group are ordered such that more primitive services are started first when going on line, and the order is reversed when going off line.

Failover groups are defined easily using Cluster Administrator. (See Chapter 9 for details.) They often consist of a collection of applications and storage devices that are used together. For example, a database exported by an application such as Microsoft SQL Server and the shared cluster disk or disks on which the database is stored may be specified in the same failover group. A failover group also can contain one or more disks without an associated application, as in the case of NTFS file service failover. When using the cluster IP address failover feature, you must place an IP failover object in each group. Each group represents an application, and must contain all the resources the group needs to run the application. If the cluster software detects a software or system failure, the entire failover group will be migrated to the alternate cluster server.

### Failover Policy

*Failover policy* is the plan of action the cluster software follows for a failover group. Each failover group is associated with a failover policy that is defined using Cluster Administrator. The ability to define failover policy for individual failover groups gives the system administrator more flexibility in workload balancing. Using Cluster Administrator, the system administrator can:

- Define which cluster server will be the primary server for each failover group.
- Disable and enable failover for individual failover groups. For example, the system administrator may decide to enable failover for mission-critical applications only.
- Define whether the failover group will automatically migrate to the primary server when the primary server returns to service in the cluster.

The system administrator may determine that the cluster workload must be redistributed. For example, a database application must be moved from one cluster server to the other because of changes in client service requirements. Workload redistribution is easily

accomplished through software reconfiguration of failover group policy for the cluster. Using Cluster Administrator, the system administrator can override the failover group policy by:

- Manually migrating the group to the other server.
- Disabling failover, forcing the group to stay on the current cluster server. This could be used to keep the group on the alternate server while the primary server is being brought up and down for maintenance and testing.
- Disabling the group, which would take the group off line temporarily while keeping its configuration unchanged in the cluster configuration database.

### Failback

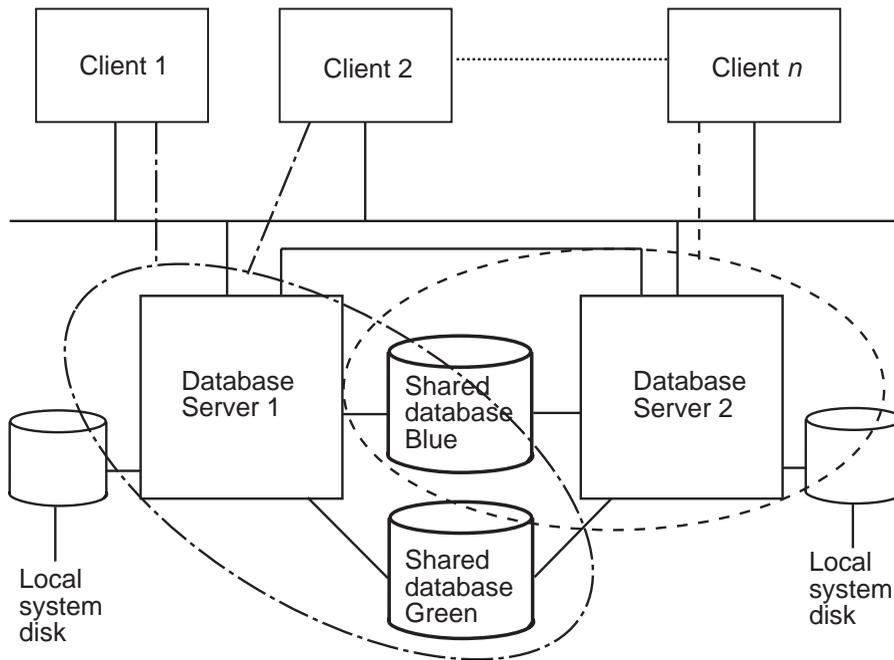
*Failback* is the action of a resource transitioning from its secondary location back to its primary location. It describes what can happen when the cluster server causing the failover returns to an operational status. Failback is also a policy attribute to control whether failback happens automatically. In Cluster Administrator, this policy attribute is called “Return to the Primary Server” (see the section Working with a Failover Group in Chapter 9). Failback can be initiated either through automatic decision, as enabled in Cluster Administrator through the policy attribute, or through a manual transfer request. There are two ways to initiate a manual transfer: using Manage Manual Failover in Cluster Administrator (see Chapter 8), and through the command line utility CLUXFER (see Appendix A). Failback is important to restore the cluster to full operation and restore the static workload partitioning of the cluster.

### Example of Database Failover

Let’s look at how database failover works in a cluster. The following figure shows a DIGITAL Clusters for Windows NT configuration. It includes two server systems sharing two buses connected to two storage boxes containing disk devices. There is a second dedicated Ethernet connection between the two servers that furnishes redundant interserver communication in the event of a primary LAN failure. Each server is running a database application—one could be running Microsoft SQL Server, the other Oracle7 Workgroup Server. The databases are located on the shared disks. One server at a time can access a database on a given disk.

## Example of Database Failover

### Database Failover Example—Normal

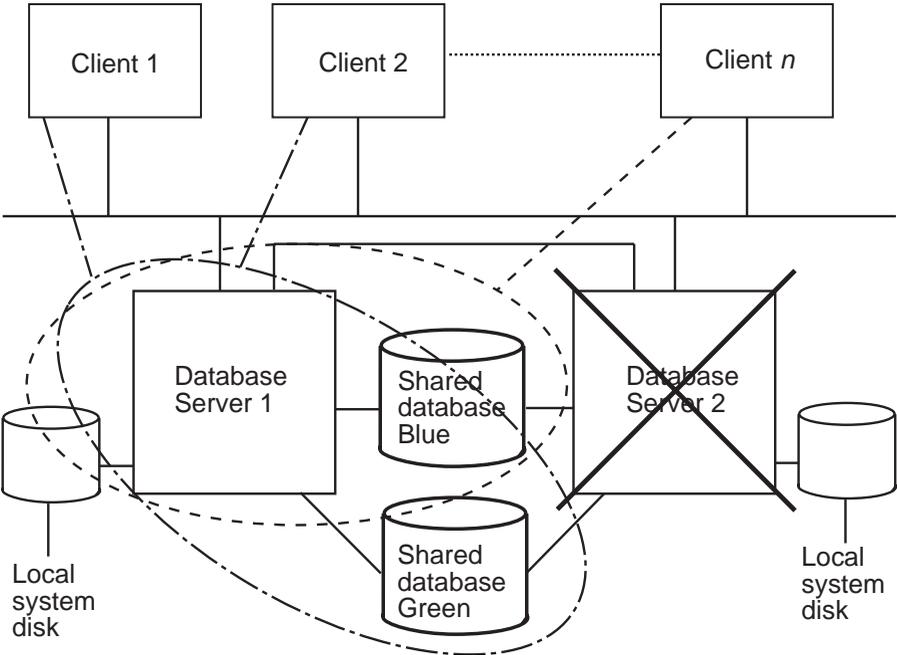


ZK-8759A-FH5

Several end-user clients are accessing the Blue database via Server 1, while other end-user clients are accessing the Green database via Server 2. In this example, we will assume that all the clients have connected to the cluster using the cluster alias. The clients do not need to know which server a database is being served by because the cluster software automatically routes them to the correct server. If one of the cluster servers should fail, the cluster software has been configured to provide the same services and resources to the clients using a backup path, the alternate cluster server.

Suppose Server 2 fails. Server 1 will assume support for its work as shown in the next figure. The cluster software will migrate the disk on which the Green database resides, start up the database server application, and redirect the client end users.

**Database Failover Example—Failure**



ZK-8760A-FH5

Note that it is necessary for the database application to be installed on both cluster servers for failover to occur. See Chapter 3 for instructions on database software installation and configuration requirements for failover. See Chapter 10 for a discussion of application considerations.

---

# Configuring Database Software for Failover

This chapter presents Microsoft SQL Server and Oracle7 Server installation and configuration requirements in the DIGITAL Clusters for Windows NT environment; recommendations; restrictions; and verification procedures.

## Microsoft SQL Server Installation and Configuration Requirements

By combining the failover features of DIGITAL Clusters for Windows NT with those of Microsoft SQL Server for Windows NT, you can have a high availability database solution. This section discusses the necessary steps to ensure high availability of SQL Server databases in the DIGITAL Clusters for Windows NT environment.

### New Features and Enhancements in DIGITAL Clusters 1.1

The following new features and enhancements to SQL Server since DIGITAL Clusters Version 1.0 are included in DIGITAL Clusters Version 1.1 and Version 1.0 with Service Pack 2:

- Support for multiple, independent SQL Server databases. You can now use multiple SQL Server databases independent of one another. You can enroll and unenroll databases in failover groups as desired. SQL Server clients now have the option of establishing connections to databases using either TCP/IP sockets or named pipes. Clients can connect to a particular group of SQL Server databases.
- Ability to run SQL Server simultaneously on both cluster servers. You can now run SQL Server on both cluster servers, each supporting a different group of SQL Server databases. If one server fails, the other server can assume support for all databases.
- Automated configuration of SQL Server databases. Cluster Administrator now provides a more automated setup sequence for SQL Server databases, eliminating manual invocations of stored procedures.

## Microsoft SQL Server Installation and Configuration Requirements

- IP address failover support. DIGITAL Clusters 1.1 provides support for failing over SQL Server clients that use the TCP/IP network protocol. A cluster IP address can be associated with the SQL Server software. If SQL Server fails over from one cluster server to the other, the associated cluster IP address will migrate with the application. SQL Server clients continue to use the same cluster IP address. The address is transparently routed to the alternate cluster server.

### Prerequisite Information

This section gives important information that you should review before using Cluster Administrator to configure for SQL Server database failover.

### Software Requirements

Software requirements are as follows:

- Acquire one or more Microsoft SQL Server licenses for each cluster server in accordance with the Microsoft SQL Server licensing requirements.
- Install Microsoft SQL Server Version 6.5 with Service Pack 1 on a local disk on each cluster server.

### Designating Primary and Failover Servers

In Microsoft SQL Server terms, the SQL Server product assumes a static definition of both a *primary server* and a *fallback server*. In contrast, the DIGITAL Clusters for Windows NT software defines a *primary server* and a *failover server* when you use Cluster Administrator to add a failover group. The cluster software uses these definitions only for the purpose of failback, in which failover group control is returned to the primary server when the primary server returns to operational status. See the section Failback in Chapter 2 for details.

To simplify the discussion, we will use the term *failover server* in references to both the Microsoft SQL Server product and the DIGITAL Clusters for Windows NT product.

### SQL Server Database Failover

*Microsoft SQL Server database failover* refers to a database failing over, not the entire application. Upon primary server failure, the failover server will service the databases on the shared disks.

To ensure failover, the servers and databases must be properly configured by invoking stored procedures supplied with the Microsoft SQL Server product. In DIGITAL Clusters 1.1, the cluster software automatically invokes these procedures when you use the Manage SQL Server Databases dialog box of Cluster Administrator and choose the Enroll button. See the section Configuring SQL Server for Failover for details.

# Microsoft SQL Server Installation and Configuration Requirements

## Access to Shared Disks

Access to shared disks is not required when installing Microsoft SQL Server. However, it is required when adding new SQL Server databases that you would like to designate as highly available. You can create the new databases either before or after installing the cluster server software.

## Configuration Requirements

Verify the following to ensure that your cluster servers and clients are configured properly to run Microsoft SQL Server with DIGITAL Clusters for Windows NT:

- Both cluster servers must reside in the same Windows NT domain.
- Neither cluster server can be a member of any other cluster.
- The cluster servers must be configured with the same network transports to enable interserver communication. For example, if you use the TCP/IP network protocol, the servers must be in the same IP subnet and NetBIOS must be installed on each.
- Cluster clients connecting to SQL Server using named pipes must be configured with the same network transports as the server systems. For example, if the servers communicate with clients using named pipes over TCP/IP, the clients and servers must be in the same IP subnet and NetBIOS must be installed on each. In comparison, clients using the NetBEUI protocol can communicate with cluster servers on the same LAN.

Cluster clients connecting to SQL Server using TCP/IP sockets do not have the same subnet restriction as named pipes connections over TCP/IP.

- Passwords must match between the SQL Server administrator (sa) account and the SQL Server failover object. Mismatched passwords can cause the DIGITAL Clusters Failover Manager to freeze.

## Configuration and Run-Time Recommendations

Before you configure Microsoft SQL Server to run with DIGITAL Clusters, review the following recommendations:

- Do not install the SQL Server software on a Windows NT Primary Domain Controller (PDC) or Backup Domain Controller (BDC). This recommendation is stated in the Microsoft SQL Server product documentation as well.
- DIGITAL strongly recommends that you install the SQL Server software on a local disk on each server system *before* installing the DIGITAL Clusters for Windows NT software.
- Before initiating a manual failover of an SQL Server database, DIGITAL strongly recommends that you close all active client connections to the database. The cluster software will close any remaining connections when the database is failed over.

# Microsoft SQL Server Installation and Configuration Requirements

## Restrictions

The following Microsoft SQL Server restrictions have been identified in DIGITAL Clusters 1.1:

- You cannot install SQL Server while the Cluster Failover Manager and Cluster Failover Management Database Server (CFMD Server) services are running.

If you have chosen to install the SQL Server software after installing the cluster software (DIGITAL strongly recommends that you install the SQL Server software *before* the cluster software), verify that the Cluster Failover Manager and CFMD Server services are not running by using the Services applet on the Windows NT Control Panel.

- When creating SQL Server databases on shared cluster disks, you must give each database in the cluster a unique name. There cannot be two databases with the same name on shared disks served by different cluster servers.

## Configuring SQL Server for Failover

This section presents the steps you must perform to configure Microsoft SQL Server for failover. Two step sequences are presented depending on when you choose to create the SQL Server databases on the shared disks. DIGITAL strongly recommends using the first step sequence because you are less likely to encounter potential problems.

In the first sequence, the steps are as follows:

1. Install the SQL Server software on a local disk on each cluster server.
2. Install the DIGITAL Clusters software on each server system.
3. Create a failover group for one or more SQL Server databases.
4. Create the SQL Server databases on shared disks.
5. Enroll the SQL Server databases for high availability.

In the second sequence, the steps are as follows:

1. Install the SQL Server software on a local disk on each cluster server.
2. Create the SQL Server databases on shared disks.
3. Install the DIGITAL Clusters software on each server system
4. Create a failover group for one or more SQL Server databases.
5. Enroll the SQL Server databases for high availability.

► **To configure the SQL Server software for failover:**

1. Install Microsoft SQL Server on a local disk on each cluster server.
2. If you choose to create your shared SQL Server databases before installing the cluster software, you must follow these steps:
  - a. Shut down and turn off the power to the failover server. This will ensure that only the primary server will be allowed access to the shared disks.

---

**Caution**

---

There is danger of disk corruption if both server systems are turned on with the shared bus connected when the cluster software is not installed.

---

- b. On the primary server, use a database administrator tool to create the SQL Server databases (and optionally, the transaction logs) on the shared disks.

---

**Caution**

---

A database may not be split between disks that will be placed in different failover groups.

---

3. Install the cluster software on each server system. See the *Configuration and Installation Guide* for instructions.
4. On the primary server, use Cluster Administrator to create a failover group for one or more SQL Server databases using the following guidelines:
  - a. Each failover group must contain *all* the shared disks on which a given SQL Server database resides.
  - b. You may have more than one database in a given failover group.
  - c. A shared disk may be specified only in one failover group.

---

**Caution**

---

A database may not be split between disks that are in different failover groups.

---

## Microsoft SQL Server Installation and Configuration Requirements

After creating a failover group, the shared disks will be placed on line on the primary server for the group. See the section *Creating a Failover Group* in Chapter 9 for instructions.

5. Repeat step 4 for additional SQL Server databases that resides on shared disks.
6. If you have not already done so, use a database administrator tool on the primary server to create the SQL Server databases (and optionally, the transaction logs) on the shared disks.

---

### Caution

---

A database may not be split between disks that are in different failover groups.

---

7. Configure an SQL Server database for high availability:
  - a. On each cluster server, use the Services applet of the Windows NT Control Panel to check if the MSSQLServer service is running. If not, start the service.
  - b. Using Cluster Administrator, choose SQL Server Databases from the Manage menu.

The Manage SQL Server Databases dialog box is displayed. It lists all the databases eligible for failover support. The list corresponds to the databases that reside on shared cluster disks that are currently on line on the primary server. Databases that already have been configured for high availability are followed by “Yes”, whereas databases that have not been configured for high availability yet are followed by “No”.

- c. Select the database that you want to configure for high availability.
  - d. Choose the Enroll button. This operation configures the SQL Server software on both cluster servers and creates the SQL Server failover object needed to enable database failover.

See the section *Managing SQL Server Databases* in Chapter 8 for further details.

8. Repeat step 7 for each additional SQL Server database that you want to configure for high availability.

## Configuring Clients to Access SQL Server Databases

You must establish your SQL Server client/server connections over one of the following supported network transports: named pipes or IP. The next procedure outlines how to do this on a Windows NT, Windows 95, or Windows for Workgroups client.

## Microsoft SQL Server Installation and Configuration Requirements

### ► To configure SQL Server connections on a Windows NT, Windows 95, or Windows for Workgroups client:

1. Start the SQL Server Client Configuration Utility as described in the section Starting the Client Configuration Utility of the *Microsoft SQL Server Administrator's Companion*.
2. Choose the Advanced option as described in the section Setting Up Server Connections of the *Microsoft SQL Server Administrator's Companion*:
  - *For Windows NT and Windows 95 Clients:*  
Choose the Advanced tab. The SQL Server Client Configuration Utility dialog box is displayed.
  - *For Windows for Workgroup Clients:*  
Choose the Advanced button. The Advanced Client Options dialog box is displayed.
3. Fill in the boxes as follows:

---

#### Note

---

A client cannot simultaneously establish connections to an SQL Server using named pipes and IP.

---

*For client connections using named pipes:*

- a. In the Server box, create an SQL Server database alias.
- b. In the DLL Name box, type the dynamic link library (DLL) name for the named pipes Net-Library. Or, type or choose from the drop-down list Named Pipes.
- c. In the Connection String box, use this syntax:

```
\\ClusterName\pipe\ObjectName\query
```

where the variables have the following meanings:

*ClusterName* Specifies the name of your cluster.

*ObjectName* Specifies the name of the SQL Server failover object that the cluster software created automatically for you. This name is the *same* as the name of the failover group you created (in step 4 of the procedure "To configure the SQL Server software for failover") that contains the shared disks on which the SQL Server database resides.

## Microsoft SQL Server Installation and Configuration Requirements

For example, suppose your cluster name is `sqlcluster`; you have created a failover group for your SQL Server database named `sqldbl`; and SQL Server is using the default named pipe name `\\.\pipe\sql\query`. Your connection string would be:

```
\\sqlcluster\pipe\sqldbl\query
```

*For client connections using IP:*

- a. In the Server box, create an alias for the server that is currently controlling the SQL Server database to which the client wants to connect. The Server name should be unique amongst all servers to which the client can connect.
- b. In the DLL Name box, type the following:
  - *For Windows NT and Windows 95 clients:*  
Type the DLL name for the TCP/IP Net-Library. Or, type or choose from the drop-down list TCP/IP Sockets.
  - *For Windows for Workgroups clients:*  
Type the DLL name `dbmssoc3`.
- c. In the Connection String box, enter either the IP address that you used when creating the IP failover object for the SQL Server database, or the name that you associated with the IP address in the `etc/hosts` file. See the section Creating an IP Failover Object in Chapter 9 for details.

---

### Note

---

The IP address contained in the SQL Server failover object must be in the same subnet as the server IP address.

---

4. Add the connection information to the appropriate configuration file on the client.
  - *For Windows NT and Windows 95 Clients:*  
Choose the Add/Modify button.
  - *For Windows for Workgroups Clients:*  
Choose the Add/Change button, and then the OK button.
5. Choose the Done button.

## Verifying SQL Server Failover

Use the next procedure to verify that the SQL Server software is failing over properly in the DIGITAL Clusters for Windows NT environment.

► **To verify that the SQL Server software is failing over properly:**

*On the Primary Server:*

1. Use the Windows NT Registry editor, `regedt32.exe`, to open and examine the Registry. Be sure to place the editor in read-only mode by enabling Read Only Mode on the Options menu. For details on using the Registry editor, either select Help from within the program or refer to the documentation packaged with your Windows NT operating system:

- a. Locate the following keys:

*For DIGITAL Clusters Version 1.0 with Service Pack 1:*

HKEY\_LOCAL\_MACHINE/System/CurrentControlSet/Services/Cfmd/Database/Pipes/SQL

*For DIGITAL Clusters Version 1.0 with Service Pack 2 and Version 1.1:*

HKEY\_LOCAL\_MACHINE/System/CurrentControlSet/Services/Cfmd/Database/Pipes/SQL-failover-object

Note that in Version 1.0 with Service Pack 2 and Version 1.1, there is one or more keys (in addition to HKEY\_LOCAL\_MACHINE/SYSTEM/CurrentControlSet/Services/Cfmd/Database/Pipes/SQL, retained for backward compatibility with Service Pack 1) with names corresponding to all the SQL Server failover objects created by enrolling databases for failover support.

- b. Expand the keys and locate the `ConnectionPoint` parameter.
  - c. Verify that the value of the `ConnectionPoint` parameter is the name of the primary server.
2. Shut down the primary server.

## Microsoft SQL Server Installation and Configuration Requirements

### *On the Failover Server:*

3. Repeat step 1, verifying that the value of the `ConnectionPoint` parameter is now the name of the failover server.

---

#### **Note**

---

It may take up to 2 minutes for the Registry on the failover server to be updated. Be sure to enable Auto Refresh on the Options menu.

---

### *On a Cluster Client:*

4. For client connections using named pipes, use the SQL Client Configuration Utility to configure an SQL Server alias. See step 3 of the section *Configuring Clients to Access SQL Server Databases* for details.
5. Using either the ISQL/w application or the Enterprise Manager Query Tool, use the SQL Server alias created in step 4, not the server name, to query an SQL Server database residing on a shared disk:
  - If the client cannot access the SQL Server database, verify that you have followed the instructions outlined earlier in this chapter.
  - If problems persist, contact your local DIGITAL Multivendor Customer Services sales specialist.

### *On the Primary Server:*

6. Bring the primary server on line again.
7. Repeat step 1, verifying that the value of the `ConnectionPoint` parameter has returned to the name of the primary server.

---

#### **Note**

---

It may take up to 2 minutes for the Registry on the primary server to be updated. Be sure to enable Auto Refresh on the Options menu.

---

### *On a Cluster Client:*

8. Repeat step 4.

## Adding a Shared Disk to a Failover Group for SQL Server

If you want to modify a failover group for an SQL Server failover object to include another shared disk, you must follow the instructions in the next procedure.

► **To modify a failover group for an SQL Server failover object to include another shared disk:**

1. Using Cluster Administrator, add the new shared disk to the failover group. Be sure to arrange the objects in the group so that all the disks are placed *before* the SQL Server failover object. See the section Modifying a Failover Group in Chapter 9 for instructions.
2. Use a database administrator tool to create additional SQL Server databases (and optionally, the transaction log) on the shared disk.
3. Configure the new databases for high availability:
  - a. On each cluster server, use the Services applet of the Windows NT Control Panel to check if the MSSQLServer service is running. If not, start the service.
  - b. Using Cluster Administrator, choose SQL Server Databases from the Manage menu.

The Manage SQL Server Databases dialog box is displayed. It lists all the databases eligible for failover support. The list corresponds to the databases that reside on shared cluster disks that are currently on line on the primary server. Databases that already have been configured for high availability are followed by “Yes”, whereas databases that have not been configured for high availability yet are followed by “No”.

- c. Select the database that you want to configure for high availability.
- d. Choose the Enroll button. This operation configures the SQL Server software on both cluster servers and creates the SQL Server failover object needed to enable database failover.

See the Section Managing SQL Server Databases in Chapter 8 for details.

## Unenrolling an SQL Server Database from Failover Support

There may be times when you will want to unenroll an SQL Server database from cluster failover support. The next procedure outlines the steps you need to perform to do so.

► **To unenroll an SQL Server database from cluster failover support:**

1. On each cluster server, use the Services applet of the Windows NT Control Panel to check if the MSSQLServer service is running. If not, start the service.

## Microsoft SQL Server Installation and Configuration Requirements

2. Using Cluster Administrator, choose SQL Server Databases from the Manage menu.

The Manage SQL Server Databases dialog box is displayed. It lists all the databases eligible for failover support. The list corresponds to the databases that reside on shared cluster disks that are currently on line on the primary server. Databases that already have been configured for high availability are followed by “Yes”, whereas databases that have not been configured for high availability yet are followed by “No”.

3. Select a database that you want to unenroll.
4. Choose the Unenroll button.

Upon completing this procedure, the SQL Server database will have returned to the same state it was in just prior to enrolling it for cluster failover support. Note that unenrolling an SQL Server database does not affect primary server access.

See the section Managing SQL Server Databases in Chapter 8 for further details.

### Expanding or Shrinking an SQL Server Database

To change the size of a database that has been enrolled for cluster failover support, follow the instructions outlined in the next procedure.

#### ► To expand or shrink an SQL Server database:

1. Execute the procedure in the previous section Unenrolling an SQL Server Database from Failover Support.
2. If you will be expanding the database onto an additional disk, check if the failover group containing the SQL Server failover object also contains this disk. If not, use Cluster Administrator to add the new disk to the failover group. See the section Modifying a Failover Group in Chapter 9 for instructions.
3. Use a database administrator tool to expand or shrink the SQL Server database.
4. Reenroll the database for failover support:
  - a. On each cluster server, use the Services applet of the Windows NT Control Panel to verify that the MSSQLServer service is running.
  - b. Using Cluster Administrator, choose SQL Server Databases from the Manage menu.

The Manage SQL Server Databases dialog box is displayed. It lists all the databases eligible for failover support. The list corresponds to the databases that reside on shared cluster disks that are currently on line on the primary server. Databases that already have been configured for high availability are followed by “Yes”, whereas databases that have not been configured for high availability yet are followed by “No”.

- c. Select the database that you want to configure for high availability.
- d. Choose the Enroll button. This operation configures the SQL Server software on both cluster servers and creates the SQL Server failover object needed to enable database failover.

See the section Managing SQL Server Databases in Chapter 8 for further details.

### Moving an SQL Server Database and Disk to a Different Failover Group

If you want to move an SQL Server database and the disk on which it resides to a different failover group, both groups must be on line on the *same* server at the time you move the database.

► **To move an SQL Server database and its disk to a different failover group:**

1. Verify on which disk the database resides.
2. Deactivate the database. See the section Managing SQL Server Databases in Chapter 8 for instructions.
3. Move the disk on which the database resides out of the current failover group. See the section Modifying a Failover Group in Chapter 9 for instructions.
4. Move the disk into the new failover group, placing it *before* the SQL Server failover object.
5. Enroll the database. See the section Managing SQL Server Databases in Chapter 8 for instructions.

### Moving SQL Server Databases to a New Cluster

If you have existing SQL Server databases that you want to move to a new cluster, for example, if you are upgrading from DIGITAL Clusters 1.0 with Clusters SP1 to DIGITAL Clusters 1.1, use the following procedure.

---

**Note**

---

The databases must be placed on a shared disk in the new cluster.

---

► **To move SQL Server databases to a new cluster:**

1. Ensure that there are no client connections to the SQL Server databases.
2. For *each* database, invoke the following stored procedure:  

```
sp_dboption DatabaseName, offline, true
```

## Oracle7 Server Installation and Configuration Requirements

3. Install the DIGITAL Clusters software. See the *Configuration and Installation Guide* for instructions.
4. Enroll *each* database for high availability following steps 7 and 8 of the procedure “To configure the SQL Server software for failover.”
5. For *each* database, invoke the following stored procedure:

```
sp_dboption DatabaseName, offline, false
```

### Resetting the Suspect Status of an Unenrolled SQL Server Database

If an unenrolled SQL Server database on a cluster share becomes suspect after a failover, use the next procedure to reset the database status.

► **To reset the suspect status of an unenrolled SQL Server database on a cluster share after a failover:**

1. Enroll the database in a failover group. See the section Managing SQL Server Databases in Chapter 8 for instructions.
2. Using the SQL Enterprise Manager Query Tool, run the following stored procedures. These procedures are listed in the master database:

- a. `sp_fallback_dec_clean`
- b. `sp_fallback_activate_svr_db PrimaryServerName, DatabaseName`

where the parameters have the following meanings:

*PrimaryServerName* Specifies the name of the server where the database's disks are on line.

*DatabaseName* Specifies the name of the database.

3. Unenroll the database from the failover group. See the section Managing SQL Server Databases in Chapter 8 for instructions.

## Oracle7 Server Installation and Configuration Requirements

Oracle failover support in the DIGITAL Clusters for Windows NT product offers a high availability database solution for Oracle7 Workgroup Servers and Oracle7 Enterprise Servers. With proper configuration, when a primary server running Oracle7 Server fails, the server instance will fail over to the failover server. This section outlines the necessary steps for configuring a highly available Oracle7 Server in the DIGITAL environment.

## New Features in DIGITAL Clusters 1.1

DIGITAL Clusters 1.1 and 1.0 with SP2 include the following new features and enhancements to Oracle7 Server since DIGITAL Clusters 1.0:

- Failover support for Oracle7 Server Version 7.3.
- IP address failover support. DIGITAL Clusters 1.1 provides support for failing over Oracle7 Server clients that use the TCP/IP network protocol. A cluster IP address can be associated with an Oracle7 Server instance. If the instance fails over from one cluster server to the other, the associated cluster IP address will migrate with the application. Oracle7 Server clients continue to use the same cluster IP address. The address is transparently routed to the alternate cluster server.

## Prerequisite Information

This section gives important information that you should review before using Cluster Administrator to configure for Oracle7 Server database failover.

### Software Requirements

Software requirements are as follows:

- Acquire one or more Oracle7 Workgroup Server or Oracle7 Enterprise Server licenses for each cluster server in accordance with the Oracle7 Server licensing requirements.
- Install Oracle7 Server Version 7.1, 7.2, or 7.3 on a local disk on each cluster server.

---

#### Note

---

Do not mix Oracle7 Server versions between cluster servers.

---

- If you want to use the DIGITAL Clusters IP failover feature with Oracle7 Server, you must use Oracle7 Server Version 7.3 with DIGITAL Clusters 1.1. IP failover requires Microsoft Windows NT 4.0, which is not supported by earlier versions of Oracle7 Server and DIGITAL Clusters.

### Oracle Server Database Failover

Throughout the DIGITAL Clusters for Windows NT documentation, *Oracle Server database failover* refers to an Oracle7 Server instance failing over. Oracle7 Server allows multiple instances to run on a single cluster server. Operations of instances that have failed over will not affect operations of other instances.

The database associated with a failover instance can be accessed only by one server system at a time. Should the primary server become unavailable, the failover server is ready to service the instance database on the shared disks, thus providing high availability.

# Oracle7 Server Installation and Configuration Requirements

## Configuration Requirements

Verify the following to ensure failover of your Oracle instance:

- Both cluster servers must reside in the same Windows NT domain.
- Neither cluster server can be a member of any other cluster.
- The cluster servers must be configured with the same network transports to enable interserver communication. For example, if you use the TCP/IP network protocol, the servers must be in the same IP subnet and NetBIOS must be installed on each.
- The shared disks used for Oracle instance failover must be assigned identical fixed drive letters on each cluster server.
- All files associated with the Oracle instance, including the data files, control files, log files, and parameter file, must reside on one or more shared disks.
- When you create or modify a failover group for Oracle Server, you *must* place the objects in the following sequence:
  1. Disk object or objects
  2. IP failover object
  3. Oracle failover object

If you fail to do this, the Oracle network listener will not start.

## Creating and Configuring an Oracle Instance for Failover

This section presents the steps you must perform to create and configure an Oracle instance for failover. Following is an overview of the step sequence. For detailed instructions, see the next procedures “Before creating an Oracle instance for failover” and “To create an Oracle instance for failover.”

1. Install the DIGITAL Clusters software on each server system.
2. Create a failover group containing all shared disks to be used for Oracle instance failover.
3. On each cluster server, assign identical fixed drive letters for the shared disks.
4. *For IP socket connections:* Create an IP failover object. Then modify the failover group to include the IP failover object.
5. *On the Primary Server:* Create an Oracle instance for failover; create an associated database on one or more shared disks; create and configure the Oracle network listener description; start the network listener; then manually fail over the group to the failover server.

## Oracle7 Server Installation and Configuration Requirements

6. *On the Failover Server:* Create an identical Oracle instance to the one you created on the primary server.
7. *On the Primary Server:* Create an Oracle failover object. Then modify the failover group to include the Oracle failover object.
8. Repeat the appropriate steps for *each* instance that you configure for failover.

### ► **Before creating an Oracle instance for failover:**

1. If you have not already done so, install the DIGITAL Clusters software on each server system.
2. Using Cluster Administrator, create a failover group containing all shared disks to be used for Oracle instance failover. Give the group a meaningful name. See the section Creating a Failover Group in Chapter 9 for instructions.

The shared disks will be placed on line on the primary server.

3. On each cluster server, assign identical fixed drive letters for the shared disks used for Oracle instance failover.

*On the Primary Server:*

- a. On the primary server for the shared disks used for Oracle instance failover, click Start and choose Programs→Administrative Tools→Disk Administrator.
- b. Select a shared disk for which you want to modify the drive-letter assignment.
- c. From the Tools menu, choose Drive Letter.

The Assign Drive Letter dialog box is displayed.

- d. Assign a fixed drive letter. To avoid conflict with network drive letters, DIGITAL strongly recommends that you select a drive letter at the end of the alphabet, such as X, Y, or Z.
- e. Repeat steps b to d for each shared disk used for Oracle instance failover.
- f. Reboot the server.
- g. Manually fail over the group to the failover server. See the section Managing Manual Failover in Chapter 8 for instructions.

*On the Failover Server:*

- h. Once failover has completed, repeat steps a to f on the failover server. You must assign the same fixed drive letters on both server systems.
- i. Manually fail back the group to the primary server. See the section Managing Manual Failover in Chapter 8 for instructions.

## Oracle7 Server Installation and Configuration Requirements

*On the Primary Server:*

- j. Using Disk Administrator, verify that the shared disks have failed over.
4. *For IP socket connections:*
    - a. Create an IP failover object, making note of the IP address that you specify for later use in creating the `listener.ora` file. See the section Creating an IP Failover Object in Chapter 9 for instructions.
    - b. Modify the failover group that you created in step 2 to include the IP failover object.

### ► **To create an Oracle instance for failover:**

Repeat this procedure for *each* instance that you configure for failover.

*On the Primary Server:*

1. Use an Oracle database administrator tool (for example, `ORADIMnn`, where *nn* corresponds to the version number [7.1, 7.2, or 7.3] of Oracle7 Server running on your cluster servers) to create an Oracle instance for failover. Note the Oracle System Identifier (SID) that you assign to the instance for later use in creating the `listener.ora` file.
2. Create an associated database on one or more shared disks. If you are working with an existing database, move the database data files to one or more shared disks. Also use a shared disk to store all files associated with the database, including the control file, log file, and parameter file.

Note that creating a database involves several steps, such as creating the parameter file, editing the associated parameters, creating the database, allocating rollback segments, and so forth. Consult the Oracle7 Server documentation for details.

3. Create and configure the network listener description.

*For Versions 7.2 and 7.3:* If you are running Oracle7 Server Version 7.2 or 7.3, create a separate network listener description for each instance using a supported Oracle7 Server tool such as Oracle Network Manager for Windows. See the Oracle Network Manager for Windows product documentation for details. Use the following guidelines when creating the listener (note that Oracle7 Server terminology is designated in quotes):

*For named pipe connections:*

- a. Create an address for the network listener by specifying Named Pipe as the “network protocol”.
- b. Specify a unique name for the named “pipe” associated with this address of the listener.

## Oracle7 Server Installation and Configuration Requirements

- c. Specify the cluster alias as the “server” name for this address.

*For IP socket connections:*

- a. Create an address for the network listener by specifying TCP/IP as the “network protocol.”
- b. Specify a unique port number for the “port” associated with this address.
- c. Specify either the IP address or the associated IP name (as translated by the name service you are using) of the IP failover object for the “host” associated with this address of the listener.

*For simultaneous use of named pipe and IP socket connections:*

- a. Create both a named pipe address and a TCP/IP address for the same network listener following the previous instructions in this step.

The end result of step 3 is a description of the network listener in the file `listener.ora`. This file contains descriptions of all the listeners used on the server system. There is only one `listener.ora` file per server system. See the next section, Example `listener.ora` File, for a sample file extract.

4. Using either the Oracle utility `LSNRCTL` or the Services icon on the Windows NT Control Panel, check if there is a network listener running. If so, stop and then restart it. Otherwise, start the network listener.

---

### Note

---

You must use the `LSNRCTL` utility to start a new network listener for the first time. The service name on the Control Panel will appear as the listener name prefixed by “OracleTNSListener”.

---

5. With the associated database mounted and open, verify that you can access it using an Oracle database administrator tool such as `SQLDBAnn` or `SVRMGRnn` (where *nn* designates the Oracle7 Server version number). Use the Oracle SID to access the database. If you cannot access the database:
  - a. In the parameter file, verify that you have modified all path specifications so that files associated with the instance are created on a shared disk.
  - b. Verify that all files associated with the instance, including the data file, control file, log file, and parameter file, are on a shared disk.
6. Shut down the instance using an Oracle database administrator tool. Be sure to leave services for the instance running.

## Oracle7 Server Installation and Configuration Requirements

7. Using either the Oracle utility LSNRCTL or the Services icon on the Windows NT Control Panel, stop the network listener.
8. Verify that services for the instance are set to restart automatically using either the Oracle Instance Manager or the Services icon on the Windows NT Control Panel. If the services are not set to “Startup Automatic”, change this.
9. Manually fail over the group to the failover server. See the section Managing Manual Failover in Chapter 8 for instructions.

### *On the Failover Server:*

10. Once the failover server can access the shared disks, repeat steps 1 and 3 to 9 to create an *identical* Oracle instance to the one you created on the primary server using the following guidelines:
  - When repeating step 1, use the same Oracle SID that you used on the primary server.
  - When repeating step 3, create an identical listener description to the one stored in the primary server’s `listener.ora` file. If you have not configured local databases on either cluster server, you can use the same `listener.ora` file by copying it from the primary server.
  - When repeating step 9, manually fail over the group to the primary server.

### *On the Primary Server:*

11. Create an Oracle failover object, supplying information about the instance that you created in steps 1 to 3. See the section Creating an Oracle Failover Object in Chapter 9 for instructions.
12. Modify the failover group that you created in step 2 of the procedure “Before creating an Oracle instance for failover” to include the Oracle failover object that you created in step 11. The Oracle failover object must be placed *last* in the failover group. See the section Modifying a Failover Group in Chapter 9 for instructions.

Modifying the failover group will start the Oracle Server instance and will make the database available to the Oracle clients, once configured.
13. Verify that the Oracle Server instance is now running on the primary server. If the instance is not running, verify that the information you supplied in step 11 to create the Oracle failover object is correct.

## Example listener.ora File

Before Oracle Server can receive connections from clients, a network listener must be active on the server system. The configuration file for the network listener is `listener.ora`.

Following is an extract from a properly configured `listener.ora` file. The extract contains an example description of a network listener named `TestList` configured for both named pipe and IP socket connections.

```
#####
# Filename.....: listener.ora
# Name.....: TestCluster.world
# Date.....: 27-JAN-97 10:47:58
#####
SQLNET.AUTHENTICATION_SERVICES = (NONE)

USE_PLUG_AND_PLAY_TestList = OFF
USE_CKPPFILE_TestList = OFF
TestList = _____ 1
  (ADDRESS_LIST =
    (ADDRESS=
      (PROTOCOL=IPC)
      (KEY= TestDB.world) _____ 2
    )
    (ADDRESS=
      (PROTOCOL=IPC)
      (KEY= TEST) _____ 3
    )
    (ADDRESS =
      (COMMUNITY = NMP.world)
      (PROTOCOL = NMP) _____ 4
      (Server = TestCluster) _____
      (Pipe = testpipe1) _____ 5
    )
    (ADDRESS =
      (COMMUNITY = TCP.world)
      (PROTOCOL = TCP)
      (Host = 16.151.8.202) _____ 6
      (Port = 1565) _____ 7
    )
  )
STARTUP_WAIT_TIME_TestList = 0
CONNECT_TIMEOUT_TestList = 10
```

## Oracle7 Server Installation and Configuration Requirements

```
TRACE_LEVEL_TestList = OFF
SID_LIST_TestList =
  (SID_LIST =
    (SID_DESC =
      (GLOBAL_DBNAME = TestDB.world)           _____ 8
      (SID_NAME = TEST)
      (PRESPAWN_MAX = 10)                     _____ 9
    )
  )
PASSWORDS_TestList = (oracle)
```

In the previous file extract:

- 1 Specifies the Oracle network listener name.
- 2 Specifies the Oracle global database name. This address is used when the Oracle client and server reside on the same system.
- 3 Specifies the Oracle SID. This name is used when the client and server reside on the same system.
- 4 Specifies the cluster alias.
- 5 Specifies a unique pipe name.
- 6 Specifies the address of the cluster IP failover object associated with this Oracle database. You can use an IP name instead provided the name can be translated to the associated IP address.
- 7 Specifies a unique port number.
- 8 Specifies the global database name.
- 9 Specifies the Oracle SID.

## Configuring Clients to Access Oracle Server Databases

You must establish your Oracle client/server connections over one of the following supported network transports: named pipes or IP. A prerequisite for clients using named pipe connections to Oracle Servers is that they must have installed the DIGITAL Clusters client software. Clients using IP connections to Oracle Servers do not need to install the DIGITAL Clusters client software.

A key goal in configuring clients is to provide the addresses of the Oracle Servers. This task must be done regardless of whether the Oracle Server software is running in a DIGITAL Clusters environment. In the previous procedure “To create an Oracle instance for failover”, you configured the servers to respond to the network addresses specified in the

## Oracle7 Server Installation and Configuration Requirements

`listener.ora` file. The network addresses also must be provided to the clients so that they can communicate with the servers. Typically, clients identify servers using an Oracle service name that gets translated into a network address.

There are several methods that a client can use to translate Oracle service names into network addresses, including:

1. Use a local client configuration file, `tnsnames.ora`.
2. Use an Oracle Names server if one is used on the network.
3. Use a name service such as Distributed Computing Environment Cell Directory Service (DCE CDS), Sun®'s Network Information Service (NIS), Novell®'s NetWare Directory Services (NDS), or Banyan®'s Street Talk™.
4. Use a combination of the methods just listed.

If you will be using a local client configuration file, Oracle7 Server supplies several utilities to assist in creating this, including the SQL\*Net Easy Configuration utility and the Oracle Network Manager for Windows. Note that the SQL\*Net Easy Configuration utility can be used only for simple configurations. See your Oracle7 Server documentation for details on configuring clients in the Oracle Server environment.

### Configuring the `tnsnames.ora` File

This section discusses proper configuration of the `tnsnames.ora` file. If you use other name services, be sure that these map the Oracle service name to the appropriate network address.

You can use Oracle Network Manager to create the `tnsnames.ora` file for specific servers or network communities. Following is an example `tnsnames.ora` file that has been created to use with the `listener.ora` file you created in the procedure "To create an Oracle instance for failover." Note that if you are using only named pipes or only IP, your `tnsnames.ora` file will contain only the transport you are using:

## Oracle7 Server Installation and Configuration Requirements

```
#####
# Filename.....: tnsnames.ora
# Name.....: LOCAL_REGION.world
# Date.....: 27-JAN-97 10:47:58
#####
TestDB.world = _____ 1
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS =
        (COMMUNITY = NMP.world)
        (PROTOCOL = NMP)
        (Server = TestCluster) _____ 2
        (Pipe = testpipel) _____ 3
      )
      (ADDRESS =
        (COMMUNITY = TCP.world)
        (PROTOCOL = TCP)
        (Host = 16.151.8.202) _____ 4
        (Port = 1565) _____ 5
      )
    )
  (CONNECT_DATA =
    (SID = TEST) _____ 6
    (GLOBAL_NAME = TestDB.world) _____ 7
  )
)
```

In the previous file extract:

- 1 Specifies the Oracle service name or Oracle database alias.
- 2 Specifies the cluster alias as configured in the `listener.ora` file.
- 3 Specifies a unique pipe name matching the one configured in the `listener.ora` file.
- 4 Specifies the IP address configured in the `listener.ora` file. You can use an IP name instead provided the name can be translated to its IP address.
- 5 Specifies a unique port matching the one configured in the `listener.ora` file.
- 6 Specifies the Oracle SID.
- 7 Specifies the Oracle global database name.

## Client Connections During A Failover

During a failover, client connections to Oracle Server via the Oracle database alias are broken. Clients are responsible for handling this. Once failover has completed, clients should reconnect to the same Oracle database alias, which will connect them to the failover server.

## Initiating Manual Failover of an Oracle Instance

If you need to manually fail over an Oracle instance, DIGITAL strongly recommends that you first shut down the instance using the database administrator tools provided with the Oracle7 Server software. This guarantees orderly shutdown of the instance. Otherwise, current transactions will be rolled back instead of committed, and client applications will be responsible for handling this.



### **To initiate manual failover of an Oracle instance:**

1. Shut down the instance using an Oracle database administrator tool.
2. Initiate a manual failover using Cluster Administrator. See the section Managing Manual Failover in Chapter 8 for instructions.

---

# Configuring Lotus Notes for Failover

This chapter provides information on installing and setting up Lotus Notes in a DIGITAL Clusters for Windows NT environment. The chapter presents examples for the following configurations:

<b>DIGITAL Clusters Version</b>	<b>Lotus Notes Version</b>
DIGITAL Clusters 1.1	Domino 4.5 server and Notes 4.5 client
DIGITAL Clusters 1.0, Service Pack 1 or later	Notes 4.11 server and client

The chapter assumes you are familiar with Lotus Notes. For detailed information on installing and using Lotus Notes, refer to your Notes documentation.

## About Lotus Notes

Lotus Notes is a networked computing environment that allows workgroups in a company to share information and develop applications. Users can share documents in a central database and communicate through electronic mail.

Lotus Notes uses the client/server model—one or more Notes servers provide services to Notes workstation clients such as shared databases, mail routing, and replication. The Domino 4.5 server includes a web server that lets a user access Notes databases from a web browser rather than the Notes workstation client. The Domino administrator determines the level of access.

Lotus Notes also uses its own domains. Notes domains are distinct from Windows NT domains, but share some similar concepts. For example, after you set up one Notes server as the registration server in a Notes domain, you must register other Notes servers and clients. You must register the clients and additional servers *before* you install them.

## Domino 4.5 Installation and Configuration (V1.1)

This section describes one method to install and configure the Domino 4.5 server in a DIGITAL Clusters 1.1 environment. Before you start the installation, read the following section that describes some clustering methods.

### Clustering Methods

If you use the TCP/IP protocol for your Domino servers, you can take advantage of the IP failover feature in DIGITAL Clusters 1.1. If the primary cluster server for a Domino server fails, another instance of the same Domino server starts up on the secondary cluster server. Clients do not have to manually reconnect to the Domino server.

If you do not run Notes over TCP/IP, you can use the failover method described for Lotus Notes 4.11 (page 4–23). This method employs pointer files to let clients access their Notes database from another Notes server after a failover. However, clients must manually connect to the other Notes server.

### IP Failover for a Single Domino Server

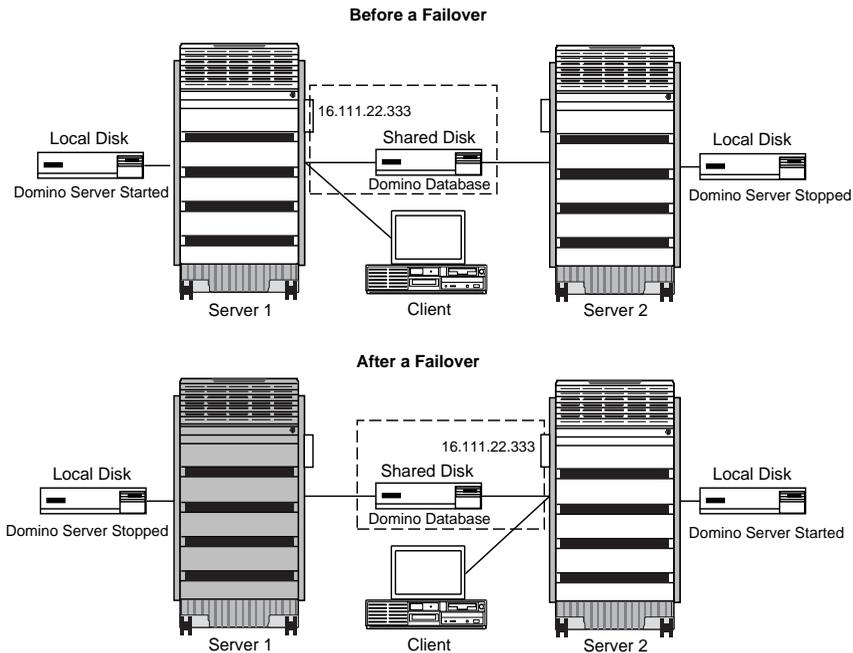
A Domino server lets you specify a fully qualified TCP/IP domain name for the server's name—for example, `srv1.dec.com`. Outside a cluster, you typically would use the domain name of the host system where the server is installed. However, by specifying a unique domain name for the Domino server, you can use the IP address for that name as a cluster IP address that can fail over from one cluster server to another. This method also provides failover support for the Domino web server (page 4–20).

The domain name and IP address must be unique. Do not use the domain name or IP address of a cluster server.

You install the Domino server program files locally on each cluster server, each pointing to the same Domino database on a shared disk. Then you place the cluster IP address failover object and shared disk in a failover group. Normally, the Domino server runs on the designated primary cluster server. If a failover occurs, the domain name and cluster IP address fail over to the secondary cluster. The failover is virtually transparent to connected clients.

The following figure shows a sample configuration for a single Domino server in a cluster.

### IP Failover Model for a Single Domino Server



You do not have to use fully qualified TCP/IP domain names to use for Domino servers. However, you must make sure that all clients can resolve the Domino server name to the correct cluster IP address, by using DNS, hosts files, or a similar method. In this case, you also must enter the selected cluster IP address in the Domino server's document in the Notes Address Book. This allows Domino servers and Notes clients to resolve the server name to an address. See the Notes documentation for more details.

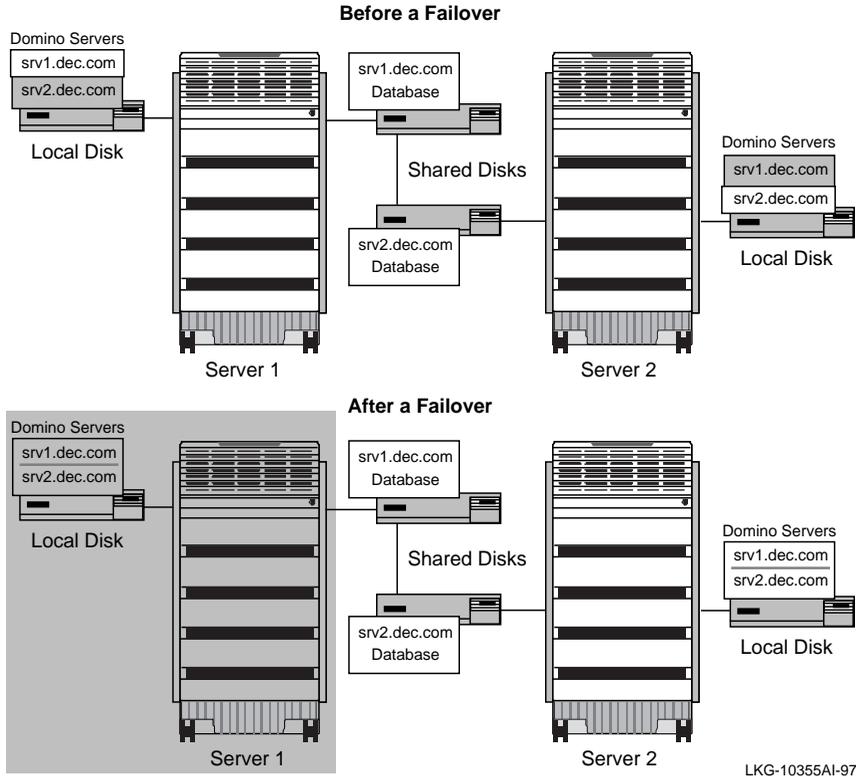
### IP Failover and Partitioned Domino Servers

Domino 4.5 allows you to install up to six partitioned Domino servers on a system. Using this method, you can run one or more Domino servers on each cluster server. You install each partitioned Domino server locally on both cluster servers. Each partitioned Domino server requires a cluster IP address.

# Domino 4.5 Installation and Configuration (V1.1)

The following figure shows a configuration for two partitioned Domino servers. When cluster server 1 fails, the Domino server `srv1.dec.com` and its database migrate to cluster server 2.

## IP Failover Model for Partitioned Domino Servers



## Before You Start

Before you start the installation, plan what groups of users you want to support from each cluster server. You can run multiple Domino servers on each cluster server.

Notes uses an organizational hierarchy. You can divide your organization into groups. Each group uses a unique certifier. When you register Notes servers and users for a group, you specify the group certifier.

The following procedure assumes you have a functional Version 1.1 cluster, with each cluster server acting as the primary server for one or more shared disks.

## Check Windows NT Settings

The Notes documentation recommends that you use the following settings in the Windows NT 4.0 Control Panel:

- Network applet → Services tab → Server → Properties → Maximum throughput for network applications
- System applet → Performance tab → Application Performance → Set to None to make foreground and background applications equally responsive.

## Create Failover Groups

In Cluster Administrator, create a failover group for each partitioned Domino server before the installation. Each group must contain:

- The shared disk where the Domino server's databases reside.
- The IP failover object that specifies the cluster IP address for the Domino server.  
This step makes the cluster IP addresses available when you register your Domino servers.

After the installation, you create a script failover object for each Domino server and add the object to the failover group.

## Installation

The following procedure describes the basic steps for installing multiple partitioned Domino servers on a cluster. The procedure installs two partitioned Domino servers on each cluster server, with one Domino server running on each cluster server.

## Software Requirements

You need the following items to install and configure Domino servers for failover:

- DIGITAL Clusters for Windows NT Version 1.1
- One or more Domino 4.5 servers and licenses for each cluster server

You can install and register one or more Domino 4.5 servers on a local disk of each cluster server, as described in this chapter. To install multiple partitioned Domino servers on a single system, you need a Domino Advanced Server license.

You need enough Notes client licenses for the number of Notes clients you plan to support. Register Notes clients as described in Notes documentation.

- A cluster IP address for each Domino server

**Recommended:** Register a fully qualified TCP/IP domain name for each cluster IP address used for Domino servers. During installation, you specify the domain name as the Domino server name.

## Domino 4.5 Installation and Configuration (V1.1)

If you do not use a fully qualified domain name, you must make sure that each Notes client can resolve the Domino server name to the cluster IP address.

- `Instsrv.exe` and `srvany.exe`, from the Windows NT 4.0 Resource Kit

You need these files to install partitioned Domino servers as Windows NT services.

### Setting Up Partitioned Domino Servers on the First Cluster Server

Setting up partitioned Domino servers on a cluster server involves four major steps:

1. Install the first Domino server and Notes client.
2. Register the additional Domino servers.
3. Install the additional Domino servers.
4. Install the Domino servers as Windows NT services.

### Install the First Domino Server and Notes Client on the First Cluster Server

This procedure assumes that you are installing the first Domino server in your Notes domain, which becomes the administrative server. If you already have a Domino administrative server on another system, you would first register your new Domino server before installing it.

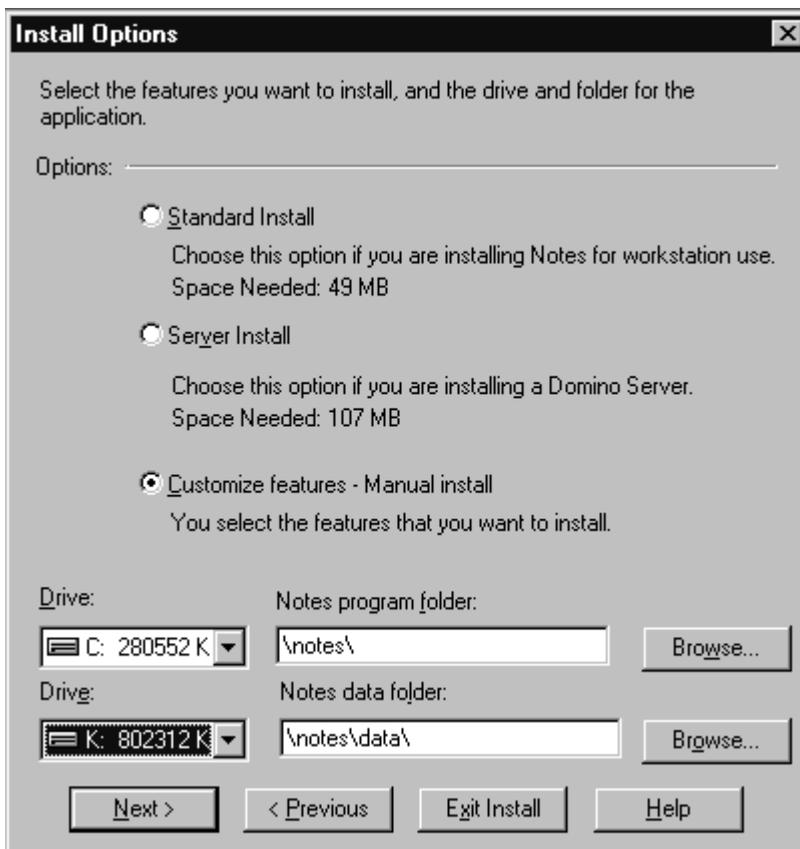
#### ► To install the first Domino server and Notes client on the first cluster server:

1. Access the Domino 4.5 CD-ROM from the server and start the installation for your platform:

```
\w32intel\install.exe or \w32alpha\install.exe
```

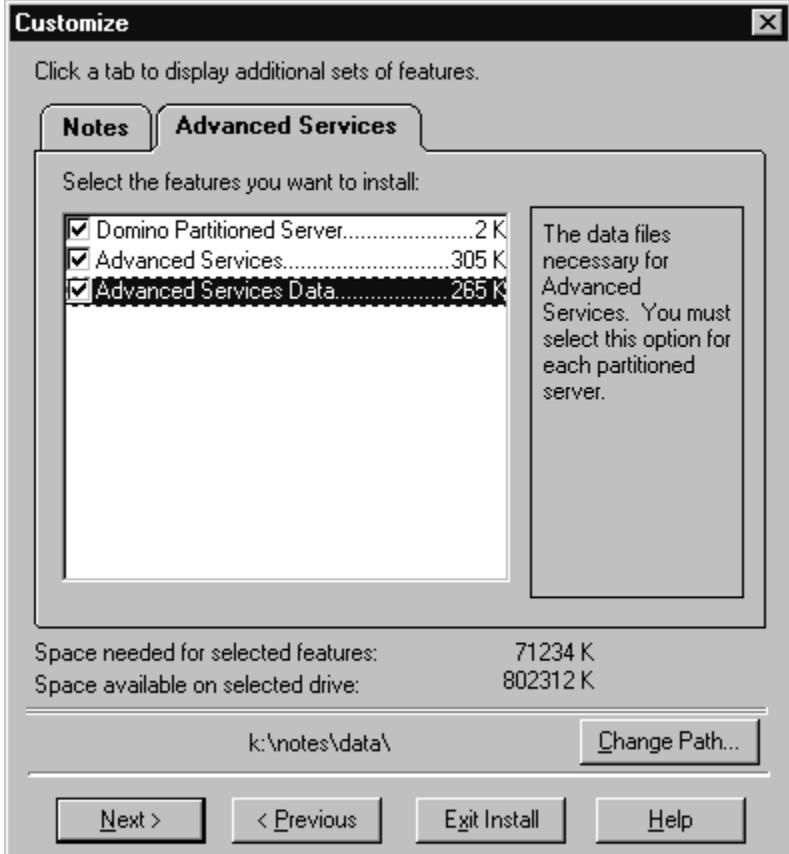
2. A dialog box prompts you to select the kits you want to install and specify where to store the Notes program folder and Notes data folder.
  - a. Check the box **Customize Features – Manual Install**.
  - b. Install the program folder on one of the cluster server's local unshared disks.
  - c. Install the data files on one of the cluster's shared disks available to this cluster server. When you create a failover group for the Domino server, you will include this disk in the group.

After you specify the drive locations, choose **Next**.



3. The Customize dialog box appears.
    - a. On the Notes tab, select the desired features. You need the Domino server software. You should also install the Notes workstation client software, if you want to administer your Notes environment from this system.
 

Some optional features, such as Single Password Logon, User Synchronization, and Notes Performance Monitor, create Windows NT services or registry settings. After the installation, you must reboot the cluster server for these settings to take effect.
    - b. On the Advanced Services tab, select all three features. These features provide the items needed to create a Domino partitioned server. When you select each feature, the software prompts you to confirm that you have an Advanced Server license.
- After you complete the Advanced tab, choose Next.



4. Select a program group, then confirm that you want to begin installing files.
5. In the program group, click the Notes workstation icon. You perform all setup steps from the Notes workstation client. The first time you run the client, a Notes Server Setup dialog box appears.
6. In the Notes Server Setup dialog box, specify that you are installing the first Notes server.
7. Fill in the First Server Setup dialog box.
  - a. Enter a name for the Domino server.

If you have registered a fully qualified TCP/IP domain name for the Domino server, enter that name. For example: `server1.dec.com`.

You do not have to use a fully qualified TCP/IP domain name. However, you must make sure that all clients can resolve the IP name to the correct cluster IP address, by using DNS, hosts files, or a similar method.

- b. Specify the administrator's name and password.
- c. Set the network type to TCP/IP.
- d. If you plan to use this system for administrative tasks, check the feature **Server is also administrator's personal workstation**. If not, you can select another Notes workstation later.

**First Server Setup**

Server name (e.g. Acme Server1):

Organization name (e.g. Acme Corp):

Administrator's last name:  First name:  M:

Administration password (case sensitive):

Network type:  Serial port:

Modem type:

Server is also administrator's personal workstation

Buttons: OK, Quit, Advanced Options..., Setup..., Script...

8. In the First Server Setup dialog box, choose Advanced Options and fill in the Advanced Server Setup Options.
  - a. Specify a Lotus Notes domain name and network name. Lotus Notes domains do not correspond to Windows NT domains.
  - b. Check the features **Create organization certifier ID**, **Create server ID**, and **Create administrator ID**. By default, this creates files named `cert.id`, `server.id`, and `user.id` in the `\notes\data` directory of the shared disk you specified in step 2.

## Domino 4.5 Installation and Configuration (V1.1)

9. Go to the Windows NT Start menu, select the Lotus Notes program group and start the Domino server. The server runs in its own command window.

If the server window displays a cache error, go to the Windows NT Control Panel and choose Network applet → Services tab → Server → Properties → Maximum throughput for network applications.

At this point, the Domino server is running manually. Later, you will install the Notes server as a Windows NT service, so it is available each time the cluster server restarts.

10. By default, your partitioned Domino server and client pairs are listed by number on the Start menu (Domino Server 1, Domino Server 2, and so on). You can use the Taskbar Properties to change these names to reflect the actual Domino server names.

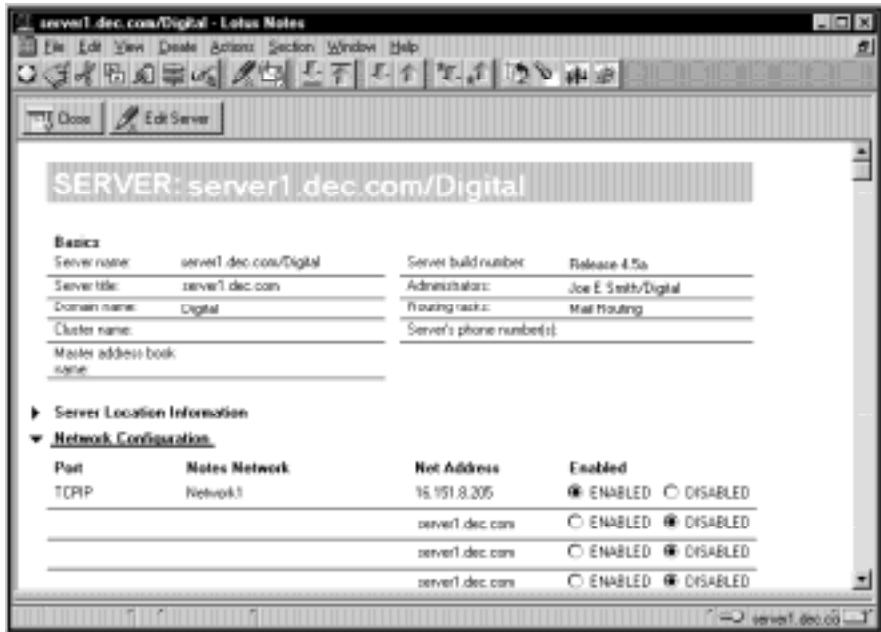
11. Add the cluster IP address for the Domino server to the server's document in the Address Book.

- a. Start the Notes workstation client, if not already open.
- b. Choose File menu → Databases → Open, and open the Public Address Book.
- c. In the Address Book, select Server → Servers. The list of known servers appears in the right pane.
- d. Double-click the name of your server to open its Server Document.
- e. In the Server Document, scroll down to the Network Configuration section. Click the triangle to view the Network Configuration fields.
- f. Choose Edit Server.

The first server entry is Enabled. In the first entry, click the Net Address field and replace the Domino server name with the cluster IP address.

Enter TCPIP in the Port field.

- g. Save the document and close the Address Book.



12. At this point, you may want to make your first Domino server the administration server for the Address Book. To do this, you choose File → Tools → Server Administration and select Database Tools. See your Notes documentation for details.

### Register Additional Domino Servers on the First Cluster Server

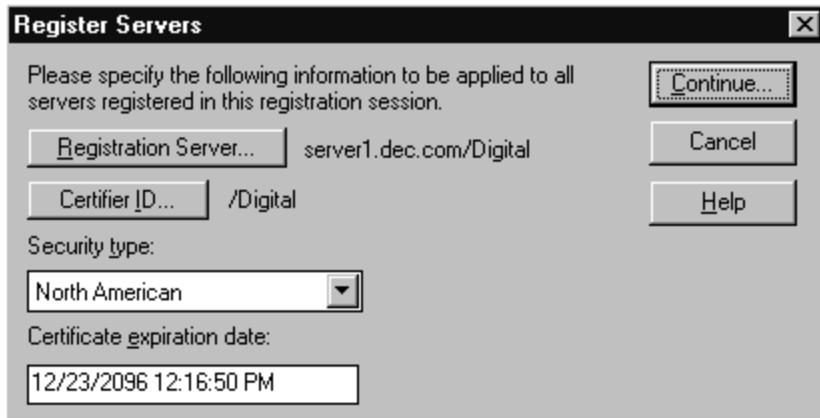
Before you install additional partitioned Domino servers on your first cluster server, you must register them. This procedure assumes that the first Domino server you installed is your administration server.

If you use fully qualified domain names for your Domino servers, the cluster IP addresses must be available at this point. Make sure you have created the failover groups and included IP failover objects. See the section Create Failover Groups (page 4–5).

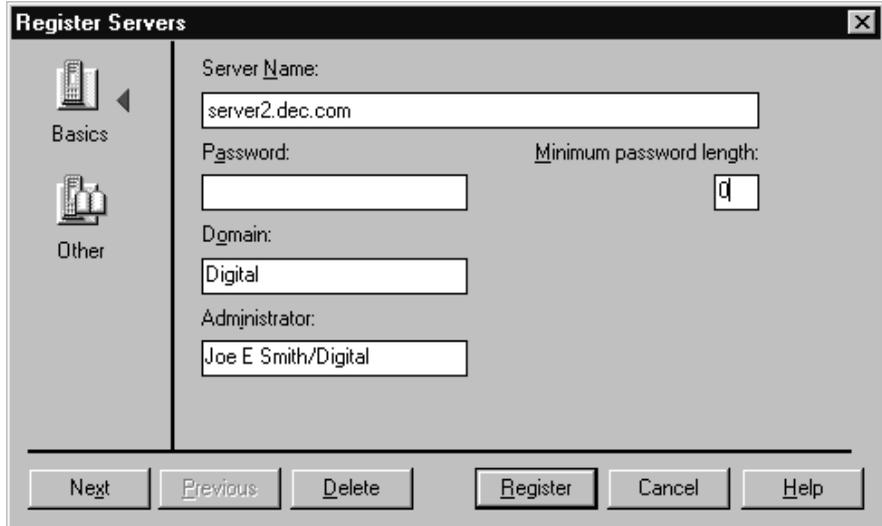
#### ► To register additional Domino servers on the first cluster server:

1. In the Windows NT Start menu, select the Lotus Notes program group and start the first Notes workstation client. In the client, pull down the File menu and choose Tools → Server Administration.
2. In the Server Administration dialog box, choose Servers → Register Servers.
3. In the Register Servers dialog box, specify the first Notes server as the Registration Server. Also specify the Certifier ID. Then choose Continue.

## Domino 4.5 Installation and Configuration (V1.1)



4. In the next Register Servers dialog box:
  - a. Specify a Server Name for the second Notes server.
  - b. Do not enter a password. Set the minimum password length to 0.
  - c. The Domain and Administrator fields are automatically filled in with Lotus domain and administrator information you specified previously.

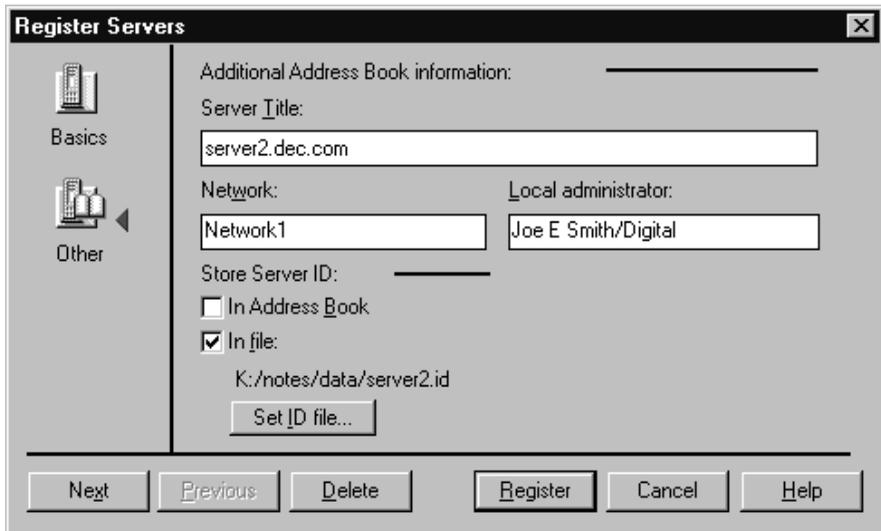


5. In the Register Server dialog box, choose Other.
  - a. If desired, enter a server title to use in the Notes Public Address Book.
  - b. Specify to store the server ID in a file. Choose Set ID file to specify the path to the ID file. For example, you might store the ID file in the \notes\data directory of the first server. By default, the file name is the first eight characters of the server name—in this example, server2.id.

You need the path to the new Domino server's ID file and the administrator's ID file (user.id) when you install the Domino server. Alternatively, you can copy the two files to a diskette.

DIGITAL does not recommend storing the server ID in the Address Book. Currently, the Address Book requires a password to access a listed Notes server. When a Notes server comes back on line after a failover, the system would stop during the failover script to prompt for the server password.

After you fill in the dialog box, choose Register.



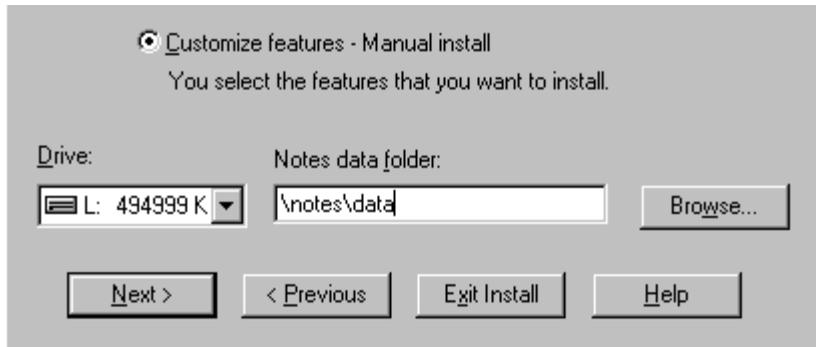
## Install a Second Domino Server on the First Cluster Server

This procedure assumes that the second partitioned Domino server will normally run on the second cluster server. For failover support, you must also install the second Domino server on the first cluster server.

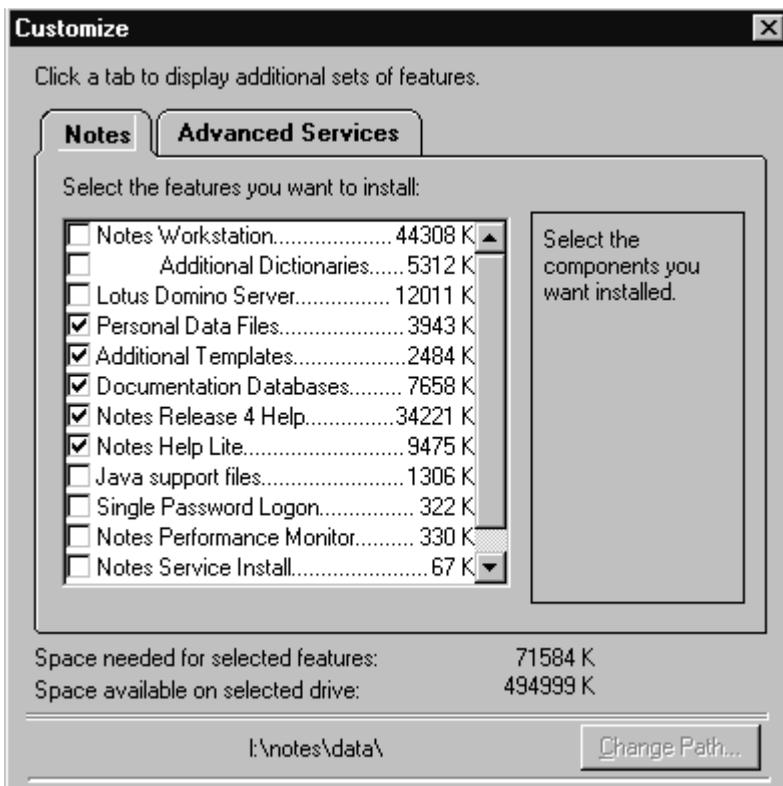
► **To install a second Domino server on the first cluster server:**

1. Before you can install the second Domino server, the first cluster server needs access to the shared disk where the second Domino server's databases will reside. If the disk is already in a failover group, go to Cluster Administrator and manually fail over the group so it is running on the first cluster server.
2. Repeat steps 1 to 5 of the installation procedure for the first Domino server (page 4–6).
  - At step 2, you only need to specify the location for the new Notes data folder, because the program files were already installed locally with the first Domino server. Specify a location on the shared disk that you failed over.

By default, Domino numbers the data directories (\data2, \data3, and so on). You can delete the number, unless you plan to install multiple server databases on the same disk.



- At step 3, on the Notes tab, do not specify the Domino server files and Notes workstation client files because they are already installed.



On the Advanced tab, specify Domino Partitioned Server and Advanced Services Data. Do not specify Advanced Services—this selection is only for the first partitioned Domino server.

- At step 5, specify that you are installing an additional Domino server.
3. Fill in the Additional Server Setup dialog box.
    - a. Enter a name for the second Domino server.

If you have registered a fully qualified TCP/IP domain name for the second Domino server, enter that name. For example: `server2.dec.com`.

If you do not use a fully qualified name, you must make sure that Notes clients can resolve the name to the cluster IP address.

- b. In the **Get Domain Address Book from** field, specify the first Domino server.

## Domino 4.5 Installation and Configuration (V1.1)

- c. Check the **New server's ID supplied in a file** and **Server is also administrator's personal workstation** features.

After you fill in the dialog box, choose Advanced Options.

**Additional Server Setup**

New server name:  
server2.dec.com

Get Domain Address Book from

Server name:  
server1.dec.com

Via network  
 Via serial port

Phone number:

Network type: TCP/IP

Serial port: (None)

Modem type: . Auto Configure (for unlisted modems, only)

New server's ID supplied in a file  
 Server is also administrator's personal workstation

Buttons: OK, Quit, Advanced Options..., Setup..., Script...

4. In the Advanced Server Setup Options, check the **Administrator's ID is supplied in a file** feature.
5. The software prompts you for the location of the second Domino server's ID file and the administrator's ID file (`user.id`), displaying a list of known drives. Specify the path where you stored the files when registering the server (page 4–13). If you stored the files on a diskette, insert the diskette.

DIGITAL recommends storing the server ID in a file rather than in the Address Book. Currently, the Address Book requires a password to access a listed Notes server. When a Notes server comes back on line after a failover, the system would stop during the failover script to prompt for the server password.

6. Add the cluster IP address for the second Domino server to the second server's document in the Address Book, as you did for the first server (page 4–10).

## Install the Domino Servers as Windows NT Services

After you install your partitioned Domino servers on the first cluster server, you must install the Domino servers as Windows NT services.

### Requirements

You need the following files from the Windows NT 4.0 Resource Kit:

- `instsrv.exe`
- `srvany.exe`

Currently, Domino servers started as Windows NT services do not survive after a logoff sequence.

### ► To install a partitioned Domino server as a Windows NT service:

1. On the first cluster server, open a command window. If the `instsrv.exe` and `srvany.exe` files are not in your current path, move to the directory where the files are located. Then enter the following command:

```
> instsrv Domino_name disk:path\srvany.exe -a domain\administrator -p password
```

where:

*Domino\_name* is the name of your partitioned Domino server as you want it to appear in the Windows NT Control Panel, Services applet.

*disk:path* is the full path to the `srvany.exe` file.

*domain*\administrator specifies the current Windows NT domain and an administrator account.

*password* specifies the password for the given administrator account.

The command window displays a message if the command is successful.

2. Go to the Windows NT Control Panel, Services applet. Your new Domino service should be listed. Select the service and choose Startup.
3. In the Service dialog box, select the following settings:
  - Startup Type: Manual
  - Log On As: System Account  
Allow Service to Interact with Desktop
4. Start the Registry Editor (`regedt32.exe`).
  - a. Under `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Domino`, create a Parameters key. *Domino* is the name you gave to the Domino service. When you create the key, leave the Class field blank.

## Domino 4.5 Installation and Configuration (V1.1)

- b. Under the Parameters key, create an Application value with the default type of REG\_SZ. Set the Application value to:

*Disk*: \notes\data\nserve.bat

where:

*Disk* is the letter of the shared disk where you installed the \notes\data files for this Domino server.

*nserve.bat* is the default command file for starting the partitioned server. The default file sets the Notes partition number, then starts the Domino executable:

```
Set NOTESPARTITION=1
Start C:\notes\nserver.exe (for Intel) or aserver.exe (for Alpha)
```

You may want to add a command in the file to display a startup message. For example:

```
net send %COMPUTERNAME% "Domino Server Name starting on %COMPUTERNAME%"
```

- c. Under the Parameters key, create an AppDirectory value with the default type of REG\_SZ. Set the AppDirectory value to:

*Disk*: \notes\data

### Setting Up Partitioned Domino Servers on the Second Cluster Server

After you set up your partitioned Domino servers on the first cluster server, you repeat the process on the second cluster server.

#### ► To set up partitioned Domino servers on the second cluster server:

1. In Cluster Administrator, manually fail over the groups that contain the shared disks where you installed your \notes\data directories when setting up the first cluster server.
2. Install the first Domino server on the second cluster server (page 4–6).

Install the Notes program folder on a local disk of the second cluster server.

Install the Notes data folder on the same shared drive and path that you used for the first Domino server on the first cluster server.

When you start the Notes client, you are not prompted to register the Domino server because you completed this step on the first cluster server.

3. Install the second Domino server on the second cluster server (page 4–14).

Install the Notes data folder on the same shared drive and path that you used for the second Domino server on the first cluster server.

When you start the second Notes client, you are not prompted to register the Domino server because you completed this step on the first cluster server.

4. Install the Domino servers as Windows NT services on the second cluster server (page 4–17). Use the same NT service names as on the first server.

### Creating Script Files and Completing the Failover Groups

After setting up the partitioned Domino servers on both cluster servers, you can:

- Create your failover script files.
- Complete your failover groups.
- Verify that your setup works by manually failing over groups.

### Creating Failover Scripts

For each Domino server, you need start and stop command files. These files contain the commands to start and stop the Domino server. Place a copy of the files for each Domino server in your designated scripts directory on each cluster server, then create a Script Failover Object in the Cluster Administrator and specify the path to these files.

For example, if your first Domino server has a service name of Domino Manufacturing, you could create the following two files:

#### **Start Domino Mfg.cmd**

```
rem Start script for the Domino Server Manufacturing, in Notes partition 1.
rem First stop the service. (See the end of the Stop script for details.)
rem Use the Windows NT service name for the Domino server.
net stop "Domino Manufacturing"
net start "Domino Manufacturing"
net send %COMPUTERNAME% "Starting Domino Server Manufacturing on %COMPUTERNAME%"
```

#### **Stop Domino Mfg.cmd**

```
rem Stop script for the Domino Server Manufacturing, in Notes partition 1.
rem Data directory is K:\notes\data
rem
net send %COMPUTERNAME% "Stopping Domino Server Manufacturing on %COMPUTERNAME%"
K:
cd \notes\data
Set NOTESPARTITION=1
rem Use "nserver" for Intel. Use "aserver" for Alpha.
c:\notes\nserver -quit
rem Do not stop the Windows NT service now. This may halt the shutdown procedure.
rem Stop the Windows NT service at the beginning of the Start script.
```

## Domino 4.5 Installation and Configuration (V1.1)

### Completing the Failover Groups

In Cluster Administrator, add the script failover object for a Domino server to the server's failover group. The group now contains:

- The shared disk where the Domino server's databases reside
- The IP failover object that specifies the cluster IP address for the Domino server
- The script failover object that specifies the start and stop commands for the server

### Adding Failover Support for the Domino Web Server

The Domino server includes an integrated web server that can let users access Notes databases from a web browser, if permitted by the Domino administrator. See the Notes documentation for a full description of Domino web server features and setup.

By default, the web pages for the Domino web server are stored in a directory under the `\notes\data` directory.

To add failover support for the web server, you simply set the web server task to start automatically whenever the Domino server starts:

1. Shut down the Domino server. You can enter `Exit` in the server command window.
2. In a text editor, open the `NOTES.INI` file.
3. At the end of the `ServerTasks=` line, add a comma and the command `http`. For example:  

```
ServerTasks=Replica, Router, Update, Stats, http
```
4. Save and close the file.
5. Restart the Domino server so the changes take effect.

#### Web Server Address

When started, the web server is available from the address of the Domino server. If you specified a fully qualified TCP/IP domain name for the Domino server, users can access the server by that name. For example:

```
http://myserver.digital.com
```

If you did not specify a fully qualified TCP/IP domain name, you must ensure that clients can resolve the server name to the cluster IP address. Otherwise, users cannot specify the server name for web access. In any case, users can enter the cluster IP address for the Domino server. For example:

```
http://11.22.33.44
```

### Running Multiple Web Servers

By default, a web server uses TCP/IP port 80 for connections. You can run multiple web servers by using a different port setting for each web server. For example, you could run web servers for two partitioned Domino servers. You should check with your network administrator and read your web server documentation before activating multiple web servers. If you use a port number other than 80, users must specify the port in the URL.

```
http://myserver.digital.com:81
```

You can specify the TCP/IP port for a Domino web server in the HTTP section of the server's document, stored in the Address Book. See the Notes documentation for details.

## Lotus Notes 4.11 Installation and Configuration (V1.0 SP1)

This section describes how to install and configure Lotus Notes 4.11 in a DIGITAL Clusters 1.0 environment. Before you start the installation, read the following sections that describe clustering and failover methods.

### Software Requirements

You need the following items to install configure Lotus Notes for failover:

- DIGITAL Clusters for Windows NT Version 1.0, Service Pack 1 or later
- Lotus Notes 4.11 server and license for each cluster server

Install and register a Lotus Notes 4.11 server on a local disk of each cluster server, as described in this chapter.

You need enough Notes client licenses for the number of Notes clients you plan to support. Register Notes clients as described in Notes documentation.

### Clustering Methods

There are two basic choices for configuring Lotus Notes 4.11 in a DIGITAL Clusters environment:

- For larger Notes environments, you can run separate Notes servers on each cluster server with their own databases.
- For smaller Notes environments, you may want to run Notes server on only one cluster server and use the other Notes server for failover purposes only.

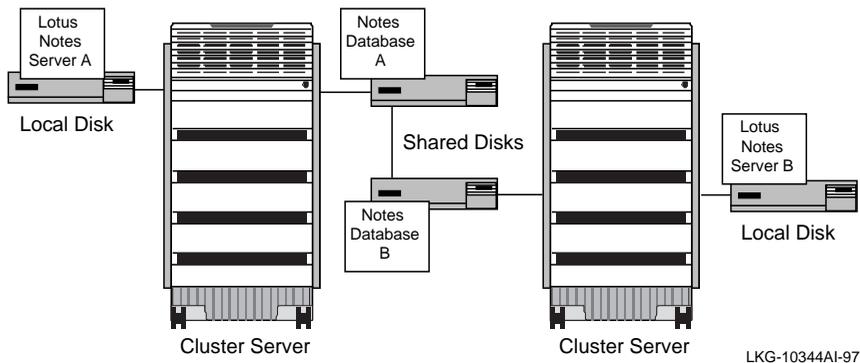
In either case, you install a Notes server locally on each cluster server.

## Setting Up Databases on Both Servers

This section describes how you might install and run separate Notes servers on each cluster server. In this case, each Notes server has its own database residing on a shared disk assigned to the cluster server for that Notes server. This approach allows you to set up separate Notes workgroups on each server, accessing their own databases.

The following figure shows a sample Notes configuration, with a Notes database on each cluster server.

### Sample Lotus Notes 4.11 Failover Configuration



For example, you might set up a Manufacturing workgroup on one server and a Finance workgroup on the other server. Each Notes server can support its own database stored on the cluster's shared disks. If one cluster server fails, the Notes database on that system can fail over to the Notes server on the other cluster server.

For smaller Notes environments where you want to run the Notes server on only one cluster, you can use a similar approach. You would still install a Notes server on each cluster server, but you would start the second Notes server only when the first Notes server failed.

Lotus Notes 4.11 does not support connections using a cluster alias. A Notes client must know the name of the server supporting its database. If a server fails, Notes users must manually connect to the other server and access their database again. You can use a failover script to notify users when a failure occurs.

## Failover Method: A Script and Notes Directory Pointer Files

You can create a simple failover method for Lotus Notes 4.11 by using brief failover scripts and Notes directory pointer files.

The Notes server keeps track of the drive location of databases, so users do not have to enter drive letters. To open a Notes database, users select from the list of available databases presented by their Notes server.

- **Directory Pointer Files.** You can add a directory pointer file to the list of Notes databases. A pointer file is an ASCII file that provides the path to a database on another disk. The pointer file appears as a folder in the list of databases that users see. When users open the folder, they see the available databases on the other disk. In a cluster, pointer files provide a simple method for accessing the database of a failed Notes server from the other Notes server.
- **Failover Scripts.** You can use failover scripts to notify Notes users when their server fails, so they can reconnect to the other server.

You also must use a failover script to ensure that a Notes server does not restart before the cluster software when the server comes back on line.

See the examples of pointer files and failover scripts later in this guide.

## Installation

The following procedure describes the basic steps for installing a separate Notes server and database on each cluster server.

### Before You Start

Before you start the installation, plan what groups of users you want to support from each cluster server. Notes uses an organizational hierarchy. You can divide your organization into groups. Each group uses a unique certifier. When you register Notes servers and users for a group, you specify the group certifier.

The following procedure assumes you have a functional cluster with one or more shared disks assigned to each cluster server.

The Lotus Notes documentation recommends that you use the following settings in the Windows NT Control Panel:

- Network applet → Server → Configure → Maximum throughput for network applications
- System applet → Tasking → Foreground and background applications equally responsive

## Install Lotus Notes on the First Cluster Server

The following procedure assumes that you are installing the first Notes server in your Notes domain. If you already have a Notes registration server on another system, you would register your new Notes server before installing it on the cluster.

### ► To install Lotus Notes on the first cluster server:

1. Access the Lotus Notes 4.11 CD-ROM from the server and start the installation for your platform:

`\w32intel\install.exe` or `\w32alpha\install.exe`

The installation prompts you to select the kits you want to install. Check the box to install both the Notes server and workstation client.

2. A dialog box asks you to specify where to store the Lotus Notes program files and data files.
  - a. Install program files on one of the cluster server's local unshared disks.
  - b. Install the data files on one of the cluster's shared disks assigned to this cluster server. When you create a failover group for the Notes server, you will include this disk in the group.

The installation creates a Lotus Notes program group.

3. In the program group, click the Workstation icon. You perform all setup steps from the Workstation client. The first time you run the Workstation client, a Notes Server Setup dialog box appears.
4. In the Notes Server Setup dialog box, specify that you are installing the first Notes server.
5. Fill in the First Server Setup dialog box.
  - a. Set the server name to the cluster server's computer name.
  - b. Specify the administrator's name and password.
  - c. Specify the network type: TCP/IP, NetBIOS, or IPX.
  - d. If you plan to use this system for administrative tasks, check the feature **Server is also administrator's personal workstation**. If not, you can select another Notes workstation later.

## Lotus Notes 4.11 Installation and Configuration (V1.0 SP1)



6. In the First Server Setup dialog box, choose Advanced Options and fill in the Advanced Server Setup Options.
  - a. Specify a Lotus Notes domain name and network name. Lotus Notes domains do not correspond to Windows NT domains.
  - b. Check the features **Create organization certifier ID**, **Create server ID**, and **Create administrator ID**. By default, this creates files named `cert.id`, `server.id`, and `user.id` in the `/notes/data` directory of the shared disk you specified in step 2.



## Lotus Notes 4.11 Installation and Configuration (V1.0 SP1)

7. Return to the Lotus Notes program group and click the Server icon to start the Lotus Notes server.

If the server window displays a cache error, go to the Windows NT Control Panel. Open the Network applet, choose Server, then choose Maximum Throughput for Network Applications.

At this point, the Notes server is running manually. You must install the Notes server as a Windows NT service, so it is available each time the cluster server restarts.

8. Go to the \notes directory on the local disk where you installed the Notes program files. Run the program `ntsvinst.exe`.

The program installs the Notes server as a Windows NT service in the Control Panel, Services applet.

---

### Caution

---

By default, the Notes server is set to start up manually and interact with the desktop. Do not change these default settings. If you set the server to start automatically, Notes will start before the cluster software is initialized and display an error message that it cannot find the database directory. Use a failover script to start the Notes server after the cluster server returns on line.

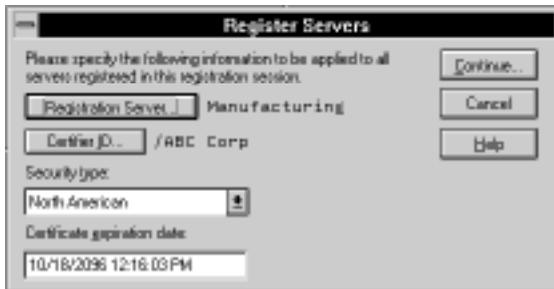
---

## Registering the Second Lotus Notes Server

Before you install Lotus Notes on your second cluster server, you must register the second Notes server on your first Notes server.

► **To register the second Lotus Notes server:**

1. From the Notes Workstation client on your first cluster server, click the File menu and choose Tools → Server Admin.
2. In the Server Administration dialog box, choose Servers.
3. In the Register Servers dialog box, specify the first Notes server as the Registration Server. Also specify the Certifier ID. Then choose Continue.



4. In the next Register Servers dialog box:
  - a. Specify a Server Name for the second Notes server.
  - b. Do not enter a password. Set the minimum password length to 0.
  - c. The Domain and Administrator fields are automatically filled in with Lotus domain and administrator information you specified previously.



5. In the Register Server dialog box, choose Other to display additional settings.
  - a. If desired, enter a server title to use in the Notes Address Book.
  - b. Specify to store the server ID in a file. Choose Set ID file to specify the path to the ID file. For example, you might store the ID file in the \notes\data directory of the first server. By default, the file name is the server name—in this case, finance.id.

You need the path to the second server's ID file and the administrator's ID file (user.id) when you install the second Notes server. Alternatively, you can copy the two files to a diskette.

DIGITAL does not recommend storing the server ID in the Address Book. Currently, the Address Book requires a password to access a listed Notes server.

## Lotus Notes 4.11 Installation and Configuration (V1.0 SP1)

When a Notes server comes back on line after a failover, the system would stop during the failover script to prompt for the server password.

After you fill in the dialog box, choose Register.



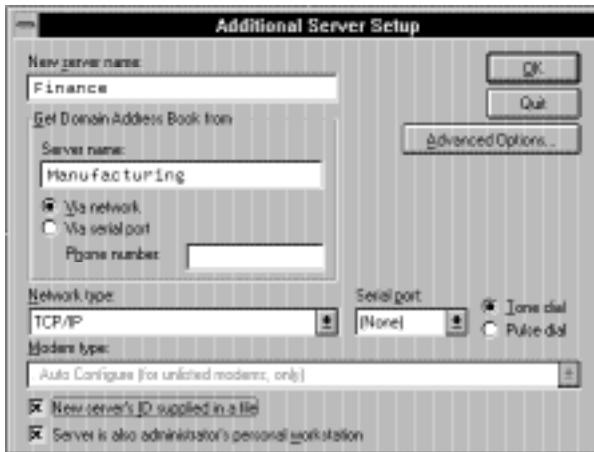
### Installing Lotus Notes on the Second Cluster Server

After you register the second Lotus Notes server, you can install it on the second cluster server.

#### ► To install Lotus Notes on the second cluster server:

1. Repeat steps 1 to 3 of the section Installing Lotus Notes on the First Cluster Server (page 4–24), but this time using the second cluster server's local disks and shared disks.  
At step 4, specify that you are installing an additional Notes server.
2. Fill in the Additional Server Setup dialog box
  - a. Specify the server name to match the second cluster server's computer name.
  - b. In the **Get Domain Address Book from** field, specify the first Notes server.
  - c. Check the **New server's ID supplied in a file** and **Server is also administrator's personal workstation** features.

After you fill in the dialog box, choose Advanced Options.



3. In the Advanced Server Setup Options, check the **Administrator's ID is supplied in a file** feature.
4. The software prompts you for the location of the second server's ID file and the administrator's ID file (`user.id`), displaying a list of known drives. Specify the path.

The shared drive for the first Notes server will not appear in the list, because it is controlled by the first cluster server. If you stored the files on the first server, choose Network and enter the path to the files (in our example, the `\notes\data` directory on the first cluster server). Then copy the second server's ID file (in this case, `finance.id`) and the `user.id` file from the first server to the `\notes\data` directory on the second cluster server.

DIGITAL recommends storing the server ID in a file rather than in the Address Book. Currently, the Address Book requires a password to access a listed Notes server. When a Notes server comes back on line after a failover, the system would stop during the failover script to prompt for the server password.

5. Repeat steps 7 and 8 in the section Lotus Notes on the First Cluster Server (page 4–24) to start the Notes server and install it as a Windows NT service.

---

**Note**

---

You should perform these steps even if you plan to use the second Notes server only as a failover backup for the first server. In this case, you can stop the second Notes server after you confirm it is configured correctly and installed as a Windows NT service.

---

## Registering Users and Installing Notes Workstation Clients

Refer to the Lotus Notes documentation to register users and install the Notes workstation client software on their system. You must register users before installing the Notes workstation software. Distribute the user IDs to the users or administrator setting up the workstation software.

Because Notes clients cannot use the cluster alias feature yet, you do not need to install cluster client software if you are running only Notes on the cluster.

You can set up a Notes Profile for a group in advance, identifying Connection documents that you want to assign to users. If you want to register a group of users at one time, you can set up a text file that contains the user names and Profile name.

## Creating the Directory Pointer Files

After setting up the two Notes servers on the cluster, you can create a directory pointer file for each server. The pointer file contains a single line with the path to the Notes server's database directory. For example, if you were creating the pointer file for a Notes server named Manufacturing whose database is on drive X, you could create the following ASCII pointer file:

**Manufacturing.dir**

```
x:\notes\data
```

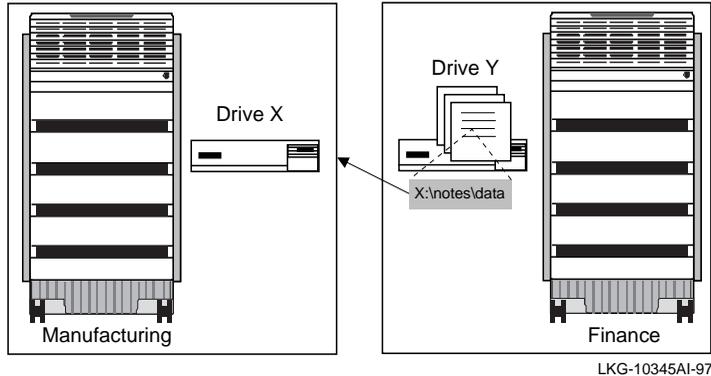
You can also restrict database access with this file by adding group names to the file.

Place the pointer file for a Notes server in the `\notes\data` directory of the other Notes server. For example, if the other Notes server were named Finance and used drive Y for its database, you would place the `Manufacturing.dir` file in `Y:\notes\data`.

When Notes users accessing the Finance server display the Open Database dialog box, the `Manufacturing.dir` pointer file name would appear as a directory in the list of databases. If users clicked the Manufacturing directory while the Manufacturing server is running, Notes would display an error message because drive X is not accessible.

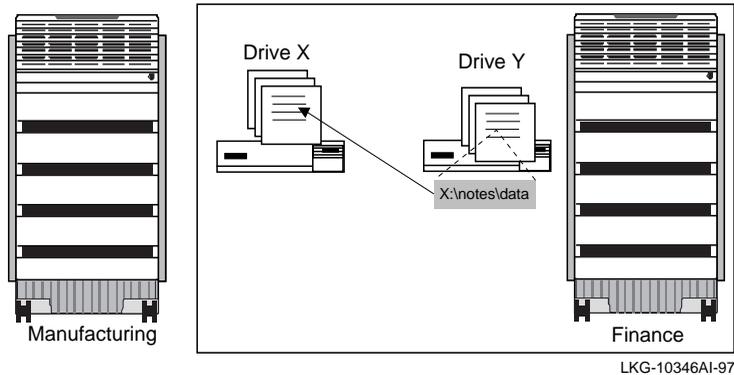
The following figure shows this situation.

### Pointer Files: Database Access Before a Failover



If the Manufacturing server database fails over, drive X is now accessible to Notes users connected to the Finance cluster server. These Notes users can now see a list of available databases in the Manufacturing database, as the following figure shows.

### Pointer Files: Database Access After a Failover



---

#### Note

If you plan to use only one Notes server, you would need only one pointer file in the database of the second, inactive Notes server.

---

## Sample Script Files

The Cluster Administrator lets you create a Script Failover Object, which consists of a pair of command scripts that run when a failover group comes on line or goes off line. Use a stop command script to notify Notes users when their database drive goes off line. Use a start command script to restart the Lotus Notes Server service when the drive comes back on line.

The following two script examples are for database drive X on the Manufacturing cluster server. You would create a similar pair of scripts for the Notes Server database on the other cluster server.

If you plan to use only one Notes server, you would create a pair of scripts for that Notes server only. Your stop command script would need an additional command to start the second Notes server at the appropriate point. Similarly, your start command script would need a command to stop the second Notes server.

### Stop Command Script: DriveX-offline.cmd

The stop command script runs when the failover group that includes drive X goes off line. The script notifies users connected to the Manufacturing cluster server that the drive is off line and the Lotus Notes Server is being stopped. Then the script stops the Lotus Notes Server.

```
REM Failover Script:      c:\cluster_scripts\driveX-offline.cmd
REM This script runs when the failover group "Database Disk X"
REM (drive X:) is going off line.
REM Drive X's primary host is Manufacturing, stop NOTES on primary host ONLY.
REM
@ set PRIM="Manufacturing"
@ set LOG=c:\cluster_logs\DriveX-status.LOG
@ if "%COMPUTERNAME%"=="Manufacturing" set otherhost="Finance"
@ if "%COMPUTERNAME%"=="Finance" set otherhost="Manufacturing"
@ now >> %LOG%
@ net send %otherhost% "Node %COMPUTERNAME% Lost connection to Drive X..."
@ hostname >> %LOG%
@ echo %0 %1 >> %LOG%
@ echo Message: Node %COMPUTERNAME% lost connection to Drive X >> %LOG%
@ if NOT "%COMPUTERNAME%"=="%PRIM%" goto skip_notes
@     echo Message: Stopping Lotus Notes Service >> %LOG%
@     echo ----- >> %LOG%
@     net stop "Lotus Notes Server" 2>> %LOG%
@ :skip_notes
@ now >> %LOG%
@ echo ----- >> %LOG%
```

### Start Command Script: DriveX-online.cmd

The online script runs when drive X comes back on line. The script notifies users connected to the Manufacturing server that drive X is back on line and the Lotus Notes Server is starting. Then the script starts the Lotus Notes Server on the Manufacturing server. Starting the Lotus Notes Server in a script ensures that the service does not start before the cluster software has initialized. If you set the service to automatically start at system startup, the service might start before the cluster software.

```

REM Failover Script:      c:\cluster_scripts\driveX-online.cmd
REM This script runs when the failover group "Database Disk X"
REM (drive X:) is brought on line by an available node.
REM Drive X's primary host is Manufacturing, start NOTES on primary host
ONLY.

REM
echo on

@ set PRIM="Manufacturing"
@ set LOG=c:\cluster_logs\DriveX-status.LOG
@ if "%COMPUTERNAME%"=="Manufacturing" set otherhost="Finance"
@ if "%COMPUTERNAME%"=="Finance" set otherhost="Manufacturing"
@ now >> %LOG%
@ net send %otherhost% "Node %COMPUTERNAME% brought Drive X on line"
@ echo
@ hostname >> %LOG%
@ echo %0 %1 %2 >> %LOG%
@ echo Message: Node %COMPUTERNAME% brought Drive X on line >> %LOG%
@ if NOT "%COMPUTERNAME%"==%PRIM% goto skip_notes
@     echo Message: Start Lotus Notes Service >> %LOG%
@     echo ----- >> %LOG%
@     net start "Lotus Notes Server" 2>> %LOG%
@ :skip_notes
@ now >> %LOG%
@ net send %otherhost% "Drive X Startup Complete"
@ echo ----- >> %LOG%
    
```

## Lotus Notes 4.11 Failover Issues

Consider the following issues when configuring Lotus Notes 4.11 for failover:

- Lotus Notes 4.11 does not support the use of a cluster alias. Notes clients must know the name of the systems they are connecting to. You do not need to install cluster client software if you are using only Lotus Notes.
- To support access to a cluster server's Notes database after a failover, you must add a directory pointer file in the other server's Notes database directory. Clients can then connect to the other server and use the pointer file to access the database.
- After a failover, Notes clients on the failed server must wait for the failover script to complete before accessing their database from the other server. The waiting period depends on the script's complexity.

Clients must manually reconnect to the other server and redo any transactions that were in progress when the system failed.

- The Notes name and address book is automatically replicated on all servers in a Notes domain, so each server maintains an updated account of authorized users.
- Failback of groups should be disabled. You don't want users abruptly terminated a second time when a failed server comes back on line.

---

# Configuring Web Servers for Failover

## 1.1

DIGITAL Clusters for Windows NT 1.1 provides support for failing over IP socket-based applications, such as web servers. This chapter describes how to install and configure two popular web servers for IP failover:

- Microsoft Internet Information Server (IIS)
- Netscape Enterprise Server

For information on configuring a Domino web server, see Chapter 4.

The configuration procedure for other web servers is similar to those documented in this chapter. See your web server documentation for general information on your web server.

## Overview

To access pages on a web server, users specify a uniform resource locator (URL) address in their browser. The URL address begins with a domain name or IP address for the server.

You can install a copy of the same web server on each cluster server and use a cluster IP address for the web server's address. You specify which cluster server the web server normally runs on. If that cluster server fails, the web server's IP address and data files can migrate to the other cluster server.

You can associate a fully qualified TCP/IP domain name with a cluster IP address—for example, `srv1.dec.com`. This allows users to specify the name rather than the cluster IP address in the URL. You should register a unique name for the web server. Do not use the name or IP address of an individual cluster server.

You do not have to use fully qualified TCP/IP domain names for web servers. However, you must make sure that all clients can resolve the name to the correct cluster IP address, by using DNS, hosts files, or a similar method.

## Overview

The basic configuration steps are similar for all web servers:

- Create or select a failover group for the web server. The failover group must contain the shared disk where you plan to install the web server's data directory.
- Install the web server and its administrative files on a local disk of each cluster server.
- Install the web server's data directory (the root directory for web documents) on a shared disk.
- Create an IP failover object that specifies the cluster IP address to use for the web server. This address cannot be the IP address of either cluster server. Add the IP failover object to the web server's failover group.
- Create a script failover object to start and stop the web server. In some cases, you need to create batch command files. Add the script failover object to the web server's failover group.

## Before You Start

Consider the following points when planning your web server configuration:

- Only one instance of each web server can run on the cluster at a time. Determine which cluster server you want to use as the primary server for your Internet/Intranet access.

You can run different web servers at the same time, if you specify that they use different TCP/IP ports. For example, you might use the IIS server on one cluster server and the Netscape Enterprise Server on the other cluster server. If one cluster server fails, the other cluster server would run both web servers.

By default, a web server uses TCP/IP port 80 for connections. If you use a port number other than 80, users must specify the port in the URL. For example:

```
http://myserver.digital.com:81
```

Check with your network administrator and read your web server documentation before activating multiple web servers. Make sure any ports you specify are not already in use by other services.

- You store the web server's data files on a shared cluster disk that is a member of a failover group. Determine what other applications you may want to install on the same disk, then create the appropriate failover objects and scripts. Add the objects and scripts to the failover group.

## Microsoft IIS Installation and Configuration

This section describes how to install and configure the Microsoft Internet Information Server (IIS) 3.0 on a cluster. For more details on configuring specific IIS features, see the IIS online help.

### Software Requirements

You need the following items to install IIS. Make sure you have the correct software for your platform (Alpha or Intel).

- DIGITAL Clusters for Windows NT Version 1.1
- Microsoft IIS for each cluster server
- Cluster IP address

**Recommended:** Register a fully qualified TCP/IP domain name for the cluster IP address. Users can specify this name as the URL for the web server.

If you do not use a fully qualified domain name, make sure that each client can resolve the web server's name to the cluster IP address.

### Installation

You must install the IIS software on both cluster servers.

#### ► To install the IIS software on both cluster servers:

1. In Cluster Administrator, select or create a failover group to use for the web server. Add the shared disk that you plan to use for web documents to the group.

Make sure the shared disk has the same fixed drive letter assigned on both cluster servers. For details, see Chapter 3 of the *Configuration and Installation Guide*. You can test this by manually failing over the group and checking that the drive appears in Windows NT Explorer or File Manager on the server that takes over the group.

2. On both cluster servers, install IIS on a local disk. Choose the default installation options, which install all files locally.
3. On both servers, go to the Windows NT Control Panel, Services applet, and set the FTP Publishing Service, Gopher Publishing Service, and World Wide Web Publishing Service to start manually.
4. Create the web server database on the shared disk. To do this, you can copy the `\InetPub` directory tree from the local disk on either server where you installed IIS.

## Microsoft IIS Installation and Configuration

5. On both servers, use the Microsoft Internet Service Manager to point the WWW, Gopher and FTP database services to the web server database on the shared disk. To do this, double-click on each service and edit the root directory entries on the Properties → Directories page to specify the shared disk drive letter.

### Configuration

After you install IIS on both cluster servers, you can configure it for cluster use.

#### ► To configure IIS for failover:

1. In Cluster Administrator, create an IP failover object. Specify a cluster IP address that you want clients to use for accessing the IIS server. This address cannot be the individual address of either cluster server.
2. In Cluster Administrator, create a failover script object.
  - a. Create a pair of start and stop batch command files for the first cluster server and store them on a local drive. The files include commands to start and stop the web, ftp, and gopher services provided by IIS. For example:

#### C:\Scripts\IIS\_Start.bat

```
REM C:\Scripts\IIS_Start.bat file.  
REM Start IIS services when the failover group comes on line.  
@echo off  
@net start w3svc  
@net start msftpsvc  
@net start gophersvc
```

#### C:\Scripts\IIS\_Stop.bat

```
REM C:\Scripts\IIS_Stop.bat file.  
REM Stop IIS services when the failover group goes off line.  
@echo off  
@net stop w3svc  
@net stop msftpsvc  
@net stop gophersvc
```

- b. Copy the start and stop command files to an identical directory path on the second server. For example, if the files are in C:\Scripts on the first server, they must be in C:\Scripts on the second server.
- c. In Cluster Administrator, create a script failover object and specify the start and stop batch command files.

Start Command: **C:\Scripts\IIS\_Start.bat**

Stop Command: **C:\Scripts\IIS\_Stop.bat**

## Netscape Enterprise Server Installation and Configuration

3. In Cluster Administrator, select the failover group containing the shared disk with the IIS web database and add:
  - The IP failover object
  - The script failover object
4. To verify the configuration:
  - a. Go to a client and open a web browser. In the browser, enter the web server's domain name or associated cluster IP address as the URL. For example:  

```
http://wbsrvr.dec.com      or      http://11.222.33.44
```

You should see the default home page of the web server.
  - b. In Cluster Administrator, manually fail over the web server's group and wait for the web server to migrate to the other cluster server. When the web server is running on the other cluster server, repeat step a to confirm that you can still access the web server.

## Netscape Enterprise Server Installation and Configuration

This section describes how to install and configure Netscape Enterprise Server 2.0 (Intel) or 2.01 (Alpha) for failover support.

### Software Requirements

You need the following items to install the Netscape Enterprise Server. Make sure you have the correct software for your platform (Alpha or Intel).

- DIGITAL Clusters for Windows NT Version 1.1
- Netscape Navigator™ or Navigator Gold web browser installed on a local drive of each cluster server
- Netscape Enterprise Server for each cluster server
- Cluster IP address

**Recommended:** Register a fully qualified TCP/IP domain name for the cluster IP address. Users can specify this name as the URL for the web server.

If you do not use a fully qualified domain name, make sure that each client can resolve the web server's name to the cluster IP address.

### Installation

You must install the Netscape Enterprise Server software on each cluster server. Before you start, you must have the Netscape Navigator or Navigator Gold web browser installed locally on each cluster server.

# Netscape Enterprise Server Installation and Configuration

## ► To install the Netscape Enterprise Server software on both cluster servers:

1. In Cluster Administrator, select or create a failover group to use for the web server. Add the shared disk that you plan to use for web documents to the group.

Make sure the shared disk has the same fixed drive letter assigned on both cluster servers. For details, see Chapter 3 of the *Configuration and Installation Guide*. You can test this by manually failing over the group and group and checking that the drive appears in Windows NT Explorer or File Manager on the server that takes over the group.

2. Install the Netscape Enterprise Server on the cluster server that currently has access to the shared drive you plan to use for web documents.
  - Install the server files on a local drive. The default location is `C:\netscape\server`.
  - When prompted, install the server's document root directory on a shared drive. For example, if your shared drive is K, you might specify the directory `K:\netscape\server\docs`.



3. In the Windows NT Control Panel, Services applet, set the **Netscape Enterprise Server https-servername** service to start manually.
4. In Cluster Administrator, manually fail over the group that contains the shared disk so the disk is available on the second cluster server.
5. Repeat steps 2 and 3 to install the Netscape Enterprise Server on the second cluster server.

# Netscape Enterprise Server Installation and Configuration

- When prompted, install the server's document root directory on the shared drive in the same directory you specified for the first server—in this example, `K:\netscape\server\docs`.

## Configuration

After you install the Netscape Enterprise Server on both cluster servers, you can configure it for cluster use.

### ► To configure the Netscape Enterprise Server for failover:

1. In Cluster Administrator, create an IP failover object. Specify a cluster IP address you want clients to use for accessing the Netscape Enterprise Server. This address cannot be the individual address of either cluster server.
2. Create a script failover object to start and stop the Netscape Enterprise Server service.
  - a. Create a pair of start and stop batch command files for the first cluster server and store them on one of the server's local drives. Note that the Netscape Enterprise Server service name includes the cluster server name.

For example, if the first server is named Manufacturing and you plan to store the files in `C:\scripts`, you could create the following command files:

#### **C:\Scripts\Netscape-server-start.bat**

```
REM C:\Scripts\Netscape-server-start.bat file
REM Start the Netscape Enterprise Server on server Manufacturing.
@echo off
@net start "Netscape Enterprise Server https-Manufacturing"
```

#### **C:\Scripts\Netscape-server-stop.bat**

```
REM C:\Scripts\Netscape-server-stop.bat file
REM Stop the Netscape Enterprise Server on server Manufacturing.
@echo off
@net stop "Netscape Enterprise Server https-Manufacturing"
```

- b. Repeat step a for the second cluster server. The path and file names on the second server must match those used on the first server.

For example, if the second server is named Finance, you could create the following command files:

#### **C:\Scripts\Netscape-server-start.bat**

```
REM C:\Scripts\Netscape-server-start.bat file
REM Start the Netscape Enterprise Server on server Finance.
@echo off
@net start "Netscape Enterprise Server https-Finance"
```

## Netscape Enterprise Server Installation and Configuration

### C:\Scripts\Netscape-server-stop.bat

```
REM C:\Scripts\Netscape-server-stop.bat file
REM Stop the Netscape Enterprise Server on server Finance.
@echo off
@net stop "Netscape Enterprise Server https-Finance"
```

- c. In Cluster Administrator, create a script failover object and specify the start and stop batch command files.

Start Command: **C:\Scripts\Netscape-server-start.bat**

Stop Command: **C:\Scripts\Netscape-server-stop.bat**

3. In Cluster Administrator, select the failover group containing the shared disk for Netscape Enterprise Server database and add
  - The IP failover object
  - The script failover object

4. To verify the configuration:

- a. Go to a client and open a web browser. In the browser, enter the web server's domain name or associated cluster IP address as the URL. For example:

`http://wbsrvr.dec.com`      or      `http://11.222.33.44`

You should see the default home page of the web server.

- b. In Cluster Administrator, manually fail over the web server's group and wait for the web server to migrate to the other cluster server. When the web server is running on the other cluster server, repeat step a to confirm that you can still access the web server.

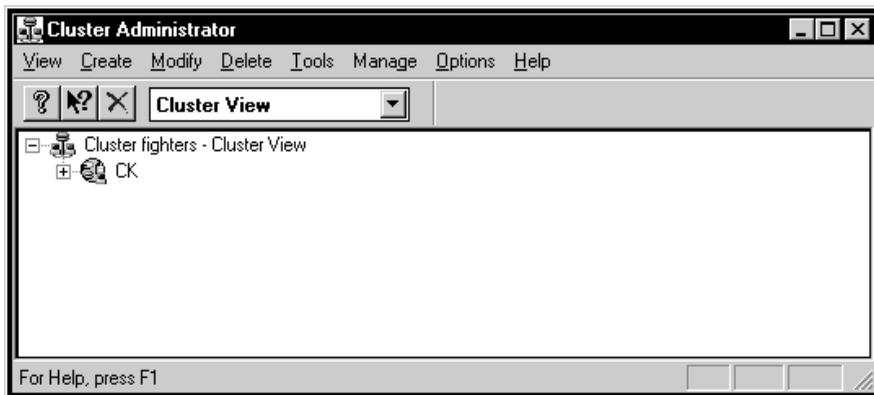
# 6

## Getting Started with Cluster Administrator

This chapter introduces Cluster Administrator. It gives step-by-step instructions on how to start and quit Cluster Administrator and how to view the cluster topology.

### Cluster Administrator Overview

Cluster Administrator is a centralized management graphical user interface for the DIGITAL Clusters for Windows NT environment. When you first start Cluster Administrator, the Cluster Administrator main window is displayed.



The main window contains a menu bar, toolbar, status bar, and shows the Cluster View by default.

The toolbar provides quick access to commands through shortcut buttons and allows you to change Views by selecting from a list.

The status bar gives a line of information related to each user-selectable action.

The toolbar and status bar are displayed by default. You can remove them from the screen by choosing the Toolbar and Status Bar options on the Options menu.

## Starting Cluster Administrator

► **To start Cluster Administrator:**

*On Windows NT 4.0:*

1. From the toolbar, click the Start button.
2. From the Program menu, choose DIGITAL Clusters for Windows NT.
3. From the Digital Clusters for Windows NT menu, choose Cluster Administrator.

The Cluster Administrator main window is displayed.

*On Windows NT 3.51:*

1. From Program Manager, choose the DIGITAL Clusters for Windows NT program group by double-clicking it. Or, select the DIGITAL Clusters for Windows NT program group and press Enter.
2. Choose the Cluster Administrator icon by double-clicking it. Or, select the Cluster Administrator icon and press Enter.

The Cluster Administrator main window is displayed.

## Quitting Cluster Administrator

► **To quit Cluster Administrator:**

*On Windows NT 4.0:*

1. Click the Close button in the upper right corner of the Cluster Administrator window. Or, from the View menu, choose Exit.

*On Windows NT 3.51:*

1. Double-click the Control-menu box or choose Close from the Control-menu box. Or, from the View menu, choose Exit.

## Displaying the Cluster Topology

Using Cluster Administrator, you can view an entire cluster in one of three perspectives:

- System View
- Cluster View
- Class View

You can use any of these views as the starting point for a cluster administration operation.

## Displaying the System View

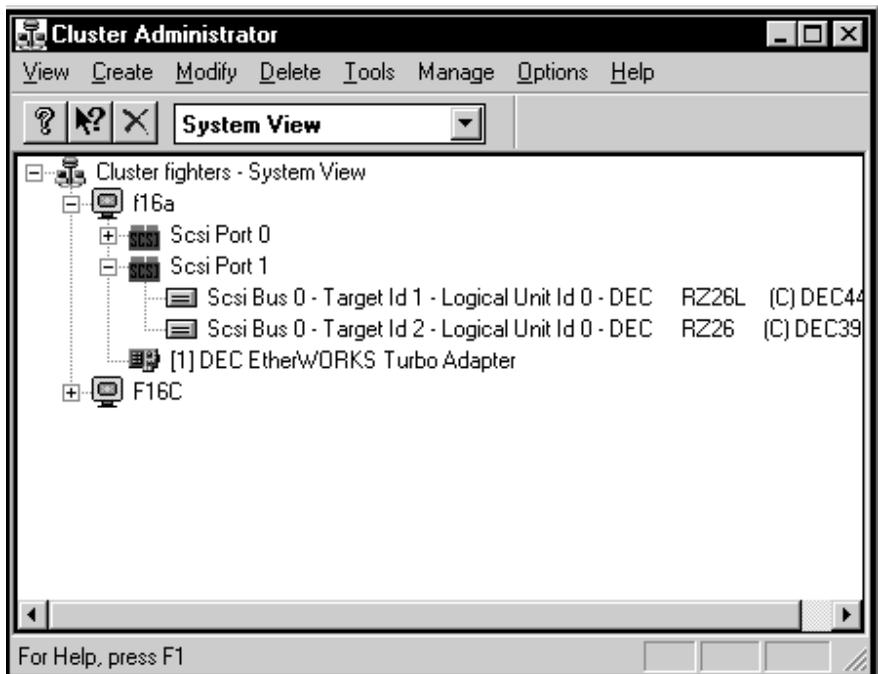
The System View allows you to see the cluster hardware. This view includes the cluster member system names, network adapters in each system, SCSI bus adapters in each system, and disks connected to each SCSI bus adapter.

You can use the System View as the starting point for a cluster administration operation. After determining which object you want to modify, choose an appropriate operation from one of the menus, for example, Manage SCSI Adapter Configuration, Manage Disk Alias, Manage Event Log, and so on.

### ► To display the System View:

1. From the toolbar, select System View from the list. Or, from the View menu, choose System.

The System View is displayed.



## Displaying the Cluster Topology

### Displaying the Cluster View

The Cluster View allows you to examine the cluster failover groups. This view includes all defined failover groups and members of each group. A failover group must contain at least one disk. It may also include recognized cluster applications such as Microsoft SQL Server and Oracle7 Workgroup Server, and scripts.

You can use the Cluster View as the starting point for a cluster administration operation by selecting an object to be modified. The default action for each cluster operation is predefined and can be invoked by double-clicking an object. Alternatively, you can select an object with a single click, then choose an appropriate operation from one of the menus.

Following is the default action for each object displayed in the Cluster View:

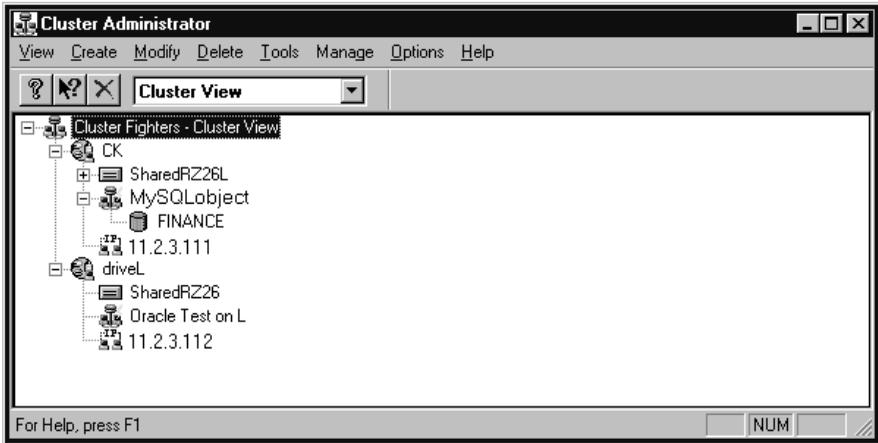
<b>Object Type</b>	<b>Default Action Dialog Box Display</b>
Group	Modify Failover Group
Disk	Manage Disk Alias
Oracle	Modify Oracle Failover Object
SQL Server	Modify SQL Server Failover Object
Script	Modify Script Failover Object
IP failover object	Modify IP Failover Object
SQL Server database	Manage SQL Server Databases

New in DIGITAL Clusters 1.1, the Cluster View displays all enrolled SQL Server databases for each shared disk. A green cylinder icon indicates that a SQL Server database is enrolled and ready for use with the cluster software.

► **To display the Cluster View:**

1. From the Toolbar, select Cluster View from the list. Or, from the View menu, choose Cluster.

The Cluster View is displayed.



## Displaying the Class View

The Class View allows you to look at the cluster from the perspective of available cluster objects without regard to physical location or failover grouping. This view, which presents all resources and services in a cluster's resource database, includes lists of the following:

- Groups
- Disks
- Shares
- Servers
- Oracle objects
- SQL Server objects
- Script objects
- IP failover objects
- SQL Server database objects

Oracle objects, SQL Server objects, and script objects are classified as Applications in the Class View.

## Displaying the Cluster Topology

You can use the Class View as the starting point for a cluster administration operation by selecting an object to be modified. The default action for each cluster operation is predefined and can be invoked by double-clicking an object. Alternatively, you can select an object with a single click and then choose an appropriate operation from one of the menus.

Following is the default action for each object displayed in the Class View:

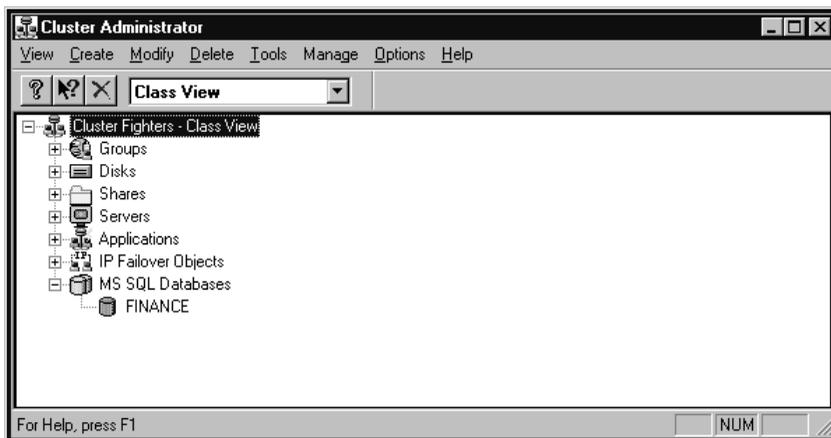
Object Type	Default Action Dialog Box Display
Group	Modify Failover Group
Disk	Manage Disk Alias
Server	None
Oracle	Modify Oracle Failover Object
SQL Server	Modify SQL Server Failover Object
Script	Modify Script Failover Object
IP failover object	Modify IP Failover Object
SQL Server database	Manage SQL Server Databases

New in DIGITAL Clusters 1.1, the Class View displays the SQL Server database class as an icon with two yellow cylinders. You can expand the SQL Server database class to view all enrolled and unenrolled SQL Server databases. A green cylinder indicates an enrolled database, while a yellow cylinder indicates an unenrolled database.

### ► To display the Class View:

1. From the Toolbar, select Class View from the list. Or, from the View menu, choose Class.

The Class View is displayed.



---

# Configuring Your Cluster

This chapter outlines the steps you need to perform to complete your initial cluster configuration.

## Configuration Steps

After you have installed the DIGITAL Clusters for Windows NT server software as instructed in the *Configuration and Installation Guide*, run Cluster Administrator to complete configuration of your cluster.

---

### Note

---

You must run Cluster Administrator from the account under which the cluster software was installed.

---

► **To complete configuration of your cluster:**

1. Start Cluster Administrator on one of the cluster servers. The Cluster View is shown by default. See the section Starting Cluster Administrator in Chapter 6 for details.
2. Display the Class View. The Class View allows you to look at the available cluster objects without regard to physical location or failover grouping. It is useful to know this information before you create a failover group. See the section Displaying the Class View in Chapter 6 for instructions

## Configuration Steps

3. If you plan to take advantage of the Microsoft SQL Server or Oracle7 Workgroup Server database failover features offered by the DIGITAL Clusters software, first you must complete the appropriate database software installation and configuration requirements outlined in Chapter 3.

Similarly, if you plan to set up Lotus Notes in a DIGITAL Clusters environment, first you must complete the installation and configuration requirements presented in Chapter 4.

And, if you are running DIGITAL Clusters 1.1 and want to configure a web server for IP failover, first you must perform the installation and configuration requirements discussed in Chapter 5.

4. If you have not already done so in step 3, create one or more new failover groups.

---

### Note

---

You can create all the failover groups from *either* cluster server, regardless of whether the current server system is the primary server for all the groups.

---

See the section Creating a Failover Group in Chapter 9 for instructions.

5. Display the Cluster View. The Cluster View allows you to examine the cluster failover groups. It includes all defined failover groups and members of each group. See the section Displaying the Cluster View in Chapter 6 for instructions.
6. Verify that the failover groups you created in step 4 are displayed in the Cluster View.
7. From the Tools menu, choose Disk Administrator. Use Disk Administrator to:
  - Verify that the shared storage for each cluster server belongs to the primary server for the failover group(s) you created in step 3 or 4. Disks not controlled by the server on which you run Disk Administrator will be designated as OFF-LINE.
  - Initiate a scan for new devices on the SCSI bus. If you add an additional disk to your storage tower, or turn on the power on a storage device that previously was turned off, the cluster software can detect this and can automatically incorporate the new device into the cluster's resource list. To initiate this action, you run Disk Administrator after adding the new storage device. Then, using Cluster Administrator, you add the device to a failover group to bring it online. See Chapters 6 and 9 for details.

- Determine which drive letters are associated with cluster disks.

---

**Note**

---

If you use Disk Administrator to change the drive letter assignments for partitions on shared storage, you *must* reboot the system for the change to take effect.

---

8. Optionally, you may want to create one or more network file shares on a cluster disk using Windows NT Explorer, File Manager, or the `net share` command. See the Windows NT documentation for details.

---

**Note**

---

You can create network file shares only from the server system that has the disk on line. Once created, the file shares automatically become cluster file shares. You do not need to take further action to make the shares known to the other cluster server.

---

---

## Managing a Cluster

This chapter presents step-by-step procedures on how to manage a cluster using Cluster Administrator.

The topics covered in this chapter include:

- Managing a SCSI adapter configuration
- Managing disk aliases
- Managing an event log
- Managing the log disk
- Managing manual failover
- Managing SQL Server databases

The general process for managing a cluster includes the following steps:

1. Create alias names for disks.
2. Create a failover group by specifying:
  - The contents of the group
  - The primary server and failover server for the group
  - Whether the group should failback to the primary server
3. Optionally, modify or delete failover groups as necessary.

See Chapter 9 for information about creating, modifying, or deleting failover objects and groups.

### Managing a SCSI Adapter Configuration

Use the Manage Adapter Configuration dialog box after adding, removing, changing, or rearranging any SCSI adapters in either of your cluster servers.

Note that when you use Cluster Administrator to manage adapter information, the server that you have selected must be running.

During installation of DIGITAL Clusters for Windows NT, the Windows NT operating system assigns logical port numbers to the cluster adapters. The operating system can change the logical port assignments if you:

- Change the physical placement of an adapter board in a server.
- Add or delete an adapter from the system.
- Add a new revision of a driver to the system.

For example, you might have an adapter in your cluster that, during installation, was assigned SCSI Port 1 and is the shared SCSI bus. If you add another SCSI controller, the first adapter could be renamed to SCSI Port 2. The cluster software detects a change but cannot detect which adapter is the shared bus, so does not use either as the shared bus. However, the cluster software indicates this by displaying the message when the system is rebooting:

```
Adapter Configuration has changed.
```

In addition, the cluster software includes a message in the event log, indicating that the cluster software cannot control any of the disks. Also, Cluster Administrator displays an error message, indicating that it cannot find information about the cluster disks.

---

#### Caution

---

By default, the cluster software is disabled after you change an adapter. When the system reboots, a warning message indicates that the adapter configuration has changed and that the cluster software is no longer controlling access to the disks.

In this state, the cluster software cannot protect files on shared disks from becoming corrupted if accessed by multiple users. Therefore, users should not access the cluster disks until the adapter configuration is updated.

---

Use Manage Adapter Configuration to check that the adapter configuration reflects the hardware configuration. Confirm that the correct adapters are being shared, then reboot the system. Or, if the wrong adapter is selected, select the correct adapter and then reboot the system.

If you remove a SCSI adapter from the configuration, any disk connected to that adapter is no longer available to be brought on line. Therefore, use Modify Group to remove those disks from any group that contains them.

---

### Note

---

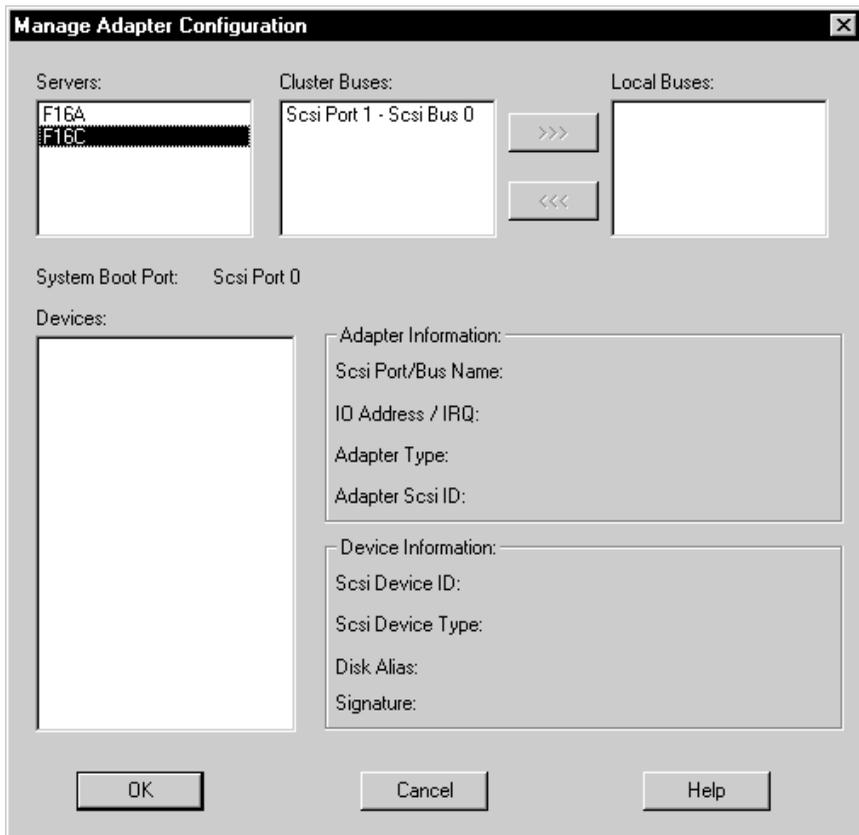
You can change the settings for only one server at a time. To change the settings for the other server, close the Manage Adapter Configuration dialog box and reopen it.

---

► **To manage a SCSI adapter configuration:**

1. From the Manage menu, choose Adapter Configuration.

The Manage Adapter Configuration dialog box is displayed.



## Managing a SCSI Adapter Configuration

In the dialog box, the Servers list box displays the computer names of the server systems in the cluster.

2. In the Servers list box, select a server for which you want to change the SCSI bus configuration. Cluster Administrator displays cluster bus information for each controller for the selected server in the following list boxes:
  - Cluster Buses—Displays the SCSI port ID and bus number for each controller that is currently configured as a clusterwide controller.
  - Local Buses—Displays the SCSI port ID and bus number for each controller that is currently configured as local to that system.
3. Select one of the controllers in either the Cluster Buses or Local Buses list box and use the arrow buttons between the list boxes to move the selected controller to the appropriate group.

When a controller is selected, Cluster Administrator displays read-only information about that device in the following informational (read only) boxes:

- System Boot Port text box—Displays the port ID of the SCSI controller that owns the system device (that is, the location where Windows NT is installed).
- Devices list box—Displays the SCSI target ID numbers for all the devices on the selected SCSI bus on the selected system.

For information about the SCSI device type and device ID or Signature for a selected device, select the device in the Devices list box. Cluster Administrator displays the information in the Device Information box.

- Adapter Information text box—Displays information about the selected SCSI adapter, including its SCSI port number and bus name, the I/O address and IRQ that it uses, and the Adapter Type (driver) and its assigned Server Adapter SCSI ID.
  - Device Information text box—Displays information about the selected SCSI target (device), including the device ID and type, the disk alias, as well as the signature of the device (disk) and the disk number as shown in Disk Administrator. Refer to your Windows NT documentation for information about the Disk Administrator.
4. When you have changed the adapter configuration, click either:
    - OK to save those changes
    - Cancel to cancel any changes before they are saved

The changes do not take effect until the system is rebooted.

Click the Help button at any time for information about this dialog box.

## Managing Disk Aliases

Use the Manage Disk Alias dialog box to create or change an alias name for a cluster disk. You can assign an alias to any disk, therefore giving it a more meaningful name than a computer-generated label.

By default, the cluster software always uses the SCSI address as the assigned disk alias when the software first discovers the disk.

---

### Note

---

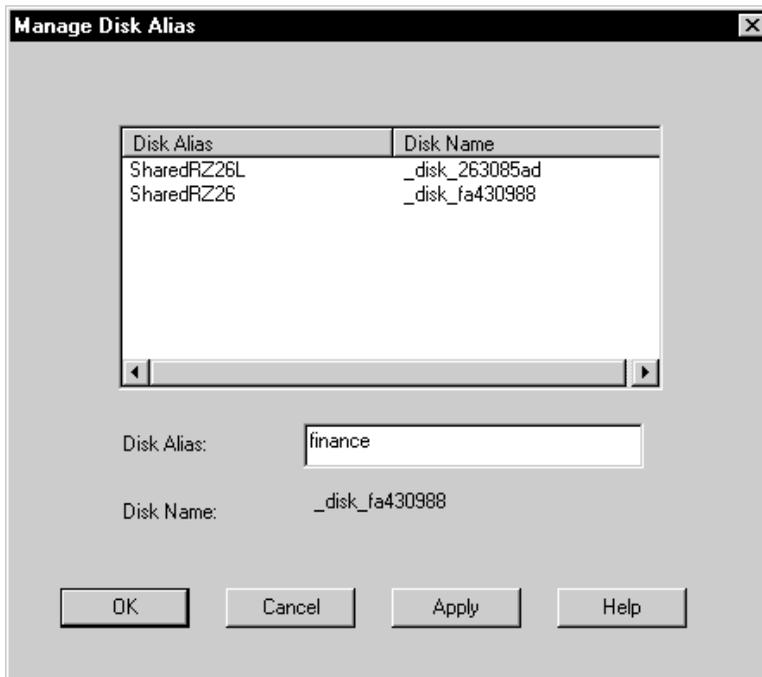
If you do not change the default disk alias, the cluster software attempts to give the alias a more meaningful value the first time the disk is brought on line, updating the disk alias with the partition volume label. Once either you or the cluster software sets the alias, the cluster software maintains the assigned alias if the volume label is changed.

---

► **To manage a cluster disk alias:**

1. From the Manage menu, choose Disk Alias.

The Manage Disk Alias dialog box is displayed.



## Managing an Event Log

The dialog box displays all disk names that are configured as part of a cluster (as compared to those disks configured as local). For each disk, the Disk Alias list box first displays the disk name automatically, and then as the disk comes on line, changes to the disk volume name.

2. Select one of the displayed disk names.
3. In the Disk Alias field, enter any unused name to represent the name of the selected disk.
4. Click OK to save the alias name, or Apply to commit the changes but leave the dialog box open.

Any view that displays the disk name will now display the alias name instead.

## Managing an Event Log

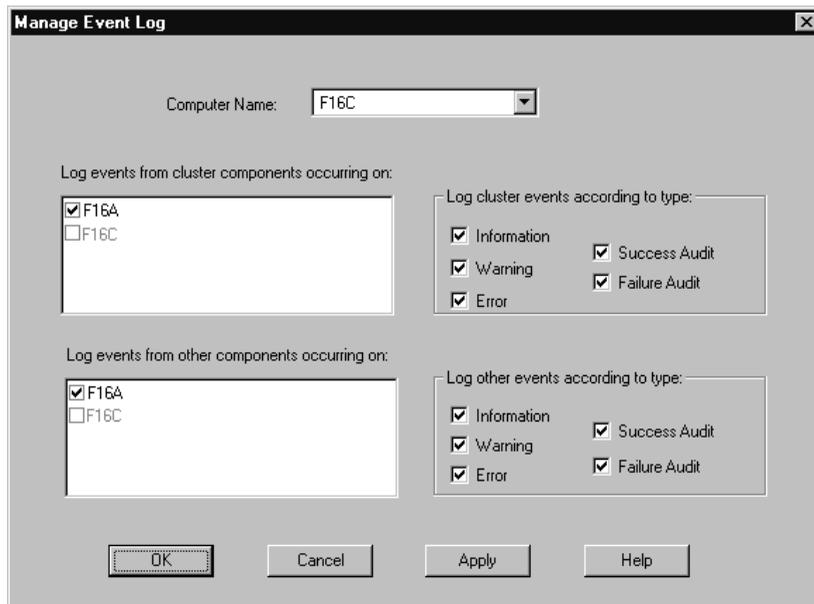
Many software components instruct the Windows NT event log service to log events. An event log can fill up quickly unless some limits are set on the type of events being logged.

By default, all type of events are logged to all members of a cluster. You can use the Manage Event Log dialog box to limit the types of events logged.

► **To manage an event log for the cluster:**

1. From the Manage menu, choose Event Log.

The Manage Event Log dialog box is displayed.



The selected computer name of the current server is displayed automatically in the Computer Name list box. You can manage event logging for that server, or select the other server to manage the event logging on the failover server.

The dialog box contains two sets of filters, which you can use to enable or disable logging of specific types of events for cluster events or for noncluster (other) events. By default, all the filters are enabled for logging.

These events correspond to Windows NT event types. For information about these event types, refer to your Windows NT documentation.

2. Specify which type of events get logged for the primary server, for:
  - Cluster components occurring on that server—In the filters for cluster log events, remove checkmarks for any of the types of log events that you do not want to include in the log. See the list following this procedure for information about cluster components.
  - Other (noncluster) components occurring on that server—In the filter for other log events, remove checkmarks for any of the types of log events that you do not want to include in the log.
3. Specify which type of events get logged for the failover server, for:
  - Cluster components occurring on that server—In the filters for cluster log events, remove checkmarks for any log events that you do not want to include in the log. See the list following this procedure for information about cluster components.
  - Other (noncluster) components occurring on that server—In the filter for other log events, remove checkmarks for any of the types of log events that you do not want to include in the log.
4. Choose one of the following:
  - OK to save the event log settings
  - Cancel to cancel any modifications to the settings
  - Apply to commit the changes but leave the dialog box open

## Managing the Log Disk

The following cluster components log events:

- Device drivers (including the DIGITAL Clusters Disk Driver, DIGITAL Clusters File Driver, and DIGITAL Clusters Port Driver)
- DIGITAL Clusters Failover Management Database Server (CFMD)
- DIGITAL Clusters Failover Manager
- DIGITAL Clusters Name Server
- DIGITAL Clusters Event Log Service

## Managing the Log Disk

The DIGITAL Clusters Failover Management Database (CFMD) uses the log disk to provide a place for the database to write its update log, thereby allowing the database to operate on one server in the absence of the other server. When the unavailable server becomes available again, this log is read, allowing the two servers to synchronize their databases.

The log disk is used whenever the Cluster Configuration Database needs to be updated while only one server is on line. For example, this situation can occur when Cluster Administrator is being used, or when the cluster software is bringing a group on line, and state information is updated in the database. In this case, the system that is restarting attempts to bring the log disk on line first so it can update the database.

For database updates to be accepted while only one server is available:

- The log disk must be a member of a group that is on line.
- The group that contains the log disk cannot be moved, deleted, or taken off line.
- The log disk cannot be changed.

---

### Note

---

Because the disk that contains the log must be both on line and cluster-available, make sure that the log disk is a member of a group that is on line. Cluster Administrator displays a warning message if the server that contains the log disk is off line, if you delete the group containing the log disk, or if you remove the log disk from the group while the remote server is down.

---

► **To manage the log disk:**

1. From the Manage menu, choose Log Disk.

The Manage Log Disk dialog box is displayed.



The Manage Log Disk dialog box contains:

- Log Disk text box—Displays the name of the disk that is currently the log disk.
- Available Disks list box—Displays the names of disks that are online to the same server as the current log disk. All shared disks are displayed if the log disk is not part of a failover group.

In general, you do not need to change the log disk. The log disk will function properly if part of a group. Only if you remove the disk designated as the log disk from the system should you assign a different log disk.

2. If you want the log to be located on the other disk, select that disk name and click OK (or double-click the selected disk name).

Note, however, that DIGITAL recommends that the disk containing the log be part of a failover group that is kept on line for database synchronization to function properly. In addition, to change to using another log disk, the old and new log disks must both be on line on the same system and both servers must be up.

## Managing Manual Failover

Use manual failover if you want to shut down a server to do routine maintenance but do not need the server to fail over when it shuts down. Or use manual failover to initiate a failover to test your cluster configuration.

After you use manual failover, the location of the group is not permanent. The group may move again. For example, if the group is placed manually on the failover server and the primary server goes on line, the group will return to the primary server.

## Managing Manual Failover

When you fail over a group manually, the failover occurs considerably faster than a failover caused by a server failure. During a manual failover, both servers take part in an orderly transfer of objects. In comparison, during a server failure, the remaining server must “sense” the failure of the failed server and obtain its objects.

---

### Note

---

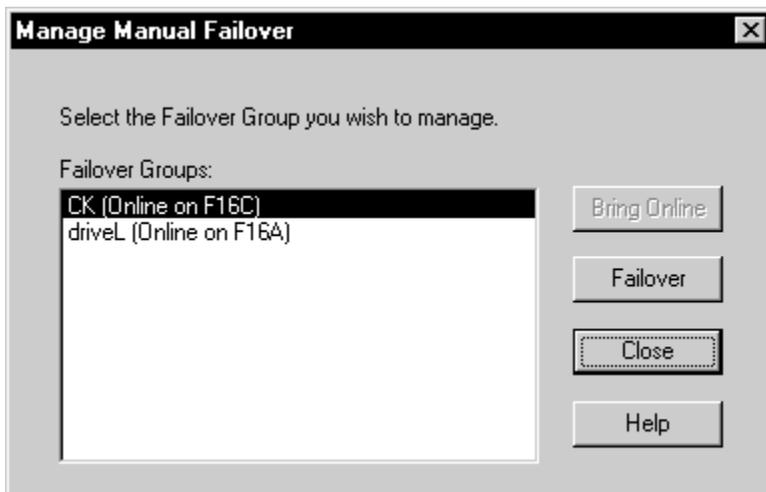
By rapidly and repeatedly failing over a group manually, you can exceed the repeated failover threshold in the DIGITAL Clusters Failover Manager. If a group fails over too many times, it is put off line to prevent the possibility of a hardware error on both servers causing an infinite failover.

---

#### ► To manage manual failover:

1. From the Manage menu, choose Manual Failover.

The Manage Manual Failover dialog box is displayed, with the first failover group is selected. by default.



2. In the Failover Groups list box, select a group and click Failover.

If the selected group is off line, it will come on line. If it is already on line on a server, the group will move to the other server.

## Managing SQL Server Databases

Use the Manage SQL Server Databases dialog box to configure your SQL Server databases in failover groups. You can:

- Enroll SQL Server databases as part of a failover group
- Unenroll SQL Server databases from a failover group
- Deactivate SQL Server databases

---

**Note**

---

To manage SQL Server databases, both cluster servers must be on line.

---

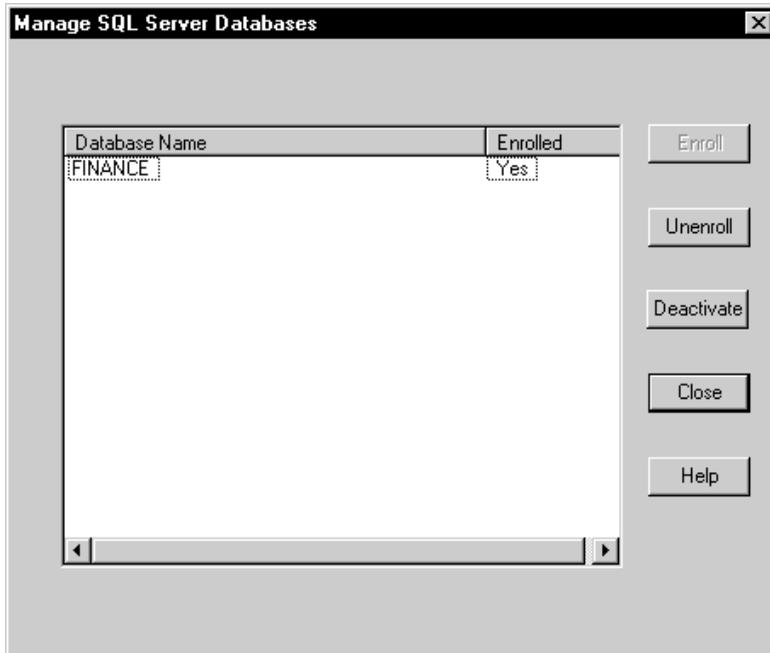
You must deactivate SQL databases before reconfiguring them in failover groups. For example, you may want to move a shared disk and its SQL Server databases from an existing failover group to a new group. For details, see Chapter 3. Deactivating a SQL Server database takes the database off line.

► **To enroll or unenroll an SQL Server database:**

1. From the Manage menu, choose SQL Server Databases. Or from the Class View or Cluster View, double-click an SQL Server database icon.

The Manage SQL Server Databases dialog box is displayed.

## Managing SQL Server Databases



The dialog box lists each SQL Server database found in the cluster, followed by “Yes” or “No” to indicate whether the database is currently enrolled as a member of an SQL Server failover group.

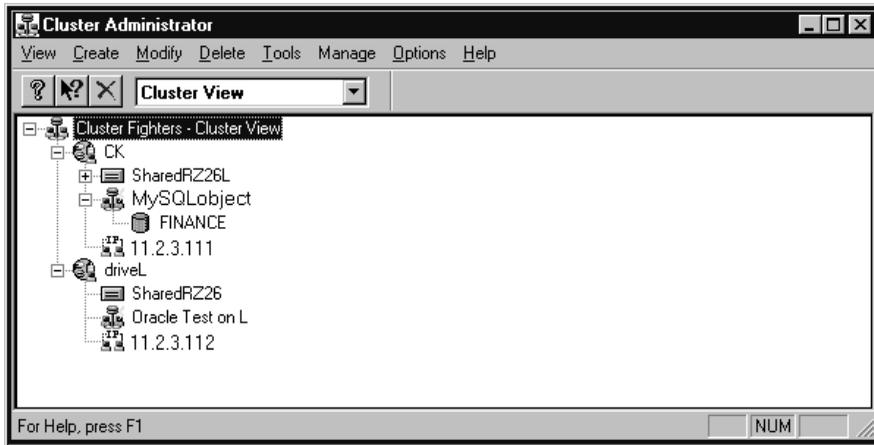
2. To enroll an available SQL Server database, select that database name and click the Enroll button. The selected SQL Server database is enrolled as a member of the group.

To unenroll an SQL Server database that is already a member of the group, select that database and click the Unenroll button. The selected SQL Server database is unenrolled as a member of the group.

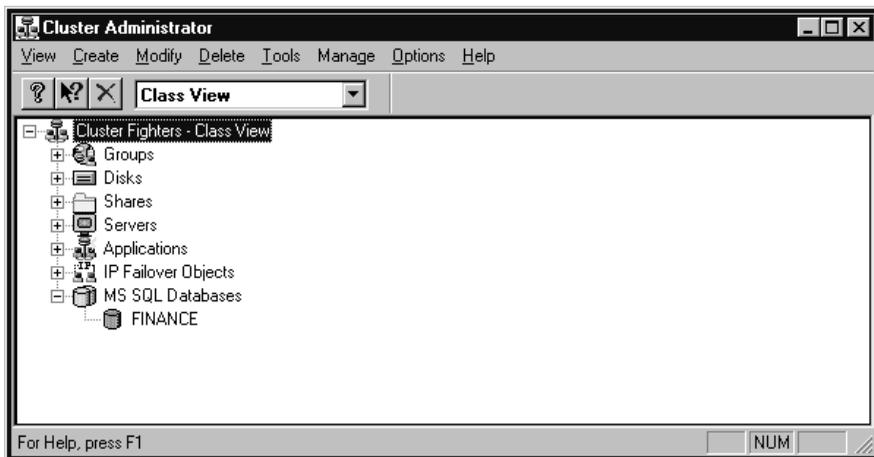
3. Click Close to close the dialog box.

You can also view the list of SQL Server databases by selecting the Cluster View or Class View in Cluster Administrator.

The Cluster View displays all enrolled SQL databases for each shared disk. A green cylinder icon indicates that a SQL database is enrolled and ready for use with cluster software.



In the Class View, two yellow cylinders represent the SQL Server database class. You can expand the database list to view all enrolled and unenrolled SQL Server databases. A green cylinder represents an enrolled database, and a yellow cylinder represents an unenrolled database.



► **To deactivate an SQL Server database:**

1. From the Manage menu, choose SQL Server Databases. Or from the Class View or Cluster View, double-click an SQL Server database icon.
2. Select the database you want to deactivate, and click the Deactivate button. The selected database is placed off line so you can reconfigure the failover group.

---

# Working with Failover Objects and Groups

This chapter describes how to create, modify, and delete failover objects and groups. The topics covered include:

- Working with an Oracle failover object
- Working with a script failover object
- Working with an SQL Server failover object
- Working with an IP failover object
- Working with a failover group

## Working with an Oracle Failover Object

An Oracle failover object allows you to provide failover for an Oracle server database through the cluster software.

When you use an Oracle server database, you need to supply Cluster Administrator with information about the database so that it can fail over the database group correctly. Use the Create Oracle failover object dialog box to supply information for Cluster Administrator to start Oracle server instances. You can create, modify, or delete an Oracle failover object with the Create, Modify, or Delete menus on the toolbar.

For an Oracle server instance to fail over correctly, the database and all the data files associated with the server instance must be located on the shared, clusterwide disk.

For more information about Oracle server failover, see Chapter 3.

### Creating an Oracle Failover Object

Create an Oracle failover object when you need to add the object to the cluster database.

---

#### Note

---

After you create a failover object as described in this section, it is available for adding to a failover group. Remember to add the object to a group by selecting Failover Group from the Create menu and moving the object from the Available Failover Objects list box into the Group Contents list box. You also can use Modify Failover Group to add the object to an existing group.

---

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

1. From the Create menu, select Oracle Failover Object.

The Create Oracle Failover Object dialog box is displayed.

The screenshot shows a dialog box titled "Create Oracle Failover Object" with the following fields and values:

Oracle Failover Object Name:	Oracle1
Comment:	Oracle failover database
Oracle Username:	internal
Oracle Password:	*****
Oracle User Type:	Normal
Oracle sid:	orcl
Oracle Pipe Name:	ORAPIFE
Oracle sid Parameter File:	H:\orant\database\initorcl.ora
Oracle Database:	Oracle
Oracle Listener Name:	Listener1

Buttons: OK, Cancel, Help

2. Supply Cluster Administrator with information about your Oracle database in the following text boxes:

- Oracle Failover Object Name—Specifies the name of the failover object. Use this name when you add this Oracle object to a failover group.
- Comment—Specifies a comment, optionally, that is displayed when this Oracle object is advertised by the cluster name service software.
- Oracle Username—Specifies the user name for the Oracle account associated with this instance of the database. Note that the Oracle user name has a length limit of 30 characters.

The user account name must be able to start and stop the Oracle database. For Oracle7 Server Version 7.1, use the default name “internal” because this is the only name that can start the database. For Oracle7 Server Version 7.2 or 7.3, you can use any account name that can start and stop the Oracle database.

- Oracle Password—Specifies the password for the Oracle account, which is used to protect access to an instance of the database. Note that the Oracle password has a length limit of 30 characters.
- Oracle User Type—Specifies one of three user types supported by Cluster Administrator: Normal (the default), Sysdba, or Sysop.
- Oracle sid—Specifies the Oracle system identifier (sid) that identifies an instance of a particular Oracle database server.
- Oracle Pipe Name—This optional field specifies:
  - For Oracle7 Server Version 7.1, the Oracle system identifier as the pipe name.
  - For Oracle7 Server Version 7.2 or 7.3, for each instance, the pipe name of the network listener associated with that instance. For example, use ORAPIPE for the default network listener.

You can check your Oracle7 Server Version 7.2 or 7.3 pipe name by using the SQL\*NET™ configuration tool. See your Oracle network configuration documentation for more information.

- Oracle sid Parameter File—Specifies the full path to the Oracle sid parameter file. The parameter file contains parameters that are applied to an instance of a database server when the instance is started.
- Oracle Database—Specifies the name of the Oracle database to be associated with the failover instance.
- Oracle Listener Name—Specifies the name included in the `listener.ora` file that identifies the Oracle TNS listener associated with this Oracle database. If you have more than one listener on a server, each listener requires a unique name.

## Working with an Oracle Failover Object

If you use IP failover, this field is required for the listener to be stopped or started as part of failover. The unique listener name ensures that only one listener (the one specifically for the database that is failing over) is running on the cluster.

If you use named pipes, this field is optional but recommended. If you leave the field empty, the failover object does not attempt to stop or start the listener service.

Valid characters for this field include: [a...z], [A...Z], and [0...9].

3. Click OK to save the settings, or click Cancel to cancel the modifications.

## Modifying an Oracle Failover Object

Modify an Oracle failover object whenever you need to change any of the information associated with that object. For example, if you change the user ID or the password for the Oracle database, you need to update this information for Cluster Administrator.

### ► To modify an Oracle failover object:

1. In the Class or Cluster View, select the Oracle failover object that you want to modify.
2. From the Modify menu, choose Oracle Failover Object.

The Modify Oracle Failover Object dialog box is displayed.

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

Field	Value
Oracle Failover Object Name	Oracle1
Comment	Oracle failover database
Oracle Username	internal
Oracle Password	*****
Oracle User Type	Normal
Oracle sid	ord
Oracle Pipe Name	ORAPIPE
Oracle sid Parameter File	H:\orant\database\initord.ora
Oracle Database	Oracle
Oracle Listener Name	Listener1

Buttons: OK, Cancel, Help

The selected Oracle failover object is displayed in the Oracle Failover Object Name text box.

By default, Cluster Administrator displays associated information in whichever fields you specified when creating or last modifying the object.

3. Update Cluster Administrator with information about changes in your Oracle database by modifying the text boxes in the dialog box.
4. Click OK to save the settings, or click Cancel to cancel the modifications.

## Working with a Script Failover Object

### Deleting an Oracle Failover Object

You can delete an Oracle failover object from a group if you no longer need that object in the cluster database.

► **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 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.

### Working with a Script Failover Object

Cluster Administrator lets you create a script failover object to run NT command scripts when a failover group comes on line or goes off line. Script failover objects are most commonly used to start and stop applications that are not directly supported by Cluster software.

One common use of a script failover object is to start or stop a third-party database application that uses data on a shared cluster disk. For example, when a shared disk is put on line on a cluster server, you may want to start a database application to read and write data from the shared disk. Or if you move the shared disk from one server to the other, you could use a script failover object to stop the database on one server and start it on the other.

Another common use of a script failover object is to notify a system administrator of a failover event. To do this, you can include `net send` commands in your failover scripts.

DIGITAL Clusters 1.1 adds support for IP failover objects, so you can use a script failover object to start or stop a web server or other socket-based application. See the section Working with an IP Failover Object (page 9–11) for more information.

The following section describes how to create a script failover object with Cluster Administrator. After you create the object, you can add it to a failover group. Use either the Create Failover Group wizard to create a new group that will contain the object, or the Modify Failover Group dialog box to place it in an existing group.

### Creating a Script Failover Object

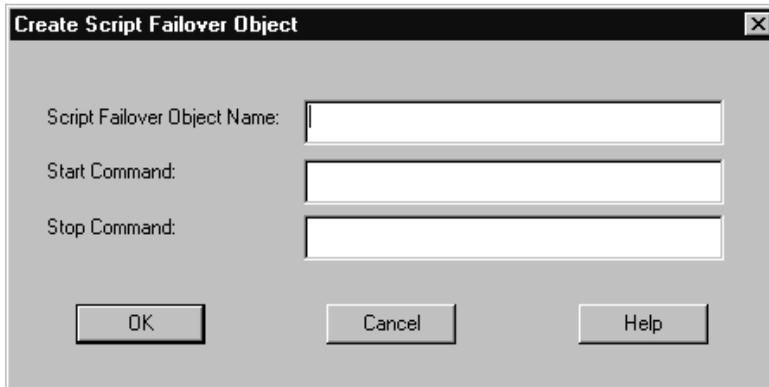
When you create a script failover object, you specify a start and stop command to run when the object's failover group comes on line or goes off line. You can specify a single command or use command files.

After you create a failover object, it becomes an available cluster resource that you can add to a failover group. Remember to add the resource to the desired failover group by using the Create Failover Group or Modify Failover Group dialog box and moving the object from Available Failover Objects to the Group Contents.

► **To create a script failover object:**

1. From the Create menu, choose Script Failover Object.

The Create Script Failover Object dialog box is displayed.



2. In the Script Failover Object text box, enter the object name associated with the specific script. The script object name can be any name that is different from other script objects.

3. In the Start Command text box, enter the script command to run when the DIGITAL Clusters Failover Manager brings the group on line.

The start command runs when the script failover object comes on line as part of a failover group. The start command can be any command recognized by the default Windows NT command interpreter (`cmd.exe`). You can use the name of a program, or a batch (`.bat`) or command (`.cmd`) file.

4. In the Stop Command text box, enter the script name to run when the DIGITAL Clusters Failover Manager takes the associated failover group off line.

The stop command runs when the script failover object goes off line as part of a failover group. The stop command follows the same rules as the start command.

Start and stop commands are optional. If you leave the command field empty, no commands are performed when the failover group comes on line or goes off line.

5. Click OK to accept the setting.

You can now add the object to a new or existing failover group.

### **Script Failover Object Command Restrictions**

Restrictions on the script failover object start and stop commands include the following:

## Working with a Script Failover Object

- Access rights for the commands and the account must match – The start or stop command runs with the same access rights as the DIGITAL Clusters Failover Manager account specified during the clusters software installation. If you do not provide a path for the commands, the cluster software uses the default path established for the DIGITAL Clusters Failover Manager account.

The script must be on storage that is accessible to the DIGITAL Clusters Failover Manager. In general, do not place the script on a network share, as network shares are unavailable to the DIGITAL Clusters Failover Manager. DIGITAL recommends that you place copies of the script on a local disk on each server in the cluster.

- Commands cannot interact with the desktop –The start and stop commands cannot display information on the screen of the server. However, a script failover object can do simple notifications by using the `net send` command to cause a pop-up notification to be displayed on any system. For example, a script failover object often is used to start an NT Service (using the `net start` command) or to start a detached process (for example, `start database_server.exe`).
- Commands should return control quickly –If the start or stop command takes time to execute, for example if it waits for user input, the failover group will fail to come on line or to go off line while the command delays. Therefore, do not use commands in your script that wait for user input, such as `pause` or `date`, or that try to start a program that displays information on the screen, such as `start winfile.exe`.

If a program started by a script failover object takes time to process (or never finishes processing), start it by using the `start` command. For example, use the command `start database_server.exe` rather than `database_server.exe`.

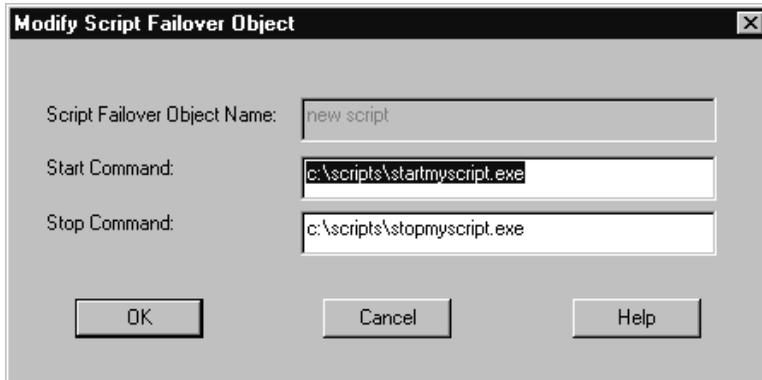
## Modifying a Script Failover Object

Modify the characteristics of a script failover object if you make changes to that object. For example, you might change the name or location of the script file.

### ► To modify a script failover object:

1. From any view, select a Script failover object, and then choose Script Failover Object from the Modify menu or double-click on a script object in a Class View. (Script objects are displayed under the Applications icon.)

The Modify Script Failover Object dialog box is displayed.



2. Edit any of the information in the text boxes to reflect the changes made to the script failover object.
3. Click OK to accept the changes, or Cancel to cancel the changes.

### Deleting a Script Failover Object

You might need to delete a script failover object if you no longer need that object in a group, or if you are reorganizing your failover groups.

► **To delete a script failover object:**

1. In a Class or Cluster View, select the script that you want to delete.
2. From the Delete menu, choose Script Failover Object, or click the Delete button.

Cluster Administrator deletes the selected script from the cluster database.

## Working with SQL Server Failover Objects and Groups

If you use a Microsoft SQL Server, you need to supply Cluster Administrator with information about the database server so that, should there be a cluster failure, Cluster Administrator correctly handles the failover group that contains the database server.

### Creating an SQL Server Failover Object

To ensure that Cluster Administrator can properly manage the failover of Microsoft SQL Server databases that reside on a shared disk, Cluster Administrator automatically creates an initial cluster SQL Server failover object which contains login information for Microsoft SQL Server and defines the password for the SQL `sa` account. The DIGITAL Clusters 1.1 software uses the password to connect to the Microsoft SQL Server process on the cluster servers.

You never need to create SQL Server failover objects manually because they are created automatically, as required, when you enroll SQL Server databases.

### Modifying an SQL Server Failover Object

Modify the characteristics of an SQL Server failover object whenever a database administrator changes the SQL system administrator password. Doing so ensures that the DIGITAL Clusters Failover Manager uses the updated password to connect to the Microsoft SQL Server process.

You need to modify only one of the SQL Server failover objects to record the new password.

► **To modify an SQL Server object:**

1. From the Class View, select the SQL Server object that you want to modify.
2. From the Modify menu, choose SQL Server Failover Object.

The Modify SQL Server Failover Object dialog box is displayed.



The screenshot shows a dialog box titled "Modify SQL Server Failover Object". It has a standard Windows-style title bar with a close button (X) in the top right corner. The dialog contains four text input fields arranged vertically. The first field is labeled "Name:" and contains the text "SQLfailoverobj". The second field is labeled "Comment:" and contains the text "This is an SQL failover object". The third field is labeled "Username:" and contains the text "sa". The fourth field is labeled "Password:" and is currently empty. At the bottom of the dialog, there are three buttons: "OK", "Cancel", and "Help", each enclosed in a rectangular box.

In the SQL Server Failover Object Name text box, the name of the selected SQL Server failover object is displayed. If you supplied a comment when first creating the SQL Server failover object, the comment is displayed also.

3. In the Comment text box, you can add an optional comment or change the comment that you provided previously. This comment is displayed when this SQL Server failover object is advertised via the cluster name service software.
4. In the Password text box, you can change the password for the SQL Server sa account, used for protecting access to an instance of the database.
5. Click OK to save the modifications to the settings, or click Cancel to cancel the modifications.

## Deleting an SQL Server Failover Object

You can delete an SQL Server failover object whenever you need to reconfigure your failover group.

► **To delete an SQL Server failover object:**

1. From the Class View, select the SQL Server object that you want to delete.
2. From the toolbar, select Delete, SQL Server Failover Object.

The SQL Server object is deleted from the cluster database, making it unavailable to be added to another group.

---

### Note

---

When you delete a group, the objects that were in the group continue to exist—you are not deleting the objects but only the group that contains them. Deleting the group frees the objects for adding to other groups.

In comparison, when you delete an object, it is deleted from the cluster database and so is no longer available to be added to a group.

---

## Working with an IP Failover Object

DIGITAL Clusters 1.1 adds support for failing over IP socket-based applications, such as web servers. As part of configuring an IP socket-based application for failover, you must create an IP failover object. When you create an IP failover object, you specify an IP address and one or more network adapters on each cluster server that can potentially enable the address. You use this address as a cluster IP address that can migrate from one server to another.

You can place the object in a failover group, along with a shared disk and selected IP socket-based applications on the disk that you want to use the IP address. Clients can access the applications using the cluster IP address. When the primary server for the group fails, the cluster IP address and associated applications migrate to the secondary server.

---

### Note

---

IP failover objects require the TCP/IP protocol and DIGITAL Clusters 1.1 running on both servers of the cluster. Otherwise, Cluster Administrator disables IP failover.

---

## Working with an IP Failover Object

An IP failover group is particularly useful for ensuring failover of IP socket-based applications such as a web server or Lotus Domino 4.5 server. For example, DIGITAL Clusters 1.1 supports failover of various web servers, including the Microsoft Internet Information Server (IIS) and the Netscape Enterprise Server.

You must install the web server software on both cluster servers for IP failover to work correctly. However, the web server runs only on one cluster server at a time. See Chapter 5 for details on setting up web servers for failover.

In general, to establish failover for an IP socket-based application, follow these steps:

1. Create an IP failover object, as described in the next section.
2. Create a failover group and add the IP failover object to the group, as described in the section *Creating a Failover Group*.
3. Add a script failover object to the group for starting and stopping the IP socket-based application. For more information, see the section *Creating a Script Failover Object* (page 9–6).

### Creating an IP Failover Object

This section describes the process for creating an IP failover object. When you create an IP failover object, you specify an IP address to use as a cluster IP address. A cluster IP address can migrate from one cluster server to another after a failover. Make sure the IP address is not currently in use on the network. Do not use the IP address of an individual cluster server or any other system.

You also specify a subnet mask and one or more network adapters on each cluster server that can activate the cluster IP address. After you create the IP failover object, add the object to the desired failover group. See the section *Creating a Failover Group* (page 9–16) or *Modifying a Failover Group* (page 9–20) for information.

---

#### Note

---

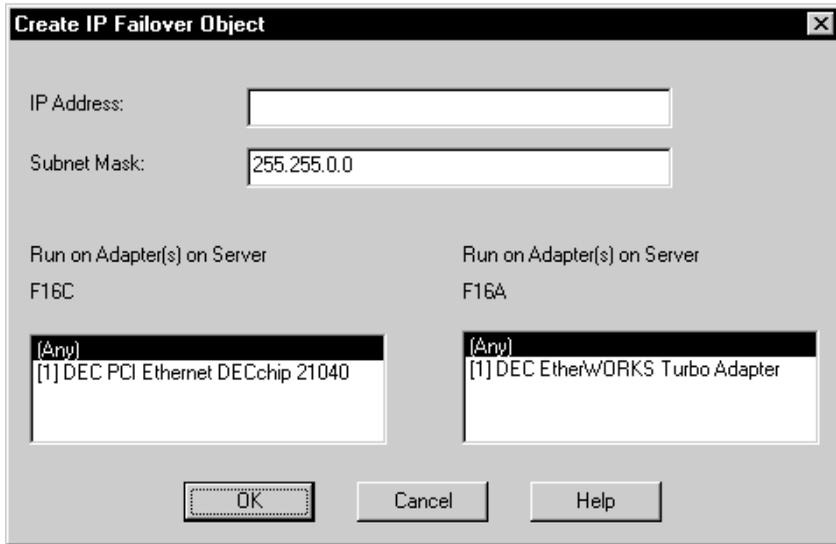
Place the IP failover object after the shared disks in the failover group.

---

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

1. From the toolbar, select *Create, IP Failover Object*.

The Create IP Failover Object dialog box is displayed.



2. In the IP Address text box, enter the IP address to use as a cluster IP address. Use this address when adding this IP object to a failover group.
3. In the Subnet Mask field, supply the subnet mask associated with the IP address.
4. Specify which network adapters on each cluster server can enable the IP address.

A list box for each server displays the list of available network adapters. The default value is Any, which allows the IP address to fail over to any adapter when a failover occurs.

Select a specific adapter only if needed for your cluster configuration. For example, if one network adapter belongs to a private local area network, you should select a different network adapter for failover.

5. Click OK to save the modifications to the settings, or click Cancel to cancel the settings.

## Modifying an IP Failover Object

If needed, you can modify an IP failover object to change the IP subnet mask or set of network adapters permitted to enable the specified cluster IP address. To change a cluster IP address, delete the current IP failover object and create a new object.

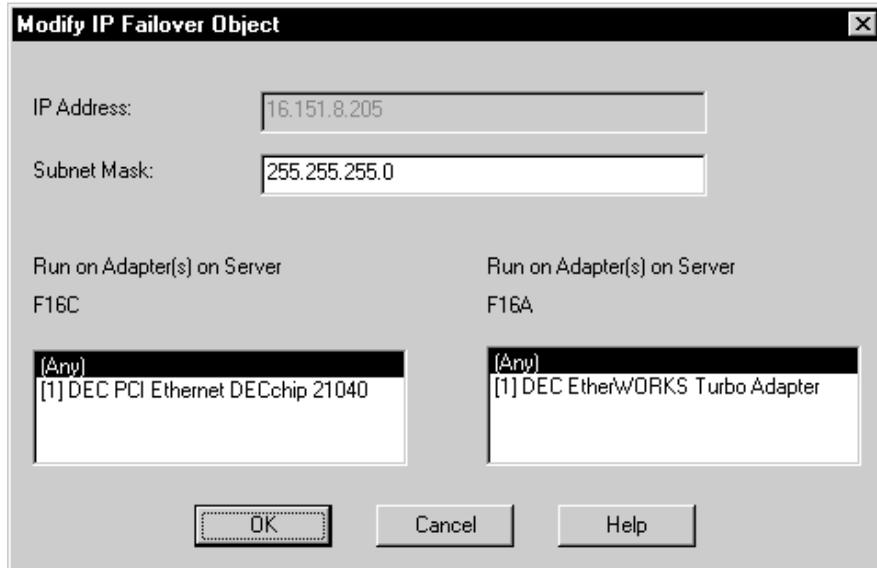
### ► To modify an IP failover object:

1. From the Class View, select the IP failover object you want to modify.

## Working with an IP Failover Object

- From the toolbar, select Modify, IP Failover Object.

The Modify IP Failover Object dialog box is displayed.



- If needed, change the subnet mask associated with the cluster IP address.
- If needed, select different network adapters that can enable the cluster IP address on the cluster servers.

The list box for each server highlights the current adapters selected. The default value when creating a new IP failover object is Any, which allows the address to fail over to any adapter on the server when a failover occurs.

Select a specific adapter only if needed for your cluster configuration. For example, if one network adapter belongs to a private local area network, you should select a different network adapter for failover.

- Click OK to save the modifications to the settings, or click Cancel to cancel the changes to the settings.

## Deleting an IP Failover Object

You can delete an IP failover object if you no longer need to use the cluster IP address specified by the object. Make sure that there are no IP socket-based applications still using the cluster IP address.

► **To delete an IP object:**

1. From the Class View, select the IP object that you want to delete.



2. From the toolbar, select Delete, IP Failover Object.

The IP object is deleted from the cluster database, making it unavailable to be added to another group.

---

### Note

When you delete a group, the objects that were in the group continue to exist—you are not deleting the objects but only the group that contains them. Deleting the group frees the objects for adding to other groups.

In comparison, when you delete an object, it is deleted from the cluster database and is no longer available to be added to a group.

---

### Working with a Failover Group

Using a failover group ensures that Cluster Administrator fails over objects together. All of the objects that you place in a failover group are treated as a unit.

A group is particularly useful when a dependency exists among the members of the group. For example, you might want to place a disk and an SQL Server database in the same group. In general, you increase cluster reliability and performance by creating many groups that contain small numbers of failover objects rather than one or two groups that contain numerous failover objects.

Because all the objects in a group are treated as a unit, all objects in a group must be on line for any of them to be on line. If any member of the group goes off line, for example because of a hardware error, all members of the group are put off line.

An object in a failover group can be on line only on one server at a time. Note that this is true for all types of failover objects, including script failover objects, even though a group does not own a script in the same way it might exclusively own a piece of hardware. A failover object cannot be shared by multiple groups.

A newly created group comes on line immediately, and modifications to a group, as made with the Modify Failover Group dialog box, take effect immediately. A deleted group goes off line immediately.

### Creating a Failover Group

Create a generic failover group to ensure that Cluster Administrator fails over a specified group of objects together. After creating the group, add to the group the specific failover objects that you want to be failed over together.

► **To create a failover group:**

1. From the Create menu, choose Failover Group. The Create Failover Group wizard prompts you to create a new group by clicking the Next button. When you click Next, the Contents dialog box is displayed.

The Contents page of the Create Failover Group contains:

- Group Name field—By default, the Group Name field is blank. Use this field to assign a name for a group of objects that you want to fail over together.
- Group Contents list box—Contains no entries because you have not yet assigned entries to the new group.
- Available Failover Objects list box—Displays all objects available for adding to the selected group (an object is available if it is not included in another failover group).

---

#### Note

---

If an object is available, Cluster Administrator lists it in the Available Failover Objects list box, and the group contents are listed in the Group Contents list box. If no failover objects are available, nothing is listed. In this case, you either need to create some failover objects or delete a group to free objects for inclusion in a different group. See the section [Deleting a Failover Group](#) (page 9–23) for information about deleting a group.

---

2. In the Group Name field, enter a new name for a failover group.

## Working with a Failover Group

3. In the Available Failover Objects list box, select any object that you want to include in this new failover group, and use the arrow between the boxes to move it into the Group Contents list box.

---

### Note

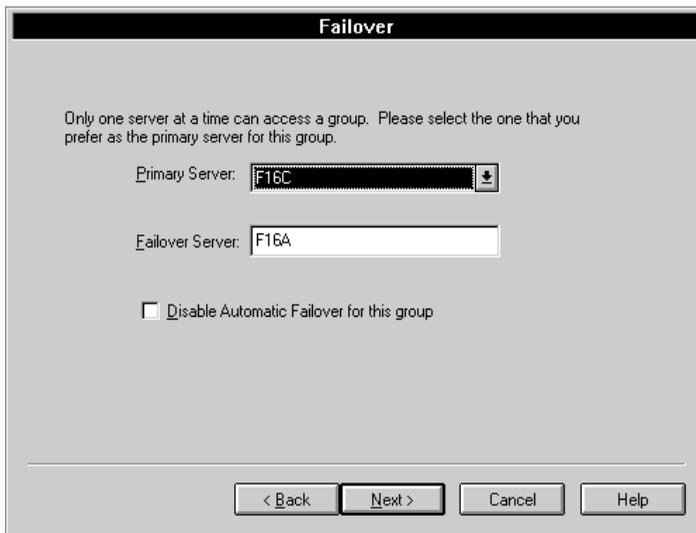
---

When adding objects to a failover group, you should list them in their starting order, as lower-level services must be started before higher-level services when a group comes on line. Therefore, add objects in the following order: disks first, followed by IP addresses, and then databases. The objects will fail over in the order they are listed.

---

4. Click the Next button.

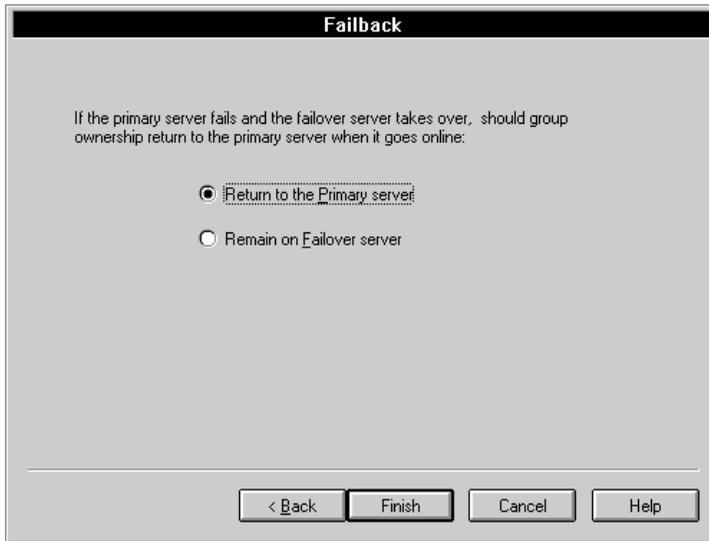
The Failover dialog box is displayed.



Use this dialog box to specify which server will be the primary server. The other server defaults to being the failover server automatically.

5. If you need to revise how you grouped your failover objects, click the Back button to return to the Contents dialog box, move the objects, then click Next to return to this dialog box. Otherwise, click Next to continue to the final step, or click Cancel to cancel any changes before they are saved.

The Failback dialog box is displayed.



6. On the Failback dialog box, specify which server should regain control of the group after a failure. Specify that the group should either:
  - Return to the Primary server — Select this radio button to specify that the group, when on line to the failover or backup server, should detect when the primary host becomes available. Upon detecting that the primary server is becoming available, the group returns to the primary server automatically.

You can provide a measure of load balancing by enabling the Return to the Primary server function, as it allows the cluster software to reestablish the static load assignments made to the servers when a second server comes on line. However, enabling this function can allow an uncontrolled failover of a failover object, returning it to its primary server while the object is in use.

---

### Caution

---

The Return to the Primary server function can allow the dismounting of a disk, regardless of the state of the disk, while a user is accessing data files on the disk. In that case, the data is vulnerable to data loss or corruption. Therefore, to redistribute objects after recovery from a power outage, enable the Remain on Failover server function and use the manual failover function.

---

## Working with a Failover Group

- Remain on Failover server — Select this radio button to specify that, when the primary server becomes available again, the group should remain on line to the failover server instead of returning to the primary server.

7. Select Finish to save the group definition.

## Modifying a Failover Group

You can modify the contents of a failover group to reorganize the objects included in the group for failover.

### ► To modify a failover group:

1. In a Cluster or Class View, select the group name you want to modify.
2. From the Modify menu, choose Failover Group.

---

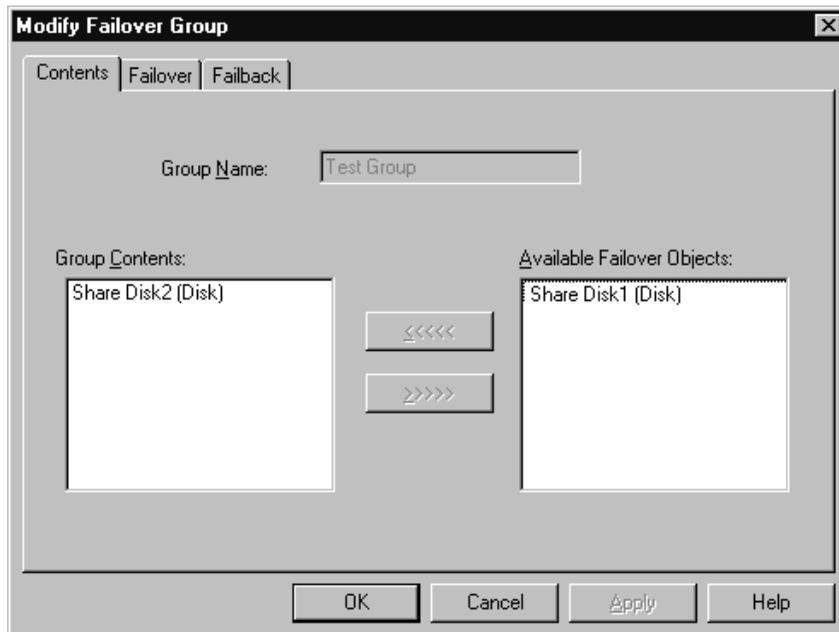
### Note

---

You can double-click on the group object to access the Modify Failover Group dialog box.

---

The Modify Failover Group tabbed dialog box is displayed. By default, the Contents page of this dialog box is displayed.



The Contents page of the Modify Failover Group dialog box contains:

- Group Name field—Displays the name of the group that you selected.
  - Group Contents list box—Displays the objects currently contained in the selected group.
  - Available Failover Objects list box—Displays any objects available for adding to the selected group.
3. In the Group Contents list box, you can select any object and remove that object from the group, so that object will no longer be failed over with the remaining objects in that group. Alternatively, you can select any object listed as available and move that into the failover group.
  4. To modify the failover behavior of the group, click on the Failover tab of the Modify Failover Group dialog box.

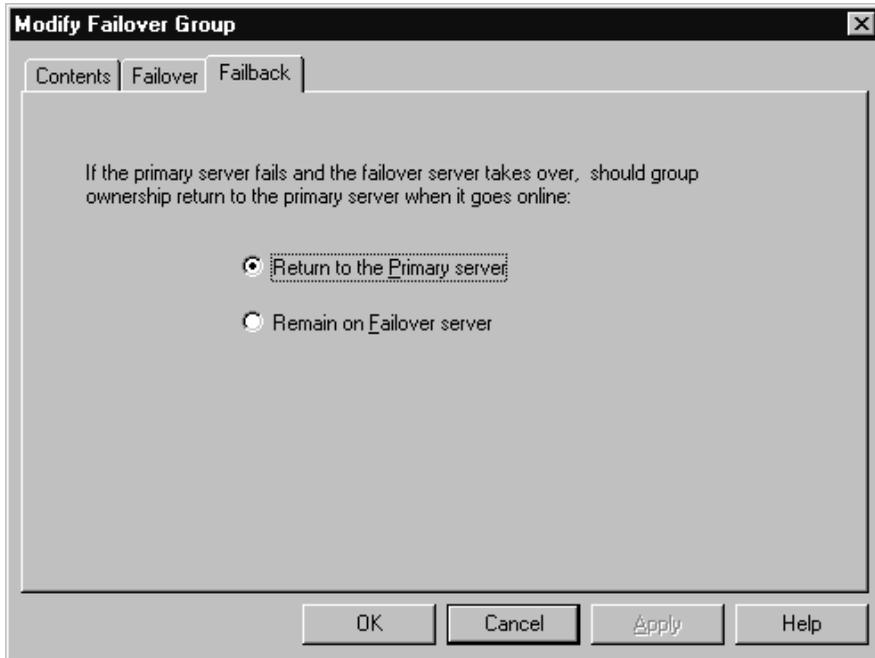


If desired, change the selected primary server for the group. The primary server is the cluster server where the group is available normally:

- Primary Server field—Select the computer name of the cluster server to act as the primary server.

## Working with a Failover Group

- Failover Server field (read only)—Displays the computer name of the server where the group will be available if the primary server fails. Because there are only two servers in a cluster, the failover server defaults to the computer not specified as the primary server.
  - Disable Failover for this group checkbox—If desired, you can disable failover for this group. To enable failover again, clear the checkbox. Disabling failover allows the group to remain on line but not move to the other server when a cluster failure occurs. Use this function to allow a group to stay on a server so that you can take down the other server without disrupting services—for example, when backing up a system or doing routine maintenance.
5. To modify the failback behavior of the group, click on the Failback tab of the Modify Failover Group dialog box.



Use this page to modify the failback behavior of a group. The page includes two radio buttons:

- **Return to the Primary server**—Select this radio button to specify that the group, when on line to the failover or backup server, should detect when the primary server becomes available. Upon detecting that the primary server is becoming available, the group returns to the primary server automatically.

You can provide a measure of load balancing by enabling the Return to the Primary server function, as it allows the cluster software to reestablish the static load assignments made to the servers when a second server comes on line. However, enabling this function can allow an uncontrolled failover of a failover object, returning it to its primary server while the object is in use.

---

### Caution

---

The Return to the Primary server function can allow the dismounting of a disk, regardless of the state of the disk, while a user is accessing data files on the disk. In that case, the data is vulnerable to data loss or corruption. Therefore, to redistribute objects after recovery from a power outage, enable the Remain on Failover server function and use the manual failover function.

---

- **Remain on Failover server**—Select this radio button to specify that, when the primary server becomes available again, the group should remain on line to the failover server instead of returning to the primary server.
6. Choose either:
- OK to save the settings and close the dialog box
  - Cancel to cancel any modifications to the settings
  - Apply to commit the changes but leave the dialog box open

## Deleting a Failover Group

You can easily delete one or more failover groups from your configuration, if needed. For example, if you want to rename a group, you can delete the group and then create it with a new name.

## Working with a Failover Group

### ► **To delete a failover group:**

1. From a Cluster or Class View, select the group name that you want to delete.
2. From the Delete menu, choose Delete Failover Group, or click the Delete button on the toolbar.

Cluster Administrator deletes the selected group after you confirm the operation.

---

#### **Note**

---

When you delete a group, the objects continue to exist—you are not deleting the objects but only the group that contains them. Deleting the group frees the objects for adding to other groups.

---

When you delete a failover group, all of the objects contained in the group are placed off line.

---

## Application Considerations

The DIGITAL Clusters for Windows NT software furnishes high server availability to cluster resources. Versions 1.1 and 1.0 with Clusters Service Pack 2 provide the following failover capabilities:

- NTFS file services and network shares (NetBEUI, TCP/IP, and IPX/SPX)
- Microsoft SQL Server 6.5
- Oracle7 Server 7.1, 7.2, and 7.3
- IP address failover support for applications typically accessed by an IP address (DIGITAL Clusters 1.1 only)
- Web servers: Microsoft Internet Information Server (IIS) and Netscape Enterprise Server (DIGITAL Clusters 1.1 only)
- Lotus Notes 4.1 and Domino 4.5 (the latter in DIGITAL Clusters 1.1 only)
- Any application that can be launched and shut down in a script

Although the cluster software supports failover of client connections to cluster resources, it does not support failover of open files. This restriction applies to client applications using named pipes for access to cluster resources as well. A *named pipe* is a network transport used for interprocess communication. Clients talk to servers by addressing a particular named pipe name. The cluster name service provides a means for a clusterwide named pipe name to refer to one server and then another.

DIGITAL Clusters supports generic application failover by allowing the system administrator to provide scripts that execute when a failover group comes on line, and when a failover group goes off line. Generic application failover scripts are most commonly used to start and stop applications that are not directly supported by cluster software, for example, custom server applications that use data on a shared cluster disk. Failover scripts also are

## Application Handling During a Failover

used to send notification messages to the system administrator in the case of a failover event.

DIGITAL Clusters 1.1 supports failing over IP socket-based applications, including the following web servers: Microsoft IIS and Netscape Enterprise Server.

## Application Handling During a Failover

Like any high-availability product, DIGITAL Clusters does not preserve application context. This holds true for both database and NTFS file service failover. The situation is the same as when a server fails and is brought back on line today. The difference lies in the speed of the repair; in a cluster, the service provided by the failed server is assumed by the other cluster server significantly faster than the normal repair time.

## Failover of Client Connections

When a client application initiates a connection to a cluster resource, the client cluster software determines which server controls the object and establishes the appropriate connection. The connection can be established using the client application of your choice. Examples of such applications include SQL Server client applications, Oracle7 Server client applications, File Manager, Explorer, Network Neighborhood, and the `net use` command.

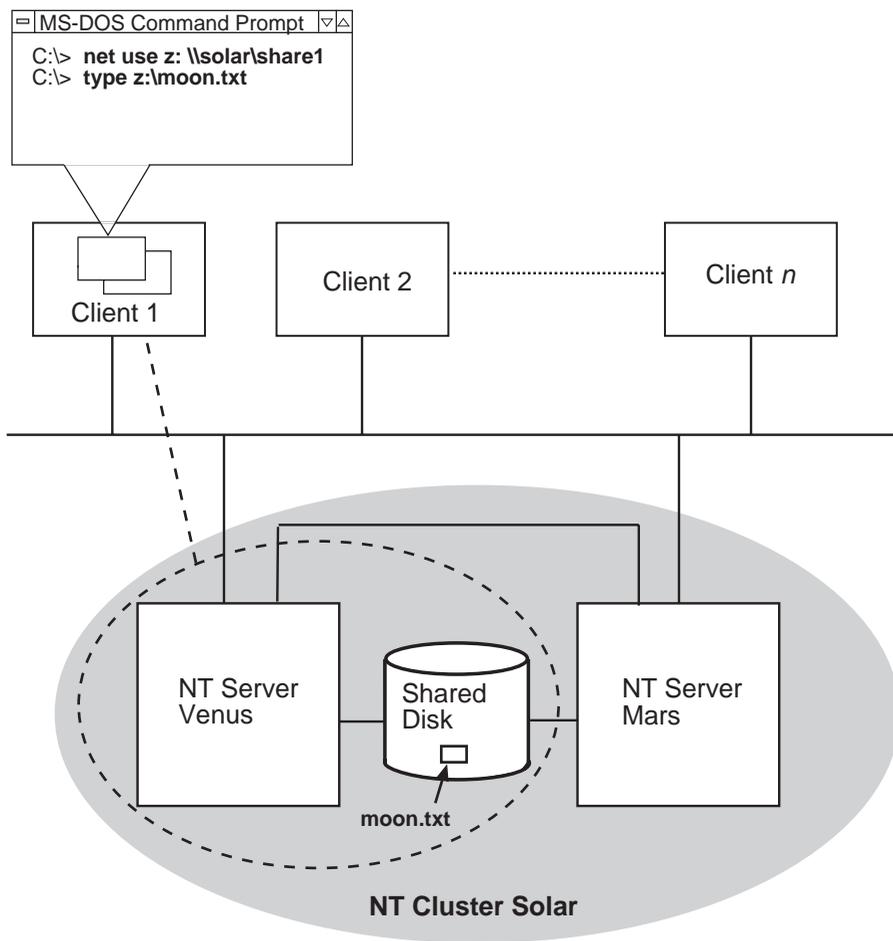
Once the connection is established, the cluster software verifies the connection with each file open command. If the resource should fail over to the failover server, the cluster client software detects this and establishes a connection to the failover server.

## Example of Client Connection Failover

This section provides an example of client connection failover using the `net use` command. The concepts also apply to other types of clients connecting to cluster resources.

The following figure shows a cluster named `solar` that is comprised of two servers, `venus` and `mars`, and a shared disk. The server `venus` controls a file share, `share1`, which resides on the cluster disk.

### NT Cluster Configuration Before Client Connection Failover



ZK-8761A-FH5

The following command establishes a connection to \\venus\share1:

```
net use z: \\solar\share1
```

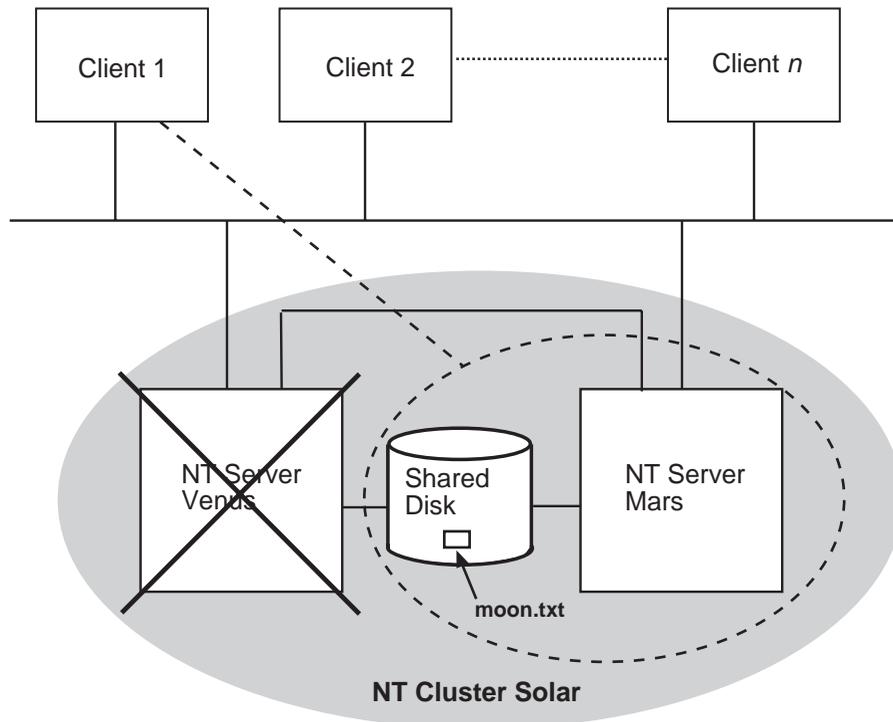
The next command opens the file star.txt on \\venus\share1:

```
type z:\star.txt
```

## Database Application Failover

If `venus` fails, as shown in the next figure, `share1` fails over to the failover server `mars`. Then, when the user reissues the `type` command, the client cluster software determines that the resource is now controlled by `mars`. The cluster software establishes a connection to `mars` and services the client application's file request.

### NT Cluster Configuration After Client Connection Failover



ZK-8762A-FH5

## Database Application Failover

DIGITAL Clusters supplies high availability to a database. If a client application is reading or writing data to a database on a disk or system that fails, the database is failed over to the failover server. The application is responsible for determining where the application resumes the operation. For SQL Server and Oracle7 Server, any in-progress transaction will have to be rolled back and restarted; one of the primary functions of database software is to provide the transactional semantics around database operations. This requirement is the same for all high-availability products for Windows NT.

In cases where there is additional server-side software (for example, a custom in-house server application), the server software needs to be failed over by mechanisms similar to those used by the underlying database. The generic application failover capability addresses many of these server applications by allowing users to provide command-line scripts for starting and stopping server applications.

### Database Client Application Failover

Client applications that are connecting to a database that fails over lose the connection and receive an I/O or connection-lost error. As is true for any Windows NT high-availability product, it is the responsibility of the application to reestablish the connection. If the client has established the connection to the cluster alias, the client simply reconnects to the cluster alias and the cluster name service automatically routes requests to the failover server.

### SQL Client Application Considerations

Microsoft SQL Server 6.5 also offers failover features for clients connecting to a specific cluster server using the Open Database Connectivity (ODBC) or DB-Library application programming interfaces (APIs). In this case, failover happens as previously described. However, failover is controlled by the ODBC or DB-Library interface instead of the cluster software. Therefore, when the client reconnects, it connects to a specific server.

## Additional Client Application Considerations

This section presents other client application considerations, including:

- Open files, named pipes, and IP socket connections
- Failover times

### Open Files, Named Pipes, and IP Socket Connections

If a cluster resource that contains an open file, named pipe, or IP socket connection fails over to the failover server, any subsequent read or write operation will fail. In this case, the client application must attempt to reconnect by closing and then reopening the file, named pipe, or IP socket connection to establish a connection to the failover server.

---

#### Note

---

Once the client application reconnects to the failover server, check to ensure that all data has been saved previously. Do not assume that all data was written to disk prior to the failover.

---

## What the User Sees During a Failover

When a file, named pipe, or IP socket connection is open during failover, it is the client application's responsibility to maintain its own context and roll back to a safe point before continuing. In the event of a failure, a user can reexecute a standard file copy application. Alternatively, the application may employ a checkpoint scheme that allows it to restart from the last checkpoint.

### Failover Times

From the client's viewpoint, failover times can vary depending on the situation. Clients cannot connect to the cluster during a cluster failover. If a new connection is initiated during failover, the client will attempt to connect to the cluster for 15 seconds before receiving an error. Once the failover is complete, subsequent attempts to connect will succeed.

Clients with open files to the cluster are a special case. The first time the client attempts to access the open file during *or after* a cluster failover, the client receives an error message from the redirector. The redirector times out, trying to access the open file with the translated universal naming convention (*UNC*) address of the cluster member that has failed over. The client name service might take up to a minute to discover, and be redirected to, the server that is providing the requested cluster service. Following the error message, the client must close and reopen the file or named pipe.

## What the User Sees During a Failover

This section describes what a user of a client connected to a cluster experiences when a failover occurs.

### Supported Clients

In the case of a failover, Windows clients using the cluster alias to access the cluster will have minimal disruption to their work. In particular, if the end user is accessing network file shares, or is using a "well-behaved" client application (one that attempts to reconnect in the event of a failover) for accessing the cluster via named pipes or Service Message Block (*SMB*), the end user may experience no disruption at all—if the user does not access the cluster until the failover is complete.

As noted previously, if a user accesses the cluster during a failover, the user will lose the connection and receive an I/O or connection lost error message. In this case, the user just clicks the Retry button to reestablish the connection to the cluster.

### Clients Not Using the Cluster Alias

Clients that cannot access the cluster through the cluster alias can still benefit from the high availability offered by DIGITAL Clusters. However, these users must know the names of the two clustered servers. In the case of a failover, the user manually reconnects the client to the failover server to continue access to the cluster service.

This chapter gives troubleshooting procedures for commonly encountered problems in the DIGITAL Clusters for Windows NT environment. Using this chapter, you will be able to resolve many problems without support intervention. In situations where support intervention is required, this chapter offers guidelines for what information you should have before calling your primary support provider.

This chapter contains several references to cluster-specific entries in the Windows NT Registry. To open and examine the Registry, use the Registry editor (located at %systemroot%\system32\regedt32.exe). For detailed information on using the Registry editor, either select Help from within the program or refer to the documentation that came with your Windows NT operating system.

## Configuration Problems

This section discusses problems you are most likely to encounter while configuring your cluster software.

When you suspect a configuration problem, you should run the cluster utility CLUIVP to verify your cluster installation. This utility runs a series of tests that check the basic setup of your cluster hardware and software. (See the section CLUIVP Utility in Appendix A for details.)

### **Where are my disks? I can't create a failover group.**

Using Cluster Administrator, if you display the Class View and there appear to be no shared disks with which to create your failover groups, check the following:

1. Did you reboot both servers after installing the cluster server software?

When the servers are rebooted, the signature of each disk in the shared storage is read and both Registries are updated with the signature information. Failure to reboot after installing the cluster server software will result in no disk signatures being read.

## Configuration Problems

2. Have you waited long enough for the Registries to be updated?

Cluster Administrator takes a snapshot of the Registry when it starts up. However, it takes about a minute after the *second* server is rebooted for the disk signatures to be written to both Registries. Wait a minute and click Refresh.

3. Have all the necessary cluster services been started?

Cluster Administrator will not see the shared disks unless all required cluster services are running. Using the Services applet on the Control Panel, verify that the DIGITAL Clusters Failover Management Database (CFMD) Server, the DIGITAL Clusters Failover Manager, and the DIGITAL Clusters Name Service have all been started.

For any services that are not running, verify that the account name is valid and that the password is correct by reentering the account information. You also should use Domain Administrator to check that the account has the advanced user right to “Log on as a service.”

4. Is your shared bus configured properly?

The majority of cluster configuration problems are the result of improper configuration of the shared SCSI bus. The most common problems are as follows:

- Your shared bus exceeds the maximum cable length.
- Your shared bus is improperly terminated.
- You are using an unsupported bus adapter or the adapter’s hardware or firmware revision level is out of date.
- You are using an unsupported disk or the disk’s hardware or firmware revision level is out of date.
- You have specified duplicate SCSI IDs on the shared bus.

Refer to the *Configuration and Installation Guide* for information on the proper way to configure your shared SCSI bus.

5. Is your shared bus specified properly?

Using the Registry editor, access the SCSI device map key:

```
HKEY_LOCAL_MACHINE\Hardware\DeviceMap\Scsi
```

Verify that the system can see the shared SCSI bus adapter and that the SCSI IDs for the adapter and disks are listed. If not, one of the following problems might have occurred during software installation:

- Your cluster has only one server.
- You attempted to join a cluster that already had two servers.

- You specified the wrong adapter for the shared bus.
  - You specified the wrong name of the second server when you specified the shared bus.
6. Has your shared SCSI bus adapter been reconfigured?

If you have moved your SCSI bus adapter to another I/O slot, added or removed bus adapters, or installed a new version of the bus adapter driver, the cluster software may not be able to access your shared disks. Look for the “Adapter configuration has changed” warning message during the blue screen phase of system reboot. If you see this message, you must use the Manage Adapter Configuration function of Cluster Administrator to respecify the cluster adapter. Then reboot the system.

7. Have there been any hardware errors or transport problems?

Use the Windows NT Event Viewer to look in the event log for disk I/O error messages or indications of problems with the communications transport. If there are no relevant error messages in the event log, open and examine the DIGITAL Clusters Failover Manager trace log.

### **My group won't come on line.**

If you run Windows NT Disk Administrator and you do not see the disk group on line to the local system, check the following:

1. Are you looking at the right disks?

If you have not labeled your disks or assigned fixed drive letters to them, you may not recognize which disks are shared and which ones are not. DIGITAL recommends that you label your disks in a meaningful manner and that you assign fixed drive letters to all partitions.

2. Is the group on line to the other server?

Using the Registry editor on the failover server, examine the DIGITAL Clusters Failover Manager's group key:

```
HKEY_LOCAL_MACHINE\System\CurrentControlSet\  
Services\Cfmd\Database\FMGroup
```

Expand the group name and look for the `ConnectionPoint` parameter. If this parameter is defined, the failover server has control of the group.

(Note that you also can use the FMSTAT utility to view graphically how the cluster resources are allocated. See the section FMSTAT Utility in Appendix A for details.)

## Failover Problems

3. Have there been any hardware problems?

Use the Windows NT Event Viewer to look in the event log for disk I/O error messages or indications of hardware problems. If there are no relevant error messages in the event log, open and examine the DIGITAL Clusters Failover Manager trace log.

## Failover Problems

This section discusses failover problems—in particular, how to determine whether the problem you are seeing is a problem with the cluster or with your client.

### Failover doesn't work.

If the cluster does not appear to be failing over properly, you need to determine whether the problem is with the cluster itself or with your client. To verify that the cluster is working properly, do the following:

1. Using the Registry editor on both servers, access the DIGITAL Clusters Failover Manager group key:

```
HKEY_LOCAL_MACHINE\System\CurrentControlSet\  
Services\Cfmd\Database\FMGroup
```

2. Expand the same group name on both servers and look for the `ConnectionPoint` parameter. The server that has this parameter defined has control of the failover group.
3. Start Cluster Administrator on the server controlling the failover group. Select the group and then choose Manual Failover from the Manage menu to transfer control to the other server.
4. Within a minute or less, the `ConnectionPoint` parameter should move to the other server, demonstrating that the cluster is failing over properly.

---

#### Note

---

Instead of opening the Registry and looking for the `ConnectionPoint` parameter, you alternatively can use the FMSTAT utility to view how the cluster resources are allocated and to observe dynamically the group fail over. See the section FMSTAT Utility in Appendix A for details.

---

If this procedure demonstrates that the cluster is functioning properly, the failover problem you are observing is on the client system. See the section My client hangs after a failover on page 11–6 for details. If you are running a database server such as Microsoft SQL Server or Oracle7 Server, you also should refer to the section Database Problems on page 11–7.

If you did not see the `ConnectionPoint` parameter move to the other server, your cluster is not operating properly. Look in the DIGITAL Clusters Failover Manager trace log for error messages that might indicate the problem. The trace log file is found in the `temp` subdirectory of the cluster destination directory, as specified during cluster server software installation. If you accepted the default directory during installation, the trace log file path is as follows on your system disk:

```
\Program Files\DIGITAL\Cluster\temp\fm#.log
```

where `#` is the sequential number of the trace log file. (A new trace log file is generated with each system reboot.)

## Client Problems

This section discusses problems that you might encounter with your client system.

### My client doesn't see any clusters.

If your client system cannot enumerate any clusters, check the following:

1. Is the cluster client software installed?

Unless you have installed the cluster client software on your client system, you will not be able to see and use the cluster aliases.

2. Is the network configured properly?

Your client system must be on the same LAN as the cluster it is trying to access. The client also must have at least one transport protocol in common with the server. (The cluster software supports the following protocols: NetBEUI, IPX/SPX, and TCP/IP.) If the client is using only the TCP/IP transport protocol, the cluster client and server systems also must be on the same subnet.

### My client doesn't see the cluster it needs.

The DIGITAL Clusters Name Service maintains a list of all clusters on the LAN. This list is found at the following Registry key:

```
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\  
ClusterNameServer\ClusterNameCache
```

When a client requests a list of available clusters, it selects a server at random to provide the list as found in its Registry. If the server's `ClusterNameCache` key is corrupted, the client will not get a complete list.

Rebooting your client may fix the problem for that client. However, you should use the Registry editor to examine the `ClusterNameCache` key for all clusters on the LAN to locate the server with the faulty list.

### **My client can't access cluster resources.**

There are several reasons your client application might not be able to access a particular shared cluster resource. To determine what the problem is, check the following:

1. Is the resource currently being failed over?

Depending on the circumstances, failover can take up to 1 minute. Be patient and try again after a reasonable length of time.

2. Does the application have access rights to the resource?

Remember that the cluster software is layered over Windows NT, and therefore, access to a resource is governed by the same rules and restrictions as imposed by the operating system.

3. Did you export the share after creating the failover group?

Again, the cluster software follows the same rules and restrictions as Windows NT. You must create a network share from your cluster failover groups for the disks to be visible across the network.

4. Does a connection point exist for the resource?

Using the Registry editor, access the `NetworkShare` key on both servers:

```
HKEY_LOCAL_MACHINE\System\CurrentControlSet\  
Services\Cfmd\Database\NetworkShare
```

In the list of resources, double-click on the resource in question and look for the `ConnectionPoint` parameter. The server that has this parameter defined controls the resource. If neither server has defined a `ConnectionPoint` parameter for the resource, then the DIGITAL Clusters Failover Manager has not put the resource on line.

### **My client hangs after failover.**

There are several reasons why a client application might hang or appear to hang after failover, including the following:

- The resource may be taking a long time to fail over. Be patient.
- The server going down is the domain controller, restricting client access.
- The resource has not failed over to the other server properly.
- The application does not know how to gracefully handle a failover.

If the application is, in fact, hanging, you need to determine whether the problem is the result of a system, cluster, or application problem. To determine this, do the following:

1. Using the LANman server name (rather than the cluster alias), attach the application to the shared resource via Microsoft Network instead of DIGITAL Clusters for Windows NT.
2. Shut down the cluster server and observe the behavior of the application. If the application still hangs, the problem is *not* a cluster problem.

To determine whether the resource is failing over properly, do the following:

1. Using the Registry editor, access the `NetworkShare` key on both servers:  

```
HKEY_LOCAL_MACHINE\System\CurrentControlSet\
  Services\Cfmd\Database\NetworkShare
```
2. In the list of resources, double-click the resource in question and look for the `ConnectionPoint` parameter.
3. If neither server has a defined `ConnectionPoint` parameter for the resource, then the DIGITAL Clusters Failover Manager has not put the resource on line. See the section `My group won't come on line` on page 11–3 for details.

If the hang occurs because the application does not know how to handle failover, do the following:

1. Try to stop and restart the application. If the application does not respond to your attempts to stop it, invoke the Task List and choose End Task.
2. If a connection to the shared resource cannot be reestablished upon application restart, you might need to manually intervene to close the old connection before establishing a new one. Depending on your client operating system, you can do this using File Manager, Explorer, or Network Neighborhood, or by issuing the `NET USE` command.

## Database Problems

This section discusses problems you might encounter while using Microsoft SQL Server or Oracle7 Server database software.

### My database isn't available.

If you get an error message indicating that your database is not available, you first should determine whether all the necessary services for your database server are running. The Microsoft SQL Server and Oracle7 Server products provide tools for the database administrator to determine whether services have been started. You also can use the Services applet of Windows NT Control Panel to check the state of the services for the database servers.

## Database Problems

If the required services are not running, check for the following:

1. Did you configure the software properly?

You must preconfigure the database software for use with the cluster software. For Oracle7 Server, you must create an Oracle failover object and add it to a failover group. For SQL Server, you must use Cluster Administrator to enroll the SQL Server database for failover support. This operation configures the SQL Server software on both cluster servers and creates the SQL Server failover object needed to enable database failover. (See Chapter 3, Configuring Database Software for Failover, for details.)

If you are running Oracle7 Server, you also must ensure that all database services are set for automatic startup. If all the required services are running but the database server itself is not started, make sure that the data you entered in the Oracle failover object dialog box is correct. Also, as database administrator, verify that you can start and stop the database server manually using the information in the failover object dialog box.

Note also that if you are using Oracle7 Server, your database drive letters must be fixed and they must be the same on both servers.

2. Have there been any startup problems?

Look in the event log or the DIGITAL Clusters Failover Manager trace log for login error messages or indications of problems during cluster startup. If you are running Microsoft SQL Server, also check for any error messages in the log maintained by the MSSQLServer service.

### **My database failover group won't come on line.**

If your database server is running but your database failover group does not come on line, you first need to determine if the problem is a cluster problem or a database problem.

Remove the database failover object from the cluster failover group and attempt a manual failover. Using the Class View of Cluster Administrator, if you do not see the shared disks on the remote system, the problem is *not* a database problem. Refer to the section Configuration Problems on page 11-1 to troubleshoot the problem.

If the cluster group fails over properly without the database failover object, run a database administrator tool for your software and enumerate your databases:

- If your database is listed as “offline,” then the database disk probably was not available when the server was started.
- If your database is listed as “suspect” (Microsoft SQL Server only), the server could not recover the database when it was brought on line. See Chapter 3 for instructions on how to recover your database.
- If you are running Microsoft SQL Server and you do not see your database listed, you probably have not enrolled the database on the local server. See Chapter 3 for details.

### **My Oracle7 Server won't fail over to the other server system.**

If you are running Oracle7 Server and your database server will not fail over to the other server, verify that:

- All the data files for the database server are located on shared disks.
- All the services for the database server are set to automatically start on *both* cluster servers.

### **My Oracle7 Server is running but the client can't access it.**

If you are running Oracle7 Server and your client application is unable to communicate with the server, the problem may be in the SQL\*Net network listener. Check to make sure that SQL\*Net Listener has been started and is configured correctly. Check also that the database alias has the correct transport information.

# A

---

## Troubleshooting Tools and Resources

This appendix describes some of the software tools and system resources that you can use to diagnose problems with your cluster.

### Software Utilities

Microsoft Windows NT and DIGITAL Clusters for Windows NT provide a variety of software utilities that are useful in troubleshooting your system.

#### regedt32

The DIGITAL Clusters software stores configuration data in the Windows NT Registry. You can use the Registry editor (located at %systemroot%\system32\regedt32.exe) to examine and modify the contents of the Registry. For more information on using the Registry editor, either select Help from within the program or refer to the documentation packaged with your Windows NT operating system.

See the section Registry on page A-6 for a description of the cluster-specific information stored in the Registry.

#### Disk Administrator

You use the Windows NT Disk Administrator to format and partition your disks and to assign fixed drive letters to the partitions. When troubleshooting disk problems, you can use Disk Administrator to determine whether a disk is on line to a particular system. If the disk is selectable under Disk Administrator, it is on line to the local system. Otherwise, if the disk is unselectable—that is, shown in gray—it is off line to the local system.

### Services Applet

You can use the Services applet of the Windows NT Control Panel to verify that the DIGITAL Clusters services are running. Following are the required DIGITAL Clusters services:

- DIGITAL Clusters Failover Management Database (CFMD) Server
- DIGITAL Clusters Failover Manager
- DIGITAL Clusters Name Service

### NET SHARE Command

You can use the NET SHARE command to verify that a particular server has exported the proper shares.

► **To use the NET SHARE command:**

1. Open an MS-DOS command prompt window.
2. Type the following command:

```
net share
```

### NET VIEW Command

You can use the NET VIEW command to verify that a particular share has been exported from the server you expected. Alternatively, you can get this information using Windows NT Explorer.

► **To use the NET VIEW command:**

1. Open an MS-DOS command prompt window.
2. Type the following command:

```
net view \\server
```

where *server* is the name of the server of interest.

► **To use Windows NT Explorer:**

1. From the Programs group, start Windows NT Explorer.
2. From the Tools menu, choose Map Network Drive.  
The Map Network Drive dialog box is displayed.
3. Browse the Shared Directories box for the server of interest.

## NETMON Network Monitor Utility

One way to determine whether a cluster problem is located on a server or a client is to use a network monitor. Using a network monitor, you can verify that:

- The cluster announcement is being sent from the server to the client. A cluster client cannot see the DIGITAL Clusters Name Server if it is unable to receive the cluster announcement.
- Requests and responses are being passed between the client and the DIGITAL Clusters Name Server.

DIGITAL Clusters for Windows NT provides an extension DLL for the System Management Server (SMS) network monitor program, NETMON, which formats the DIGITAL Clusters Name Service protocol. To enable this extension DLL, follow the next procedure on one or both of your cluster servers.

### ► To enable the DIGITAL Clusters extension DLL for NETMON:

1. Put the file `cnsmon.dll` in the `parsers` subdirectory of NETMON.
2. Using a text editor, edit file `parser.ini` as follows:
  - a. Add `cnsmon.dll` to the list of DLLs at the beginning of the file by adding the following line:
 

```
CNSMON.DLL = 0: CNS
```
  - b. For each NetBIOS protocol listed in the file, add CNS to the following set list by adding the following to the end of each protocol line:
 

```
, CNS
```
  - c. Add the following section to the end of the file:
 

```
[CNS]
      Comment = "DIGITAL Clusters Name Service protocol"
      FollowSet      =
      HelpFile       =
```
3. Using the Network applet on the Windows NT Control Panel, verify that a matching protocol is enabled on the client system. If the client is running TCP/IP, ensure that it is in the same subnet as the server system.

### CLUIVP Utility

The cluster distribution CD-ROM contains an unsupported utility, CLUIVP, that you can use to verify your cluster installation. The CLUIVP utility performs a series of simple checks on the cluster components to verify that the cluster software is installed and operating properly. It also checks the network connection and the SCSI hardware configuration.

Before running CLUIVP, you must have installed the cluster software on both server systems. After rebooting both servers, you must stop the DIGITAL Clusters Failover Manager on each server. This puts all shared disks off line.

CLUIVP is located in the clusters program directory on the distribution CD-ROM. You run the program from the MS-DOS command line. It requires no special parameters or options.

You must run CLUIVP concurrently on both servers. On one server, the program acts as the requester, directing the tests and displaying the results. On the other server, the program acts as the responder, carrying out test operations and reporting the results to the requester.

CLUIVP tests and reports on operations in the following areas:

- Remote Procedure Call (RPC) communication
- Cluster infrastructure (based on NetBIOS) communication
- Updates to the DIGITAL Clusters Failover Management Database (CFMD)
- Disk response and conformance to SCSI protocols required by the cluster software

### CLUXFER Utility

DIGITAL Clusters Version 1.1 includes a new command line utility, CLUXFER. This utility allows you to initiate a manual failover of a cluster failover group either from within a batch file or from an application. With CLUXFER, you can transfer a cluster failover group to a specific cluster server.

► **To use the CLUXFER utility:**

1. Open an MS-DOS command prompt window.
2. Type the following command:

```
cluxfer FailoverGroup [Server]
```

where the variables have the following meanings:

*FailoverGroup*      Specifies the name of the failover group to transfer.

*Server*

An optional parameter that specifies the server to which you would like the failover group to transfer.

If you specify the *Server* parameter, the cluster software tries to place the failover group on line on this server.

If you do not specify the *Server* parameter and the failover group is currently on line, the group is placed off line on one cluster server and brought on line on the other. Alternately, if the failover group is currently off line, it is brought on line on its primary server.

## Cluster Monitor Utility

The Cluster Monitor utility is a simple cluster status display program that can be run on any Windows NT system. It displays the status of all failover groups in a cluster from the perspective of the cluster servers.

The Cluster Monitor may be run both from cluster servers and from systems that are not cluster members. When run from cluster clients, the DIGITAL Clusters client software is not required.

To run the Cluster Monitor from a cluster server, you must be running from the same cluster account that you specified during cluster software installation. To run the Cluster Monitor from a system that is not a cluster member, simply copy the file `fmstat.exe` to the desired system and run it as described next.

The Cluster Monitor utility can be invoked in several ways: from the DIGITAL Clusters for Windows NT Program Group, from the Tools menu of Cluster Administrator, from the MS-DOS command line, or from the Run dialog box.

► **To use the Cluster Monitor utility from the DIGITAL Clusters for Windows NT program group:**

1. From the DIGITAL Clusters for Windows NT program group, choose Cluster Monitor.  
The Cluster Monitor main window is displayed.

► **To invoke the Cluster Monitor utility from Cluster Administrator:**

1. Using Cluster Administrator, choose Cluster Monitor from the Tools menu.  
The Cluster Monitor main window is displayed.

## Software Utilities

### ► To invoke the Cluster Monitor utility from an MS-DOS command prompt:

1. Open an MS-DOS command prompt window.
2. Type the following command:

```
fmstat [ClusterMember]
```

where *ClusterMember* is an optional parameter that you use if you want to:

- Run the Cluster Monitor from a noncluster member.
- Examine the status of another cluster.

3. Press the Return key.

The Cluster Monitor main window is displayed.

### ► To invoke the Cluster Monitor utility from the Run dialog box:

1. Choose Run from the Start menu.

The Run dialog box is displayed.

2. In the Open box, type the following:

```
fmstat [ClusterMember]
```

where *ClusterMember* is an optional parameter that you use if you want to:

- Run the Cluster Monitor from a noncluster member.
- Examine the status of another cluster.

3. Choose the OK button.

The Cluster Monitor main window is displayed.

## Cluster Monitor Utility Display

The Cluster Monitor utility displays a row for each failover group and a column for each of the two cluster servers. For example, in the following display, General Disk Group, Oracle Group, and Test Group are failover groups and F16a and F16C are cluster servers.



Within this matrix, the current status of each failover group is color coded as follows:

Color Code	Failover Group Status	Description
Blue	Server is down	Indicates that the server for the failover group is down or not responding to RPC requests. You may see a server is down status for a short period after starting Cluster Monitor while the program is gathering cluster status information.
Red	Failover group is off line	Indicates that the failover group is off line on this server and has never been online on this server.
Red, Green, or Yellow	Date/time stamped	<p>Red Indicates that the failover group is off line on this server. The date/time stamp indicates the last time the failover group was online.</p> <p>Green Indicates that the failover group is on line on this server and that this is the primary server for the failover group. The date/time stamp indicates the time when the failover group came on line on this server.</p> <p>Yellow Indicates that the failover group is on line on this server and that this is the failover server for the failover group. The date/time stamp indicates the time when the failover group came on line on this server.</p>

Click the left mouse button on any status entry for more detailed information about a failover group.

## Registry

By default, the display is refreshed every 5 seconds. You can change how often the program polls the cluster servers by using the Display icon of the Control Panel. Choose the Settings tab and select an alternate Refresh Frequency value.

## Registry

The cluster software stores the following types of cluster-specific information in the Windows NT Registry:

- Configuration information used to start cluster drivers and services at system startup. You can use this information to verify that the proper software components have been installed and are ready to run.
- Configuration information used to control the operation of the cluster software. You can use this information to verify that the cluster is configured properly.
- Dynamic state information describing the current state of the cluster software. You can use this information to verify that the cluster resources are allocated as you expect.

## DIGITAL Clusters Registry Keys

The registry keys used by the DIGITAL Clusters for Windows NT software are described briefly in the following sections. Refer to Appendix B for an example of the contents of each key.

### DIGITAL Clusters Failover Management Database (CFMD) Key

The DIGITAL Clusters Failover Management Database (CFMD) key contains entries that describe the failover objects configured for the cluster, as well as information about various software components used to manage the cluster.

The CFMD key is found at:

```
\Registry\Machine\System\CurrentControlSet\Services\Cfmd
```

### DIGITAL Clusters Port Driver (CluPort) Key

The DIGITAL Clusters Port Driver (CluPort) key contains entries that describe the cluster's shared buses. There is one entry for each shared SCSI bus adapter.

The CluPort key also contains values used by the Windows NT operating system to start the cluster SCSI port driver when the system is booted.

The CluPort key is found at:

```
\Registry\Machine\System\CurrentControlSet\Services\CluPort
```

### **DIGITAL Clusters Disk Driver (CluDisk) Key**

The DIGITAL Clusters Disk Driver (`CluDisk`) key contains values used by the Windows NT operating system to start the DIGITAL Clusters SCSI disk driver when the system is booted.

The `CluDisk` key is found at:

```
\Registry\Machine\System\CurrentControlSet\Services\CluDisk
```

### **DIGITAL Clusters File System (CFS) Key**

The DIGITAL Clusters File System (`CFS`) key contains values used by the Windows NT operating system to start the DIGITAL Clusters File System driver when the system is booted.

The `CFS` key is found at:

```
\Registry\Machine\System\CurrentControlSet\Services\Cfs
```

### **DIGITAL Clusters Log Watch Key**

The DIGITAL Clusters Log Watch (`LogWatch`) key contains values used by the Windows NT operating system to start the Log Watch server when the system is booted. The Log Watch server is responsible for replicating event log entries between the two cluster servers.

The `LogWatch` key is found at:

```
\Registry\Machine\System\CurrentControlSet\Services\LogWatch
```

### **DIGITAL Clusters Failover Manager Key**

The DIGITAL Clusters Failover Manager (`ClusterFailoverManager`) key contains entries that are used to control the timing of certain cluster operations, as well as entries specifying the location of the DIGITAL Clusters trace log files. (Some of these entries are explained in the section DIGITAL Clusters Registry Key Tuning Parameters, later in this appendix.)

The `ClusterFailoverManager` key also contains values used by the Windows NT operating system to start the DIGITAL Clusters Failover Manager when the system is booted.

The `ClusterFailoverManager` key is found at:

```
\Registry\Machine\System\CurrentControlSet\  
Services\ClusterFailoverManager
```

### **DIGITAL Clusters Name Service Key**

The DIGITAL Clusters Name Service (`ClusterNameService`) key contains the DIGITAL Clusters name cache—a list of the names of all the clusters in the LAN.

The `ClusterNameService` key is found at:

## Registry

```
\Registry\Machine\System\CurrentControlSet\Services\  
ClusterNameService
```

### SCSI Device Map Key

Windows NT stores a description of the hardware configuration in the Registry every time the system reboots. To diagnose problems with your cluster hardware, you can examine the description of the SCSI devices configured on your system.

The description of your SCSI devices is found in the Registry at the following key:

```
\Registry\Machine\Hardware\Devicemap\Scsi
```

### DIGITAL Clusters Registry Key Tuning Parameters

The Registry contains several parameters that affect the behavior of your cluster. Some of these parameters are described in the following sections.

DIGITAL Clusters parameters are stored in a subkey (*Parameters*) of the keys for their respective services: DIGITAL Clusters Failover Management (CFMD) Server and DIGITAL Clusters Failover Manager. Use the Registry editor to examine and modify these parameters.

Note that any modifications you make to the Registry entries do not take effect until you reboot your system.

### CfmdTrace

The *CfmdTrace* parameter of the DIGITAL *Cfmd* Registry key is used to control the location of the CFMD trace file. The trace file contains information used by DIGITAL support personnel.

Initially, the *CfmdTrace* parameter is not defined and tracing is not enabled. Follow the next procedure to enable tracing.

#### ► To enable the *CfmdTrace* parameter of the DIGITAL CFMD Registry key:

1. Using the Registry editor, access the following key:

```
\Registry\Machine\System\CurrentControlSet\Services\  
Cfmd\Parameters
```

---

**Note**

---

It is possible that the `\Registry\...\Cfmd\Parameters` subkey may not exist when you attempt to create the `CfmdTrace` parameter. If not, you must create the subkey:

- a. From the Edit menu, choose Add Key.
  - b. Leave the Class Name field blank.
- 

2. Create the `CfmdTrace` parameter:
  - a. From the Edit menu, choose Add Value.
  - b. Set the data type to `REG_SZ`.
3. Set the parameter value equal to the full path specification of the trace log file.

Note that you can use wildcard sequencing characters (#) in the file name. If you do this, the CFMD Server creates a new trace file each time the service is started, numbering the files sequentially. For example, specifying a file name of `cfmd###.log` causes the DIGITAL CFMD server to create log files named `cfmd001.log`, `cfmd002.log`, and so on.

Follow the next procedure to disable CFMD tracing.

► **To disable the `CfmdTrace` parameter of the DIGITAL CFMD Registry key:**

1. Using the Registry editor, access the following key:

```
\Registry\Machine\System\CurrentControlSet\Services\Cfmd\Parameters
```
2. Select the `CfmdTrace` parameter value.
3. From the Edit menu, choose Delete Value.

**CisTrace**

The `CisTrace` parameter of the DIGITAL `ClusterFailoverManager` Registry key is used to control the location of the cluster infrastructure (CIS) communication trace file. The trace file contains information used by DIGITAL support personnel.

Initially, the `CisTrace` parameter is not defined and tracing is not enabled. Follow the next procedure to enable cluster infrastructure communication tracing.

## Registry

► **To enable the CisTrace parameter of the DIGITAL ClusterFailoverManager Registry key:**

1. Using the Registry editor, access the following key:  

```
\Registry\Machine\System\CurrentControlSet\Services\  
ClusterFailoverManager\Parameters
```
2. Create the CisTrace parameter:
  - a. From the Edit menu, choose Add Value.
  - b. Set its data type to REG\_SZ.
3. Set the parameter value equal to the full path specification of the trace log file. Note that you can use wildcard sequencing characters (#) in the file name. If you do this, the DIGITAL Clusters Failover Manager creates a new trace file each time the service is started, numbering the files sequentially. For example, specifying a file name of cis###.log causes the DIGITAL Clusters Failover Manager to create log files named cis001.log, cis002.log, and so on.

Follow the next procedure to disable cluster infrastructure communication tracing.

► **To disable the CisTrace parameter of the DIGITAL ClusterFailoverManager Registry key:**

1. Using the Registry editor, access the following key:  

```
\Registry\Machine\System\CurrentControlSet\Services\  
ClusterFailoverManager\Parameters
```
2. Select the CisTrace parameter value.
3. From the Edit menu, choose Delete Value.

## ConnectionTimeout

The ConnectionTimeout parameter of the DIGITAL ClusterFailoverManager Registry key is used to control how long (in milliseconds) the communications infrastructure waits before declaring the network connection between the servers as being down. The shorter the value, the more likely that a network glitch or momentary load on the server CPU will cause a false indication of failure.

Initially, the parameter is not defined and a default value of 30,000 milliseconds (30 seconds) is used. Follow the next procedure to change the cluster communications infrastructure timeout value.

► **To modify the ConnectionTimeout parameter of the DIGITAL ClusterFailoverManager Registry key:**

1. Using the Registry editor, access the following key:  

```
\Registry\Machine\System\CurrentControlSet\Services\  
ClusterFailoverManager\Parameters
```
2. If the ConnectionTimeout parameter does not exist, create it:
  - a. From the Edit menu, choose Add Value.
  - b. Set its data type to REG\_DWORD.
3. Set the ConnectionTimeout parameter value to the desired timeout, in milliseconds.

► **To revert to the default ConnectionTimeout parameter value of the DIGITAL ClusterFailoverManager Registry key:**

1. Using the Registry editor, access the following key:  

```
\Registry\Machine\System\CurrentControlSet\Services\  
ClusterFailoverManager\Parameters
```
2. Select the ConnectionTimeout parameter value.
3. From the Edit menu, choose Delete Value.

### **FailoverEvaluateDelay**

The FailoverEvaluateDelay parameter of the DIGITAL ClusterFailoverManager Registry key is used to control the delay (in milliseconds) on failing over a group to allow the group containing the log disk to come on line first.

Initially, the parameter is not defined and a default value of 20,000 milliseconds (20 seconds) is used. Follow the next procedure to modify the timeout value of the FailoverEvaluateDelay Registry tuning parameter.

► **To modify to the default FailoverEvaluateDelay parameter value of the DIGITAL ClusterFailoverManager Registry key:**

1. Using the Registry editor, access the following key:  

```
\Registry\Machine\System\CurrentControlSet\Services\  
ClusterFailoverManager\Parameters
```

## Registry

2. If the `FailoverEvaluateDelay` parameter does not exist, create it:
  - a. From the Edit menu, choose Add Value.
  - b. Set its data type to `REG_DWORD`.
3. Set the `FailoverEvaluateDelay` parameter value to the desired timeout, in milliseconds.

► **To revert to the default `FailoverEvaluateDelay` parameter value of the DIGITAL ClusterFailoverManager Registry key:**

1. Using the Registry editor, access the following key:

```
\Registry\Machine\System\CurrentControlSet\Services\  
ClusterFailoverManager\Parameters
```
2. Select the `FailoverEvaluateDelay` parameter value.
3. From the Edit menu, choose Delete Value.

## ReconnectWait

The `ReconnectWait` parameter of the DIGITAL ClusterFailoverManager Registry key is used to control how long (in seconds) the cluster communications infrastructure waits before trying to reconnect to a server that is down. The parameter determines how quickly one server detects that the other server has come back up.

Initially, the parameter is not defined and a default value of 15 seconds is used. Follow the next procedure to modify the `ReconnectWait` parameter value.

► **To modify the `ReconnectWait` parameter value of the DIGITAL ClusterFailoverManager Registry key:**

1. Using the Registry editor, access the following key:

```
\Registry\Machine\System\CurrentControlSet\Services\  
ClusterFailoverManager\Parameters
```
2. If the `ReconnectWait` parameter does not exist, create it:
  - a. From the Edit menu, choose Add Value.
  - b. Set its data type to `REG_DWORD`.
3. Set the parameter value to the desired delay, in seconds.

► **To revert to the default ReconnectWait parameter value of the DIGITAL ClusterFailoverManager Registry key:**

1. Using the Registry editor, access the following key:

```
\Registry\Machine\System\CurrentControlSet\Services\  
ClusterFailoverManager\Parameters
```

2. Select the ReconnectWait parameter value.
3. From the Edit menu, choose Delete Value.

## LogLevel

The LogLevel parameter of the DIGITAL ClusterFailoverManager Registry key is used to control what types of messages are to be logged by the cluster communications infrastructure. The value indicates the minimum severity level of logged messages, as follows:

- 4 Success messages
- 3 Informational messages
- 2 Warning messages
- 1 Error messages

Note that the lower the number, the higher the severity of the message.

By setting the LogLevel parameter to 2, you can suppress the logging of success and informational messages.

Initially, the parameter is undefined and a default of 4 is used, enabling the logging of all messages. Follow the next procedure to change the types of messages logged.

► **To modify the LogLevel parameter value of the DIGITAL ClusterFailoverManager Registry key:**

1. Using the Registry editor, access the following key:

```
\Registry\Machine\System\CurrentControlSet\Services\  
ClusterFailoverManager\Parameters
```

2. If the LogLevel parameter does not exist, create it:
  - a. From the Edit menu, choose Add Value.
  - b. Set its data type to REG\_DWORD.
3. Set the LogLevel parameter value to the desired level (4, 3, 2, or 1).

## Registry

► **To revert to the default LogLevel parameter value of the DIGITAL ClusterFailoverManager Registry key:**

1. Using the Registry editor, access the following key:  

```
\Registry\Machine\System\CurrentControlSet\Services\  
ClusterFailoverManager\Parameters
```
2. Select the LogLevel parameter value.
3. From the Edit menu, choose Delete Value.

### ClusterName

The ClusterName parameter of the DIGITAL ClusterFailoverManager Registry key is used to store the cluster name. You can examine this parameter to determine that both servers have the same cluster name, but you must not change the value.

The ClusterName parameter is stored at the following location in the Registry:

```
\Registry\Machine\System\CurrentControlSet\Services\  
ClusterFailoverManager\Parameters\ClusterName
```

### DisableFailoverNetDelay

The DisableFailoverNetDelay parameter of the DIGITAL ClusterFailoverManager Registry key is used to turn off the stabilization delay imposed by the DIGITAL Clusters Failover Manager after a server is declared down and before failing over to the other server. This delay allows the communications infrastructure to fail over to a second network adapter, if one is present, or to recover from a transient outage that would otherwise cause a false failure.

Initially, the DisableFailoverNetDelay parameter is not defined, thereby enabling the stabilization delay. Follow the next procedure to disable the stabilization delay imposed by the DIGITAL Clusters Failover Manager.

► **To modify the DisableFailoverNetDelay parameter value of the DIGITAL ClusterFailoverManager Registry key:**

1. Using the Registry editor, access the following key:  

```
\Registry\Machine\System\CurrentControlSet\Services\  
ClusterFailoverManager\Parameters
```
2. If the DisableFailoverNetDelay parameter does not exist, create it:
  - a. From the Edit menu, choose Add Value.
  - b. Set its data type to REG\_DWORD.
3. Set the DisableFailoverNetDelay parameter value to any nonzero number.

► **To reenable the `DisableFailoverNetDelay` parameter value of the `DIGITAL ClusterFailoverManager` Registry key:**

1. Using the Registry editor, access the following key:  
`\Registry\Machine\System\CurrentControlSet\Services\  
ClusterFailoverManager\Parameters`
2. Select the `DisableFailoverNetDelay` parameter value.
3. From the Edit menu, choose Delete Value.  
Or, set the value to 0.

---

**Note**

---

The length of the stabilization delay is fixed; it cannot be modified.

---

### **DiskArbitrationInterval**

The `DiskArbitrationInterval` parameter of the `DIGITAL ClusterFailoverManager` Registry key is used to specify (in seconds) how frequently the server polls the disk to verify its ownership and how long the failover server will wait before seizing the disk after the primary server goes down.

Initially, the parameter is not defined and a default value of 30 seconds is used. Follow the next procedure to change the disk arbitration interval.

► **To modify the `DiskArbitrationInterval` parameter value of the `DIGITAL ClusterFailoverManager` Registry key:**

1. Using the Registry editor, access the following key:  
`\Registry\Machine\System\CurrentControlSet\Services\  
ClusterFailoverManager\Parameters`
2. If the `DiskArbitrationInterval` parameter does not exist, create it:
  - a. From the Edit menu, choose Add Value.
  - b. Set its data type to `REG_DWORD`.
3. Set the `DiskArbitrationInterval` parameter value to the desired interval, in seconds.

## Registry

► **To revert to the default `DiskArbitrationInterval` parameter value of the DIGITAL ClusterFailoverManager Registry key:**

1. Using the Registry editor, access the following key:  
`\Registry\Machine\System\CurrentControlSet\Services\  
ClusterFailoverManager\Parameters`
2. Select the `DiskArbitrationInterval` parameter value.
3. From the Edit menu, choose Delete Value.

### **DiskErrorThreshold**

The `DiskErrorThreshold` parameter of the DIGITAL ClusterFailoverManager Registry key is used to control the number of consecutive errors that can be generated by a disk device before that disk is taken off line.

Initially, the parameter is not defined and a default value of 3 is used. Follow the next procedure to change the disk error threshold value.

► **To modify the `DiskErrorThreshold` parameter value of the DIGITAL ClusterFailoverManager Registry key:**

1. Using the Registry editor, access the following key:  
`\Registry\Machine\System\CurrentControlSet\Services\  
ClusterFailoverManager\Parameters`
2. If the `DiskErrorThreshold` parameter does not exist, create it:
  - a. From the Edit menu, choose Add Value.
  - b. Set its data type to `REG_DWORD`.
3. Set the `DiskErrorThreshold` parameter value to the desired threshold.

► **To revert to the default `DiskErrorThreshold` parameter value of the DIGITAL ClusterFailoverManager Registry key:**

1. Using the Registry editor, access the following key:  
`\Registry\Machine\System\CurrentControlSet\Services\  
ClusterFailoverManager\Parameters`
2. Select the `DiskErrorThreshold` parameter value.
3. Using the Edit menu, choose Delete Value.

## DiskErrorSeparation

The `DiskErrorSeparation` parameter of the `DIGITAL ClusterFailoverManager` Registry key is used to specify (in seconds) the time in which two errors must occur to be considered consecutive.

Initially, the parameter is not defined and a default value of 300 seconds (5 minutes) is used. Follow the next procedure to change the disk error separation interval.

► **To modify the `DiskErrorSeparation` parameter value of the `DIGITAL ClusterFailoverManager` Registry key:**

1. Using the Registry editor, access the following key:  
`\Registry\Machine\System\CurrentControlSet\Services\  
ClusterFailoverManager\Parameters`
2. If the `DiskErrorSeparation` parameter does not exist, create it:
  - a. From the Edit menu, choose Add Value.
  - b. Set its data type to `REG_DWORD`.
3. Set the `DiskErrorSeparation` parameter value to the desired interval, in seconds.

► **To revert to the default `DiskErrorSeparation` parameter value of the `DIGITAL ClusterFailoverManager` Registry key:**

1. Using the Registry editor, access the following key:  
`\Registry\Machine\System\CurrentControlSet\Services\  
ClusterFailoverManager\Parameters`
2. Select the the `DiskErrorSeparation` parameter value.
3. From the Edit menu, choose Delete Value.

## FmTraceOutput

The `FmTraceOutput` parameter of the `DIGITAL ClusterFailoverManager` Registry key is used to control the location of the `DIGITAL Clusters Failover Manager` trace log. The parameter is a bit mask with the following value definitions:

- 1 Log to file (as named by the `FmTrace` parameter)
- 2 Log to console window
- 4 Log to kernel debugger
- 8 Log to a new console window

Initially, the parameter is undefined and a default of 1 is used, directing the log output to the file specified by the `FmTrace` parameter. Follow the next procedure to change the trace log output.

## Registry

► **To modify the FmTraceOutput parameter value of the DIGITAL ClusterFailoverManager Registry key:**

1. Using the Registry editor, access the following key:  

```
\Registry\Machine\System\CurrentControlSet\Services\  
ClusterFailoverManager\Parameters
```
2. If the FmTraceOutput parameter does not exist, create it:
  - a. From the Edit menu, choose Add Value.
  - b. Set its data type to REG\_DWORD.
3. Set the FmTraceOutput parameter value to the desired number (1, 2, 4, or 8).

► **To revert to the default FmTraceOutput parameter value of the DIGITAL ClusterFailoverManager Registry key:**

1. Using the Registry editor, access the following key:  

```
\Registry\Machine\System\CurrentControlSet\Services\  
ClusterFailoverManager\Parameters
```
2. Select the FmTraceOutput parameter value.
3. From the Edit menu, choose Delete Value.

## FmTrace

The FmTrace parameter of the DIGITAL ClusterFailoverManager Registry key is used to specify the full path of the DIGITAL Clusters Failover Manager trace log if trace log output is being directed to a file.

Initially, the FmTrace parameter is set using a file name of fm###.log, where # is a wildcard sequencing character. Specifying a file name containing these characters causes the DIGITAL Clusters Failover Manager to create a new trace file each time the service is started, numbering the files sequentially. For example, a file name of fm###.log causes the DIGITAL Clusters Failover Manager to create trace log files named fm001.log, fm002.log, and so on.

Follow the next procedure to change the name of the DIGITAL Clusters Failover Manager trace log file.

► **To change the FmTraceOutput parameter value of the DIGITAL ClusterFailoverManager Registry key:**

1. Using the Registry editor, access the following key:

```
\Registry\Machine\System\CurrentControlSet\Services\  
ClusterFailoverManager\Parameters
```

2. Set the FmTrace parameter value equal to the full path specification of the trace log file.

### **FmLogLevel**

The FmLogLevel parameter of the DIGITAL ClusterFailoverManager Registry key is used to control what types of messages are to be logged by the DIGITAL Clusters Failover Manager. The value indicates the minimum severity level of logged messages, as follows:

- 4 Success messages
- 3 Informational messages
- 2 Warning messages
- 1 Error messages

Note that the lower the number, the higher the severity of the message. By setting the FmLogLevel parameter to 2, you can suppress the logging of success and informational messages.

Initially, the parameter is undefined and a default of 4 is used, enabling the logging of all messages. Follow the next procedure to change the types of messages logged.

► **To change the FmLogLevel parameter value of the DIGITAL ClusterFailoverManager Registry key:**

1. Using the Registry editor, access the following key:

```
\Registry\Machine\System\CurrentControlSet\Services\  
ClusterFailoverManager\Parameters
```

2. If the FmLogLevel parameter does not exist, create it:
  - a. From the Edit menu, choose Add Value.
  - b. Set its data type to REG\_DWORD.
3. Set the FmLogLevel parameter value to the desired level (4, 3, 2, or 1).

## Registry

► **To revert to the default FmLogLevel parameter value of the DIGITAL ClusterFailoverManager Registry key:**

1. Using the Registry editor, access the following key:  
`\Registry\Machine\System\CurrentControlSet\Services\  
ClusterFailoverManager\Parameters`
2. Select the FmLogLevel parameter value.
3. From the Edit menu, choose Delete Value.

### FmTraceVerbosity

The FmTraceVerbosity parameter of the DIGITAL ClusterFailoverManager Registry key is used to control the priority of messages logged to the Failover Manager trace log. The value indicates the priority level, as follows:

- 1 Errors and major events
- 2, 3 Minor events
- 4, 5 Individual communication protocol messages
- 6 Data structure dumps

Initially, the FmTraceVerbosity parameter is undefined and a default of 3 is used, enabling the logging of all major and minor events. Follow the next procedure to change the level of detail of messages logged.

► **To change the FmTraceVerbosity parameter value of the DIGITAL ClusterFailoverManager Registry key:**

1. Using the Registry editor, access the following key:  
`\Registry\Machine\System\CurrentControlSet\Services\  
ClusterFailoverManager\Parameters`
2. If the FmTraceVerbosity parameter does not exist, create it:
  - a. From the Edit menu, choose Add Value.
  - b. Set its data type to REG\_DWORD.
3. Set the FmTraceVerbosity parameter value to the desired level (1 to 6).

► **To revert to the default FmTraceVerbosity parameter value of the DIGITAL ClusterFailoverManager Registry key:**

1. Using the Registry editor, access the following key:  
`\Registry\Machine\System\CurrentControlSet\Services\  
ClusterFailoverManager\Parameters`

2. Select the `FmTraceVerbosity` parameter value.
3. From the Edit menu, choose Delete Value.

### DIGITAL Clusters Client Tuning Parameters

This section presents client tuning parameters for the advanced user. These parameters should *not* be adjusted casually.

#### Windows 95 Clients

This section discusses Windows 95 client tuning parameters.

##### Adjusting the Time for Trying Network Connections

Requests for network connections, such as network drives and browsing, can take up to 15 seconds to fail. You can adjust this time period by changing settings in the Windows 95 Registry.

When the Windows 95 client receives network requests, it communicates with the cluster name server to determine if the request refers to a cluster alias. By default, the client waits 5 seconds for a response, and then retries the request up to two times.

You can adjust the default timeout period (5 seconds) and retry count (3 times).

---

#### Note

---

In the Windows 95 client implementation, the initial request to the cluster name server is counted as the first retry count.

---

For instance, remote users may want to increase the timeout period to allow for a slow network link such as RAS. LAN users may want to decrease the retry count because the need for retries is low unless the name server is unavailable.

#### ► To change the timeout period or retry count for a Windows 95 client:

1. Using the Registry editor, access the following key:  
`\HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\VxD\DCLNAMC`
2. If the `Parameters` subkey does not exist, create it by choosing Add Key from the Edit menu. Leave the Class Name field blank.

## DIGITAL Clusters Client Tuning Parameters

3. Create and set the `TimeoutPeriod` and `RetryCount` parameters:
  - a. From the Edit menu, choose Add Value.
  - b. Set the data type for each parameter to `DWORD`.
  - c. Set the parameters to the desired values.
4. Reboot the client system.

## Windows NT Clients

This section presents Windows NT client tuning parameters.

### Adjusting the Time for Trying Network Connections

Requests for network connections, such as network drives and browsing, can take up to 15 seconds to fail. You can adjust this time period by changing settings in the Windows NT Registry.

### TranslationResponseTimeout Parameter

When the Windows NT client receives network requests, it communicates with the cluster name server to determine if the request refers to a cluster alias. By default, the client waits 5 seconds for a response, and then retries the request up to two times.

You can adjust the default timeout period of 5000 milliseconds (5 seconds). For example, remote users may want to increase the timeout period to allow for a slow network link such as RAS.

#### ► To change the timeout period for a Windows NT client:

1. Using the Registry editor, access the following key:  
`\HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Cfs`
2. If the `Parameters` subkey does not exist, create it by choosing Add Key from the Edit menu. Leave the Class Name field blank.
3. Create and set the `TranslationResponseTimeout` parameter:
  - a. From the Edit menu, choose Add Value.
  - b. Set the data type to `DWORD`.
  - c. Set the parameter to the desired value.
4. Reboot the client system.

### **SolicitNameServerTimeout**

The `SolicitNameServerTimeout` parameter of the DIGITAL Cfs Registry key is used to adjust the time delay between soliciting a cluster name server request and waiting for a name server response. Initially, the parameter is not defined and a default value of 5,000 milliseconds (5 seconds) is used. Remote users may want to increase the `SolicitNameServerTimeout` value to allow for a slow network link such as RAS. Follow the next procedure to modify this client Registry tuning parameter.

► **To change the `SolicitNameServerTimeout` parameter for a Windows NT client:**

1. Using the Registry editor, access the following key:  
`\HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Cfs`
2. If the `Parameters` subkey does not exist, create it by choosing Add Key from the Edit menu. Leave the Class Name field blank.
3. Create and set the `SolicitNameServerTimeout` parameter:
  - a. From the Edit menu, choose Add Value.
  - b. Set the data type to `DWORD`.
  - c. Set the parameter to the desired value.
4. Reboot the client system.

### **Windows for Workgroups Clients**

This section discusses Windows for Workgroups client tuning parameters.

#### **Adjusting the Time for Trying Network Connections**

Requests for network connections, such as network drives and browsing, can take up to 15 seconds to fail. You can adjust this time period by changing settings in the file `system.ini`.

When the Windows for Workgroups client receives network requests, it communicates with the cluster name server to determine if the request refers to a cluster alias. By default, the client waits 5 seconds for a response, and then retries the request up to 2 times.

You can adjust the default timeout period (5 seconds) and retry count (3 times).

---

#### **Note**

---

In the Windows for Workgroups client implementation, the initial request to the cluster name server is counted as the first retry count.

---

## Trace Log

For example, remote users may want to increase the timeout period to allow for a slow network link such as RAS. LAN users may want to decrease the retry count because the need for retries is low unless the cluster name server is unavailable.

### ► To change the timeout period or retry count for a Windows for Workgroups client:

1. Add the following section to the `system.ini` file:

```
[DCLNAMC]
TimeoutPeriod=5
RetryCount=3
```
2. Set the `TimeoutPeriod` and `RetryCount` to the desired values.
3. Reboot the client system.

## Trace Log

The DIGITAL Clusters Failover Manager maintains a trace log that contains information about significant events that occur during the operation of the cluster. This information includes the following:

- DIGITAL Clusters software version number and start time
- Names of shared disks discovered by the DIGITAL Clusters software
- Messages exchanged between cluster servers to arbitrate access to shared resources
- Connection state of the remote server, as perceived by the local server
- Failover group online and offline transitions
- Cluster event delays
- Device errors on shared disks
- Cluster Administrator activity, such as the creation of new failover groups

The location of the DIGITAL Clusters Failover Manager trace log file is specified by the `FmTrace` parameter. (See the section DIGITAL Clusters Registry Key Tuning Parameters on page A-10 for more information.) By default, the file is found in the `temp` subdirectory of the cluster destination directory as specified when the cluster server software was installed. If you accepted the default directory during installation, the trace log path and file name is as follows on your system disk:

```
\Program Files\DIGITAL\Clusters\temp\fm###.log
```

where `###` is the sequential number of the trace log. (Note that a new trace log file is generated with each system reboot.)

Appendix C shows an annotated example of a typical trace log.

## Event Log

The DIGITAL Clusters software instructs the Windows NT event log service to log significant events and error conditions encountered while the software is running. You can examine this log to determine the history of the cluster operations and see how the cluster is allocating resources.

DIGITAL Clusters components that run as services put their error messages in the Application portion of the event log. These components include the following:

- DIGITAL Clusters Failover Manager
- DIGITAL Clusters Communication Infrastructure (CluCis)
- DIGITAL Clusters Failover Management Database (CFMD) Server
- DIGITAL Clusters Failover Manager Disk DLL

DIGITAL Clusters components that run as drivers put their error messages in the System portion of the event log. These components include the following:

- DIGITAL Clusters Port Driver (CluPort)
- DIGITAL Clusters Disk Driver (CluDisk)

You can read the event log using the Windows NT Event Viewer. Appendix D shows an example of a typical event log.

## Blue Screen Messages

During the early stages of a system reboot, before Windows NT has started, various messages pertinent to cluster operation are displayed on the blue screen. The messages that can be displayed are as follows:

- Cluster Adapter: \Device\ScsiPort $n$  Bus 0

This message indicates that the shared SCSI bus adapter is located on bus 0 or port  $n$ . It is a normal startup message. You can use it to verify that the cluster is using the proper adapter.

## Blue Screen Messages

- `Warning! No Cluster Adapters found`

This message indicates that no shared SCSI bus adapters have been detected during the scan of the hardware. The message is displayed if no cluster adapter was selected during software installation. It indicates that the cluster software is not controlling any of the devices. The DIGITAL Clusters Failover Manager will not start.

---

### **Caution**

---

This is a dangerous condition. The shared storage is not protected from simultaneous access by both servers.

---

This message is also displayed if the `CluPort` Registry key is missing or modified.

- `Adapter configuration has changed`

This message indicates that one or more bus adapters have been added or removed since the last time the system was booted. Choose Adapter Configuration from the Manage menu of Cluster Administrator to respecify the cluster adapter and reboot the system.

---

### **Caution**

---

This is a dangerous condition. The shared storage is not protected from simultaneous access by both servers.

---

- `Warning! Cannot attach to ScsiPortn`

This message indicates that the DIGITAL Clusters port driver encountered an error while trying to take control of the specified SCSI bus adapter. Consult the event log for more information.

---

## Registry Snapshots

This appendix contains snapshots of those parts of the Windows NT Registry that pertain to Digital Clusters for Windows NT.

### DIGITAL Clusters Failover Management Database (CFMD)

```
\registry\machine\system\currentcontrolset\services\cfmd
  Type = REG_DWORD 0x00000010
  Start = REG_DWORD 0x00000002
  ErrorControl = REG_DWORD 0x00000001
  ImagePath = REG_EXPAND_SZ C:\Program Files\Digital\Cluster\cfmdsrv.exe
  DisplayName = Cfmd Server
  ObjectName = ntsgwest\lees
  Database
    FMDisk
      _disk_343e0c4a
        Subglobal = REG_MULTI_SZ
        .Sequence = REG_DWORD 0x00000004
        DriveLetters = REG_MULTI_SZ "F:"
        Global
          Signature = REG_DWORD 0x343e0c4a
          AliasName = disk_4
          LogPartitionNumber = REG_DWORD 0x00000001
          LogPath = \_digitalclusterlog
      _disk_37a611de
        Subglobal = REG_MULTI_SZ
        .Sequence = REG_DWORD 0x00000002
        DriveLetters = REG_MULTI_SZ "U:"
        Global
          Signature = REG_DWORD 0x37a611de
          AliasName = disk_7
```

## DIGITAL Clusters Failover Management Database (CFMD)

```
_disk_3d364641
  Subglobal = REG_MULTI_SZ
  .Sequence = REG_DWORD 0x00000002
  Global
    Signature = REG_DWORD 0x3d364641
    AliasName = disk_1
_disk_49cef72c
  Subglobal = REG_MULTI_SZ
  .Sequence = REG_DWORD 0x00000002
  Global
    Signature = REG_DWORD 0x49cef72c
    AliasName = disk_2
_disk_e03e73f1
  Subglobal = REG_MULTI_SZ
  .Sequence = REG_DWORD 0x00000002
  DriveLetters = REG_MULTI_SZ "V:"
  Global
    Signature = REG_DWORD 0xe03e73f1
    AliasName = disk6
_disk_f42cb51b
  Subglobal = REG_MULTI_SZ
  .Sequence = REG_DWORD 0x00000002
  DriveLetters = REG_MULTI_SZ "T:"
  Global
    Signature = REG_DWORD 0xf42cb51b
    AliasName = disk_5
_disk_f4ab9d20
  Subglobal = REG_MULTI_SZ
  .Sequence = REG_DWORD 0x00000002
  Global
    Signature = REG_DWORD 0xf4ab9d20
    AliasName = disk_0
FMGroup
  group100
    Subglobal = REG_MULTI_SZ
    .Sequence = REG_DWORD 0x00000002
    ServerAvailability = REG_DWORD 0x00000000
    Global
      NodeList = REG_MULTI_SZ "NTCLUA1" \
        "ntclua2"
      ObjectList = REG_MULTI_SZ "FMDisk\_disk_f4ab9d20"
      Reevaluate = REG_DWORD 0x00000001
      PolicyType = REG_DWORD 0x00000001
      Comment = comment \
        generated \
        by the Cluster Administrator (aka. the UI)
```

## DIGITAL Clusters Failover Management Database (CFMD)

```
RunMeFirst = REG_DWORD 0x00000000
group101
  Subglobal = REG_MULTI_SZ
  .Sequence = REG_DWORD 0x00000002
  ServerAvailability = REG_DWORD 0x00000000
  Global
    NodeList = REG_MULTI_SZ "NTCLUA1" \
                             "ntclua2"
    ObjectList = REG_MULTI_SZ "FMDisk\_disk_3d364641"
    Reevaluate = REG_DWORD 0x00000001
    PolicyType = REG_DWORD 0x00000001
    Comment = comment \
              generated \
              by the Cluster Administrator (aka. the UI)
    RunMeFirst = REG_DWORD 0x00000000
group102
  Subglobal = REG_MULTI_SZ
  .Sequence = REG_DWORD 0x00000002
  ServerAvailability = REG_DWORD 0x00000000
  Global
    NodeList = REG_MULTI_SZ "NTCLUA1" \
                             "ntclua2"
    ObjectList = REG_MULTI_SZ "FMDisk\_disk_49cef72c"
    Reevaluate = REG_DWORD 0x00000001
    PolicyType = REG_DWORD 0x00000001
    Comment = comment \
              generated \
              by the Cluster Administrator (aka. the UI)
    RunMeFirst = REG_DWORD 0x00000000
group104
  Subglobal = REG_MULTI_SZ
  .Sequence = REG_DWORD 0x00000002
  ServerAvailability = REG_DWORD 0x00000001
  LastOnlineReason = REG_DWORD 0x00000002
  LastOfflineReason = REG_DWORD 0x00000003
  ConnectionPoint = \\NTCLUA2\group104
  Global
    NodeList = REG_MULTI_SZ "NTCLUA2" \
                             "ntclua1"
    ObjectList = REG_MULTI_SZ "FMDisk\_disk_343e0c4a"
    Reevaluate = REG_DWORD 0x00000001
    PolicyType = REG_DWORD 0x00000001
    Comment = comment \
              generated \
              by the Cluster Administrator (aka. the UI)
    RunMeFirst = REG_DWORD 0x00000001
```

## DIGITAL Clusters Failover Management Database (CFMD)

```
group105
  Subglobal = REG_MULTI_SZ
  .Sequence = REG_DWORD 0x00000002
  ServerAvailability = REG_DWORD 0x00000001
  LastOnlineReason = REG_DWORD 0x00000002
  LastOfflineReason = REG_DWORD 0x00000003
  ConnectionPoint = \\NTCLUA2\group105
  Global
    NodeList = REG_MULTI_SZ "NTCLUA2" \
                          "ntclual"
    ObjectList = REG_MULTI_SZ "FMDisk\_disk_f42cb51b"
    Reevaluate = REG_DWORD 0x00000001
    PolicyType = REG_DWORD 0x00000001
    Comment = comment \
              generated \
              by the Cluster Administrator (aka. the UI)
    RunMeFirst = REG_DWORD 0x00000000
group106
  Subglobal = REG_MULTI_SZ
  .Sequence = REG_DWORD 0x00000002
  ServerAvailability = REG_DWORD 0x00000001
  LastOnlineReason = REG_DWORD 0x00000002
  LastOfflineReason = REG_DWORD 0x00000003
  ConnectionPoint = \\NTCLUA2\group106
  Global
    NodeList = REG_MULTI_SZ "NTCLUA2" \
                          "ntclual"
    ObjectList = REG_MULTI_SZ "FMDisk\_disk_e03e73f1"
    Reevaluate = REG_DWORD 0x00000001
    PolicyType = REG_DWORD 0x00000001
    Comment = comment \
              generated \
              by the Cluster Administrator (aka. the UI)
    RunMeFirst = REG_DWORD 0x00000000
group107
  Subglobal = REG_MULTI_SZ
  .Sequence = REG_DWORD 0x00000002
  ServerAvailability = REG_DWORD 0x00000001
  LastOnlineReason = REG_DWORD 0x00000002
  LastOfflineReason = REG_DWORD 0x00000003
  ConnectionPoint = \\NTCLUA2\group107
  Global
    NodeList = REG_MULTI_SZ "NTCLUA2" \
                          "ntclual"
    ObjectList = REG_MULTI_SZ "FMDisk\_disk_37a611de"
    Reevaluate = REG_DWORD 0x00000001
```

## DIGITAL Clusters Failover Management Database (CFMD)

```
PolicyType = REG_DWORD 0x00000001
Comment = comment \
        generated \
        by the Cluster Administrator (aka. the UI)
RunMeFirst = REG_DWORD 0x00000000

FMOracle
FMScript
FMSql
FMType
    FMDisk
        .Sequence = REG_DWORD 0x00000000
        Global
        DllName = fmdisk.dll
        DependsOn =
    FMOracle
        .Sequence = REG_DWORD 0x00000000
        Global
        DllName = fmoracle.dll
        DependsOn = fmdisk
    FMScript
        .Sequence = REG_DWORD 0x00000000
        Global
        DllName = fmscript.dll
        DependsOn =
    FMSql
        .Sequence = REG_DWORD 0x00000000
        Global
        DllName = fmsql.dll
        DependsOn = fmdisk

MGEventLog
NetworkShare
Nodes
    ntclua1
    NTCLUA2
Pipe
    Sql
        .Sequence = REG_DWORD 0x00000000
        Global
        Comment = SQL Pipe
ResourceTypes
    NetworkShare
        .Sequence = REG_DWORD 0x00000000
        Global
        Comment = Exportable shares
```

## DIGITAL Clusters Failover Management Database (CFMD)

```
Nodes
    .Sequence = REG_DWORD 0x00000000
    Global
        Comment = List of cluster nodes
Pipe
    .Sequence = REG_DWORD 0x00000000
    Global
        Comment = Allow apps to failover via Named Pipes
RXACT
    Revision = REG_DWORD 0x00000001
    Initialize = REG_DWORD 0x00000001
    FMDisk
    FMGroup
    FMType
Management
    BootPort = Scsi Port 0
Security
    Security = REG_BINARY 0x000000d8
    0x80140001 0x000000c0 0x000000cc 0x00000014 0x00000034
0x00200002 0x00000001 0x00188002
    0x000f01ff 0x00000101 0x01000000 0x00000000 0x00000220
0x008c0002 0x00000005 0x00180000
    0x0002018d 0x00000101 0x01000000 0x00000000 0x00760072
0x001c0000 0x000201fd 0x00000201
    0x05000000 0x00000020 0x00000223 0x00630069 0x001c0000
0x000f01ff 0x00000201 0x05000000
    0x00000020 0x00000220 0x00630069 0x001c0000 0x000f01ff
0x00000201 0x05000000 0x00000020
    0x00000225 0x00630069 0x00180000 0x000201fd 0x00000101
0x05000000 0x00000012 0x00000225
    0x00000101 0x05000000 0x00000012 0x00000101 0x05000000
0x00000012
```

**DIGITAL Clusters Port Driver (CluPort)**

```

\registry\machine\system\currentcontrolset\services\cluport
  Type = REG_DWORD 0x00000001
  Start = REG_DWORD 0x00000000
  ErrorControl = REG_DWORD 0x00000001
  ImagePath = REG_EXPAND_SZ System32\drivers\cluport.sys
  DisplayName = Cluster Port Driver
  Group = port
  DependOnService = REG_MULTI_SZ
  DependOnGroup = REG_MULTI_SZ "SCSI miniport"
  Parameters
    Scsi
      Scsi Port 1
        Scsi Bus 0
          Initiators = REG_MULTI_SZ "NTCLUA2" \
            "ntclual"
          Name = NTCLUA2tontclualbus0
          Checksum = REG_DWORD 0x00011395
        Scsi Port 2
          Scsi Bus 0
            Initiators = REG_MULTI_SZ "NTCLUA2" \
              "ntclual"
            Name = NTCLUA2tontclualbus1
            Checksum = REG_DWORD 0x00320371
    Security
      Security = REG_BINARY 0x000000d8
        0x80140001 0x000000c0 0x000000cc 0x00000014 0x00000034
0x00200002 0x00000001 0x00188002
        0x000f01ff 0x00000101 0x01000000 0x00000000 0x00000220
0x008c0002 0x00000005 0x00180000
        0x0002018d 0x00000101 0x01000000 0x00000000 0x00650065
0x001c0000 0x000201fd 0x00000201
        0x05000000 0x00000020 0x00000223 0x00000073 0x001c0000
0x000f01ff 0x00000201 0x05000000
        0x00000020 0x00000220 0x00000073 0x001c0000 0x000f01ff
0x00000201 0x05000000 0x00000020
        0x00000225 0x00000073 0x00180000 0x000201fd 0x00000101
0x05000000 0x00000012 0x00000225
        0x00000101 0x05000000 0x00000012 0x00000101 0x05000000
0x00000012

```

## DIGITAL Clusters Disk Driver (CluDisk)

### DIGITAL Clusters Disk Driver (CluDisk)

```
\registry\machine\system\currentcontrolset\services\cludisk
  Type = REG_DWORD 0x00000001
  Start = REG_DWORD 0x00000000
  ErrorControl = REG_DWORD 0x00000001
  Tag = REG_DWORD 0x00000001
  ImagePath = System32\drivers\cludisk.sys
  DisplayName = Cluster Disk Driver
  Group = filter
  Security
    Security = REG_BINARY 0x000000d8
      0x80140001 0x000000c0 0x000000cc 0x00000014 0x00000034
0x00200002 0x00000001 0x00188002
      0x000f01ff 0x00000101 0x01000000 0x00000000 0x00000220
0x008c0002 0x00000005 0x00180000
      0x0002018d 0x00000101 0x01000000 0x00000000 0x00000220
0x001c0000 0x000201fd 0x00000201
      0x05000000 0x00000020 0x00000223 0x00740072 0x001c0000
0x000f01ff 0x00000201 0x05000000
      0x00000020 0x00000220 0x00740072 0x001c0000 0x000f01ff
0x00000201 0x05000000 0x00000020
      0x00000225 0x00740072 0x00180000 0x000201fd 0x00000101
0x05000000 0x00000012 0x00000225
      0x00000101 0x05000000 0x00000012 0x00000101 0x05000000
0x00000012
```

**DIGITAL Clusters File System (CFS)**

```

\registry\machine\system\currentcontrolset\services\cfs
  Type = REG_DWORD 0x00000002
  Start = REG_DWORD 0x00000002
  ErrorControl = REG_DWORD 0x00000001
  ImagePath = REG_EXPAND_SZ System32\drivers\cfs.sys
  DisplayName = Cluster File System
  DependOnService = REG_MULTI_SZ
  DependOnGroup = REG_MULTI_SZ "NetworkProvider"
  Linkage
    Bind = REG_MULTI_SZ "\Device\NwlnkNb" \
      "\Device\NetBT_DC21X41" \
      "\Device\Nbf_DC21X41"

  NetworkProvider
    Name = Digital Clusters for Windows NT
    ProviderPath = REG_EXPAND_SZ %SystemRoot%\System32\clunsapi.dll
  Security
    Security = REG_BINARY 0x000000d8
      0x80140001 0x000000c0 0x000000cc 0x00000014 0x00000034
0x00200002 0x00000001 0x00188002
      0x000f01ff 0x00000101 0x01000000 0x00000000 0x00000220
0x008c0002 0x00000005 0x00180000
      0x0002018d 0x00000101 0x01000000 0x00000000 0x00000220
0x001c0000 0x000201fd 0x00000201
      0x05000000 0x00000020 0x00000223 0x00000000 0x001c0000
0x000f01ff 0x00000201 0x05000000
      0x00000020 0x00000220 0x00000000 0x001c0000 0x000f01ff
0x00000201 0x05000000 0x00000020
      0x00000225 0x00000000 0x00180000 0x000201fd 0x00000101
0x05000000 0x00000012 0x00000225
      0x00000101 0x05000000 0x00000012 0x00000101 0x05000000
0x00000012

```

## DIGITAL Clusters Failover Manager

### DIGITAL Clusters Failover Manager

```
\registry\machine\system\currentcontrolset\services\ClusterFailoverManager
Type = REG_DWORD 0x00000010
Start = REG_DWORD 0x00000002
ErrorControl = REG_DWORD 0x00000001
ImagePath = REG_EXPAND_SZ C:\Program Files\Digital\Cluster\fmcore.exe
DisplayName = Cluster Failover Manager
DependOnService = REG_MULTI_SZ "Cfmd"
DependOnGroup = REG_MULTI_SZ
ObjectName = ntsqwext\lees
Parameters
    ClusterName = AlphaCluster
    FMTrace = REG_EXPAND_SZ C:\Program Files\Digital\Cluster\temp\fm###.log
    FMTraceVerbosity = REG_DWORD 0x00000003
Security
    Security = REG_BINARY 0x000000d8
        0x80140001 0x000000c0 0x000000cc 0x00000014 0x00000034
0x00200002 0x00000001 0x00188002
        0x000f01ff 0x00000101 0x01000000 0x00000000 0x00000220
0x008c0002 0x00000005 0x00180000
        0x0002018d 0x00000101 0x01000000 0x00000000 0x0015fcf8
0x001c0000 0x000201fd 0x00000201
        0x05000000 0x00000020 0x00000223 0x00000000 0x001c0000
0x000f01ff 0x00000201 0x05000000
        0x00000020 0x00000220 0x00000000 0x001c0000 0x000f01ff
0x00000201 0x05000000 0x00000020
        0x00000225 0x00000000 0x00180000 0x000201fd 0x00000101
0x05000000 0x00000012 0x00000225
        0x00000101 0x05000000 0x00000012 0x00000101 0x05000000
0x00000012
```

**DIGITAL Clusters Name Server**

```

\registry\machine\system\currentcontrolset\services\ClusterNameServer
  Type = REG_DWORD 0x00000010
  Start = REG_DWORD 0x00000002
  ErrorControl = REG_DWORD 0x00000001
  ImagePath = REG_EXPAND_SZ C:\Program Files\Digital\Cluster\cns.exe
  DisplayName = Cluster Name Service
  ObjectName = LocalSystem
  ClusterNameCache
    AlphaCluster
    Jrm4Clu
    labhxp
    ntprioris
    raw400
    will
  Security
    Security = REG_BINARY 0x000000d8
      0x80140001 0x000000c0 0x000000cc 0x00000014 0x00000034
0x00200002 0x00000001 0x00188002
      0x000f01ff 0x00000101 0x01000000 0x00000000 0x00000220
0x008c0002 0x00000005 0x00180000
      0x0002018d 0x00000101 0x01000000 0x00000000 0x00650078
0x001c0000 0x000201fd 0x00000201
      0x05000000 0x00000020 0x00000223 0x000f0000 0x001c0000
0x000f01ff 0x00000201 0x05000000
      0x00000020 0x00000220 0x000f0000 0x001c0000 0x000f01ff
0x00000201 0x05000000 0x00000020
      0x00000225 0x000f0000 0x00180000 0x000201fd 0x00000101
0x05000000 0x00000012 0x00000225
      0x00000101 0x05000000 0x00000012 0x00000101 0x05000000
0x00000012

```

## Log Watch

### Log Watch

```
\registry\machine\system\currentcontrolset\services\LogWatch
  Type = REG_DWORD 0x00000010
  Start = REG_DWORD 0x00000002
  ErrorControl = REG_DWORD 0x00000001
  ImagePath = REG_EXPAND_SZ C:\Program Files\Digital\Cluster\logwatch.exe

  DisplayName = Cluster Event Log
  ObjectName = ntsgwest\lees
  Security
    Security = REG_BINARY 0x000000d8
      0x80140001 0x000000c0 0x000000cc 0x00000014 0x00000034
0x00200002 0x00000001 0x00188002
      0x000f01ff 0x00000101 0x01000000 0x00000000 0x00000220
0x008c0002 0x00000005 0x00180000
      0x0002018d 0x00000101 0x01000000 0x00000000 0x00000101
0x001c0000 0x000201fd 0x00000201
      0x05000000 0x00000020 0x00000223 0x05000000 0x001c0000
0x000f01ff 0x00000201 0x05000000
      0x00000020 0x00000220 0x05000000 0x001c0000 0x000f01ff
0x00000201 0x05000000 0x00000020
      0x00000225 0x05000000 0x00180000 0x000201fd 0x00000101
0x05000000 0x00000012 0x00000225
      0x00000101 0x05000000 0x00000012 0x00000101 0x05000000
0x00000012
```

## SCSI Device Map

```

\registry\machine\hardware\devicemap\scsi
  Scsi Port 0
    Interrupt = REG_DWORD 0x0000000d
    IOAddress = REG_DWORD 0x00011000
    Driver = nrcrc810
    Scsi Bus 0
      Initiator Id 7
      Target Id 0
        Logical Unit Id 0
          Identifier = DEC      RZ28M      (C) DEC0568
          Type = DiskPeripheral
        Target Id 4
          Logical Unit Id 0
            Identifier = DEC      RRD45      (C) DEC 1645
            Type = CdRomPeripheral
  Scsi Port 1
    Interrupt = REG_DWORD 0x00000009
    IOAddress = REG_DWORD 0x00011100
    Driver = aic78xx
    Scsi Bus 0
      Initiator Id 6
  Scsi Port 2
    Interrupt = REG_DWORD 0x00000005
    IOAddress = REG_DWORD 0x00320000
    Driver = deckzpsx
    Scsi Bus 0
      Initiator Id 6
      Target Id 0
        Logical Unit Id 0
          Identifier = DEC      SWXRC-04      X06Z
          Type = DiskPeripheral
        Logical Unit Id 1
          Identifier = DEC      SWXRC-04      X06Z
          Type = DiskPeripheral
        Logical Unit Id 2
          Identifier = DEC      SWXRC-04      X06Z
          Type = DiskPeripheral
        Logical Unit Id 3
          Identifier = DEC      SWXRC-04      X06Z
          Type = DiskPeripheral
        Logical Unit Id 4
          Identifier = DEC      SWXRC-04      X06Z
          Type = DiskPeripheral
        Logical Unit Id 5

```

## SCSI Device Map

```
Identifier = DEC      SWXRC-04      X06Z
Type = DiskPeripheral
Logical Unit Id 6
Identifier = DEC      SWXRC-04      X06Z
Type = DiskPeripheral
Logical Unit Id 7
Identifier = DEC      SWXRC-04      X06Z
Type = DiskPeripheral
```

# C

---

## DIGITAL Clusters Failover Manager Trace Log Example

This appendix contains an annotated example of a typical DIGITAL Clusters Failover Manager trace log.

The trace log header contains useful information such as the product version, build number, cluster name, cluster server name, and log start date:

```
Digital Clusters for Windows NT(TM) 1.0-6041 BETA 2+ (Build 6041)
Digital Equipment Corporation
Windows NT(TM) is a trademark of Microsoft Corporation.
Cluster Failover Manager Trace File
Opened on cluster AlphaCluster node NTCLUA2 at 6/4/96 12:20:51 PM by
program C:\Program Files\Digital\Cluster\fmcore.exe
```

The next section logs the loading of the DLLs. There is one DLL for each type of resource managed by the DIGITAL Clusters Failover Manager. If there is a problem loading the DLL, an error will be logged.

```
12:20:51.578 tid=94 trace started
12:20:51.679 tid=94 Cfmd Server is not ready, waiting...
12:21:24.687 tid=94 Loading failover object dll fmdisk.dll...
12:21:24.835 tid=94 Disk Failover Object DLL Process Attach
12:21:24.867 tid=94 Disk Failover Object DLL Establish Linkage
12:21:24.945 tid=94 Loading failover object dll fmoracle.dll...
12:21:25.109 tid=94 Oracle Failover Object DLL Process Attach
12:21:25.171 tid=94 Oracle Failover Object DLL Establish Linkage
12:21:25.226 tid=94 Loading failover object dll fmscript.dll...
12:21:25.335 tid=94 Script Failover Object DLL Process Attach
12:21:25.390 tid=94 Script Failover Object DLL Establish Linkage
12:21:25.460 tid=94 Loading failover object dll fmsql.dll...
```

## DIGITAL Clusters Failover Manager Trace Log

```
12:21:25.562 tid=94  Sql Failover Object DLL Process Attach
12:21:25.632 tid=94  Sql Failover Object DLL Establish Linkage
```

Once the DLLs are loaded, output from one task typically is interspersed with that from others. First the existing failover groups in the database are loaded. In this example, there are seven failover groups. Note the `tid` (thread ID) field. It can be used to follow threads of activity. This is very important when multiple disks are going on line simultaneously.

```
12:21:47.140 tid=94  Created GROUP nexus "group104"
12:21:47.210 tid=94  Created GROUP nexus "group100"
12:21:47.210 tid=129 Monitoring group "group104"
12:21:47.281 tid=94  Created GROUP nexus "group101"
12:21:47.281 tid=130 Monitoring group "group100"
12:21:47.421 tid=94  Created GROUP nexus "group102"
12:21:47.421 tid=131 Monitoring group "group101"
12:21:47.554 tid=94  Created GROUP nexus "group105"
12:21:47.554 tid=132 Monitoring group "group102"
12:21:47.695 tid=94  Created GROUP nexus "group106"
12:21:47.695 tid=133 Monitoring group "group105"
12:21:47.835 tid=94  Created GROUP nexus "group107"
12:21:47.843 tid=134 Monitoring group "group106"
12:21:47.984 tid=94  Script DLL Initialize
12:21:47.984 tid=135 Monitoring group "group107"
```

In the next section, the disk DLL is initialized. Note the phases marked steps I, II, III, and IV. These are the phases of disk discovery. This process should result in disk signatures being read and event threads being created. In this example, one of the disks has a duplicate signature so it is not counted as a valid disk.

```
12:21:48.109 tid=94  Disk Failover Object DLL Initialize
12:21:54.093 tid=94  Step Ib: Forcing Scsi disk driver to discover
                    new devices
12:21:54.218 tid=94  Step II: Identifying shared devices by probing
                    Dos physical drives
12:21:54.289 tid=94  \Device\Harddisk0 is scsi address (2, 0, 0, 0)
                    product id 'DEC      SWXRC-04      X06Z'
12:21:54.351 tid=94  \Device\Harddisk1 is scsi address (2, 0, 0, 1)
                    product id 'DEC      SWXRC-04      X06Z'
12:21:54.429 tid=94  \Device\Harddisk2 is scsi address (2, 0, 0, 2)
                    product id 'DEC      SWXRC-04      X06Z'
12:21:54.507 tid=94  \Device\Harddisk3 is scsi address (2, 0, 0, 3)
                    product id 'DEC      SWXRC-04      X06Z'
12:21:54.570 tid=94  \Device\Harddisk4 is scsi address (2, 0, 0, 4)
                    product id 'DEC      SWXRC-04      X06Z'
```

## DIGITAL Clusters Failover Manager Trace Log

```
12:21:54.640 tid=94  \Device\Harddisk5 is scsi address (2, 0, 0, 5)
product id 'DEC      SWXRC-04      X06Z'
12:21:54.710 tid=94  \Device\Harddisk6 is scsi address (2, 0, 0, 6)
product id 'DEC      SWXRC-04      X06Z'
12:21:54.781 tid=94  \Device\Harddisk7 is scsi address (2, 0, 0, 7)
product id 'DEC      SWXRC-04      X06Z'
12:21:54.851 tid=94  The driver for \Device\ScsiPort2 is deckzpsx
12:21:54.859 tid=130 Waited long enough for group "group100" nexus
to come up
12:21:54.867 tid=131 Waited long enough for group "group101" nexus
to come up
12:21:54.867 tid=132 Waited long enough for group "group102" nexus
to come up
12:21:54.882 tid=129 Waited long enough for group "group104" nexus
to come up
12:21:54.882 tid=133 Waited long enough for group "group105" nexus
to come up
12:21:54.882 tid=134 Waited long enough for group "group106" nexus
to come up
12:21:54.890 tid=135 Waited long enough for group "group107" nexus
to come up
12:21:54.921 tid=94  The initiator id for \Device\ScsiPort2 is 6
12:21:55.000 tid=139 Listening for data for group "group100"
12:21:55.070 tid=62  Listening for data for group "group101"
12:21:55.140 tid=142 Listening for data for group "group102"
12:21:55.210 tid=145 Listening for data for group "group104"
12:21:55.281 tid=147 Listening for data for group "group105"
12:21:55.351 tid=149 Listening for data for group "group106"
12:21:55.421 tid=151 Listening for data for group "group107"
12:21:55.492 tid=94  Step IIIa: bus excl. prot. read partition
table/sig
12:21:56.453 tid=153 PhysicalDrive7 has signature 0x37a611de
12:21:56.546 tid=153 PhysicalDrive6 has signature 0xe03e73f1
12:21:56.640 tid=153 PhysicalDrive5 has signature 0xf42cb51b
12:21:56.742 tid=153 PhysicalDrive4 has signature 0x343e0c4a
12:21:56.867 tid=153 PhysicalDrive3 has signature 0x343e0c4a
12:21:56.953 tid=153 PhysicalDrive2 has signature 0x49cef72c
12:21:57.046 tid=153 PhysicalDrive1 has signature 0x3d364641
12:21:57.132 tid=153 PhysicalDrive0 has signature 0xf4ab9d20
12:21:57.585 tid=94  Step IV: Match physical devices with signatures
with CFMD objects
12:21:57.648 tid=154 _disk_37a611de (\\.\PhysicalDrive7): Event
thread waiting
12:21:57.648 tid=155 _disk_e03e73f1 (\\.\PhysicalDrive6): Event
thread waiting
```

## DIGITAL Clusters Failover Manager Trace Log

```
12:21:57.648 tid=156 _disk_f42cb51b (\\.\PhysicalDrive5): Event
thread waiting
12:21:57.656 tid=157 _disk_343e0c4a (\\.\PhysicalDrive4): Event
thread waiting
12:21:58.679 tid=94 warning: two disks with same signature!
12:21:58.781 tid=94 Unable to create--already exists. File:
E:\NEWBUILD\src\fm\fm disk\diskdata.c Line: 155
12:21:58.914 tid=159 _disk_49cef72c (\\.\PhysicalDrive2): Event
thread waiting
12:21:58.914 tid=160 _disk_3d364641 (\\.\PhysicalDrive1): Event
thread waiting
12:21:58.914 tid=94 Found logging disk _disk_343e0c4a partition 1
path \_digitalclusterlog
12:21:58.914 tid=153 _disk_f4ab9d20 (\\.\PhysicalDrive0): Event
thread waiting
12:21:59.187 tid=94 Disk device PhysicalDrive3 (signature
0x343e0c4a) has no corresponding entry in the cluster disk
configuration database. It will be ignored and will not be eligible
to be brought ON LINE. File: E:\NEWBUILD\src\fm\fm disk\device.c
Line: 1328
```

In the next section, the DIGITAL Clusters Failover Manager finishes initialization. In this example, neither SQL Server nor Oracle Server software is installed. If they were, any initialization errors would be logged here.

```
12:21:59.656 tid=94 ShareInitialize
12:22:00.140 tid=94 Sql DLL Initialize
12:22:00.203 tid=94 Reading sqlListHead...
12:22:00.273 tid=94 SQLDLL: Opening SC Manager...
12:22:00.335 tid=94 SQLDLL: Opening SC MS SQL Service...
12:22:00.406 tid=94 Can't open SQL service: 2010841088
Error: The specified service does not exist as an installed
service.
File: E:\NEWBUILD\src\fm\sqldll\methods.c Line: 260
12:22:00.617 tid=94 The application to be failed over is not
installed.
File: E:\NEWBUILD\src\fm\fmcore\typedata.c Line: 572
12:22:00.804 tid=94 Oracle DLL Initialize
12:22:00.875 tid=94 RegOpenKeyEx failed: 2
12:22:00.953 tid=94 The application to be failed over is not
installed.
File: E:\NEWBUILD\src\fm\fmcore\typedata.c Line: 572
12:22:01.101 tid=94 The cluster manager is operational on this
system.
File: E:\NEWBUILD\src\fm\fmcore\fmstartup.c Line: 357
```

## DIGITAL Clusters Failover Manager Trace Log

In the next section, the DIGITAL Clusters Failover Manager puts the failover groups on line. In this case, four of the failover groups are primary on this server—groups 104, 106, 105, and 107. Each group has one disk. The DIGITAL Clusters Failover Manager tells the disk DLL to put the disk on line. Because each disk is in a separate failover group, the disk-online activities occur in parallel. You can follow the progress by using the `tid` field, as follows:

Group	tid
104	146
106	150
105	148
107	152

In this example, failover group 104 contains the log disk. This group is brought on line first, and the DIGITAL Clusters Failover Manager reads the log. (See the message “Resynchronizing FM with Cfmd.”) Note also that disk F: is the log disk. (See the message “Cfmd Log Path set to F:\\_digitalclusterlog.”)

After the log disk is read, the remaining failover groups are brought on line. The steps for putting a disk on line are reflected in the log messages: starting, state change, drive letter assignment, file system structure check, success. Any errors are logged as well, along with an indication that the disk failed to go on line.

```
12:22:03.421 tid=146 Putting group "group104" Online
12:22:03.492 tid=146 Disk Failover Object DLL _disk_343e0c4a
goOnline
12:22:03.562 tid=146 Disk PhysicalDrive4 has gone from *OFFLINE* to
*ONLINE*
12:22:03.648 tid=146 Assign drive letters for disk sig 0x343e0c4a
part 4
12:22:04.195 tid=146 Disk partition \Device\Harddisk4\Partition1 was
assigned default drive letter F:.
File: E:\NEWBUILD\src\fm\fm\disk\letter.c Line: 407
12:22:04.335 tid=146 F: => \Device\Harddisk4\Partition1
12:22:04.406 tid=146 _disk_343e0c4a => \Device\Harddisk4\Partition0
12:22:04.617 tid=146 Usage (Mb): total 1000 used 15 (1.54%) free 985
(98.46%)
12:22:04.750 tid=146 Cfmd Log Path set to F:\_digitalclusterlog
12:22:07.281 tid=146 The cluster manager has put group group104 ON
LINE on this system. Reason: Failback to primary server.
File: E:\NEWBUILD\src\fm\fm\core\fmgroup.c Line: 1473
12:22:07.421 tid=146 Resynchronizing FM with Cfmd
12:22:23.984 tid=150 Putting group "group106" Online
12:22:24.046 tid=150 Disk Failover Object DLL _disk_e03e73f1
goOnline
```

## DIGITAL Clusters Failover Manager Trace Log

```
12:22:24.117 tid=150 Disk PhysicalDrive6 has gone from *OFFLINE* to
*ONLINE*
12:22:24.226 tid=150 Assign drive letters for disk sig 0xe03e73f1
part 4
12:22:24.437 tid=150 V: => \Device\Harddisk6\Partition1
12:22:24.507 tid=150 _disk_e03e73f1 => \Device\Harddisk6\Partition0
12:22:24.726 tid=150 Usage (Mb): total 1000 used 255 (25.48%) free
745 (74.52%)
12:22:25.476 tid=152 Putting group "group107" Online
12:22:25.601 tid=152 Disk Failover Object DLL _disk_37a611de
goOnline
12:22:25.625 tid=148 Putting group "group105" Online
12:22:25.750 tid=152 Disk PhysicalDrive7 has gone from *OFFLINE* to
*ONLINE*
12:22:25.882 tid=148 Disk Failover Object DLL _disk_f42cb51b
goOnline
12:22:26.046 tid=152 Assign drive letters for disk sig 0x37a611de
part 4
12:22:26.976 tid=152 U: => \Device\Harddisk7\Partition1
12:22:27.085 tid=152 _disk_37a611de => \Device\Harddisk7\Partition0
12:22:27.382 tid=152 Usage (Mb): total 1000 used 195 (19.52%) free
805 (80.48%)
12:22:27.937 tid=150 The cluster manager has put group group106 ON
LINE on this system. Reason: Failback to primary server.
File: E:\NEWBUILD\src\fm\fmcore\fmgroup.c Line: 1473
12:22:28.898 tid=148 Disk PhysicalDrive5 has gone from *OFFLINE* to
*ONLINE*
12:22:29.000 tid=148 Assign drive letters for disk sig 0xf42cb51b
part 4
12:22:29.945 tid=148 T: => \Device\Harddisk5\Partition1
12:22:30.078 tid=148 _disk_f42cb51b => \Device\Harddisk5\Partition0
12:22:30.414 tid=148 Usage (Mb): total 1000 used 81 (8.15%) free 919
(91.85%)
12:22:31.656 tid=152 The cluster manager has put group group107 ON
LINE on this system. Reason: Failback to primary server.
File: E:\NEWBUILD\src\fm\fmcore\fmgroup.c Line: 1473
12:22:33.507 tid=148 The cluster manager has put group group105 ON
LINE on this system. Reason: Failback to primary server.
File: E:\NEWBUILD\src\fm\fmcore\fmgroup.c Line: 1473
```

---

## DIGITAL Clusters Event Logs

This appendix contains an example of typical DIGITAL Clusters for Windows NT system and application event logs.

### System Event Log

```

6/4/96 12:21:54 PM  hszdisk      Error           None           1             N/A           NTCLUA2
    HSZDISK failed to find any devices.
6/4/96 12:20:34 PM  hszdisk      Information     None           0             N/A           NTCLUA2
    HSZDISK, SCSI Disk class driver exclusively for use with Digital
    Equipment Corporation's family of RAID storage controllers, has
    successfully loaded.
6/4/96 12:20:34 PM  CluPort      Information     None           514          N/A           NTCLUA2
    Device: The Cluster Port Driver has attached a filter device to port
    \Device\ScsiPort2.
6/4/96 12:20:34 PM  CluPort      Information     None           514          N/A           NTCLUA2
    Device: The Cluster Port Driver has attached a filter device to port
    \Device\ScsiPort1.
6/4/96 12:20:31 PM  Dhcp         Error           None           1004         N/A           NTCLUA2
    DHCP IP address lease 16.64.48.15 for the card with network address
    0000F820442E has been denied.
6/4/96 12:20:22 PM  EventLog     Information     None           6005         N/A           NTCLUA2
    The Event log service was started.

```

## Application Event Log

### Application Event Log

```
6/4/96 12:22:33 PM Failover Manager Information None 515 N/A NTCLUA2
The cluster manager has put group group105 ON LINE on this system.
Reason: Failback to primary server.
6/4/96 12:22:31 PM Failover Manager Information None 515 N/A NTCLUA2
The cluster manager has put group group107 ON LINE on this system.
Reason: Failback to primary server.
6/4/96 12:22:27 PM Failover Manager Information None 515 N/A NTCLUA2
The cluster manager has put group group106 ON LINE on this system.
Reason: Failback to primary server.
6/4/96 12:22:25 PM Failover Manager Information None 516 N/A NTCLUA1
The cluster manager has put group group107 OFF LINE on this system.
Reason: Failback to primary server.
6/4/96 12:22:25 PM Failover Manager Information None 516 N/A NTCLUA1
The cluster manager has put group group105 OFF LINE on this system.
Reason: Failback to primary server.
6/4/96 12:22:24 PM Failover Manager Information None 516 N/A NTCLUA1
The cluster manager has put group group106 OFF LINE on this system.
Reason: Failback to primary server.
6/4/96 12:22:07 PM Failover Manager Information None 515 N/A NTCLUA2
The cluster manager has put group group104 ON LINE on this system.
Reason: Failback to primary server.
6/4/96 12:22:04 PM FM Disk DLL Information None 516 N/A NTCLUA2
Disk partition \Device\Harddisk4\Partition1 was assigned default drive
letter F:.
6/4/96 12:22:03 PM Failover Manager Information None 516 N/A NTCLUA1
The cluster manager has put group group104 OFF LINE on this system.
Reason: Failback to primary server.
6/4/96 12:22:01 PM Failover Manager Information None 513 N/A NTCLUA2
The cluster manager is operational on this system.
6/4/96 12:21:59 PM FM Disk DLL Warning None 790 N/A NTCLUA2
Disk device PhysicalDrive3 (signature 0x343e0c4a) has no corresponding
entry in the cluster disk configuration database. It will be ignored and
will not be eligible to be brought ON LINE.
6/4/96 12:21:58 PM Failover Manager Error None 1024 N/A NTCLUA2
The Failover Manager encountered an unexpected error: Unable to create--
already exists.

6/4/96 12:21:38 PM CluCis Warning None 304 N/A NTCLUA2
Listener - refusing connection from "NTCLUA1 \xDA".
6/4/96 12:21:23 PM CluCis Warning None 304 N/A NTCLUA2
Listener - refusing connection from "NTCLUA1 \xDA".
6/4/96 12:21:08 PM CluCis Warning None 304 N/A NTCLUA2
Listener - refusing connection from "NTCLUA1 \xDA".
```

## Application Event Log

6/4/96 12:20:54 PM Failover Manager Information None 515 N/A NTCLUA1  
The cluster manager has put group group105 ON LINE on this system.  
Reason: Automatic Failover.

6/4/96 12:20:54 PM Failover Manager Information None 515 N/A NTCLUA1  
The cluster manager has put group group106 ON LINE on this system.  
Reason: Automatic Failover.

6/4/96 12:20:54 PM Failover Manager Information None 515 N/A NTCLUA1  
The cluster manager has put group group107 ON LINE on this system.  
Reason: Automatic Failover.

6/4/96 12:20:54 PM Failover Manager Information None 515 N/A NTCLUA1  
The cluster manager has put group group104 ON LINE on this system.  
Reason: Automatic Failover.

6/4/96 12:20:54 PM Failover Manager Information None 517 N/A NTCLUA1  
This system has lost connectivity to node ntclua2. Failover policy will  
be invoked.

6/4/96 12:20:54 PM Failover Manager Information None 517 N/A NTCLUA1  
This system has lost connectivity to node ntclua2. Failover policy will  
be invoked.

6/4/96 12:20:54 PM Failover Manager Information None 517 N/A NTCLUA1  
This system has lost connectivity to node ntclua2. Failover policy will  
be invoked.

6/4/96 12:20:54 PM Failover Manager Information None 517 N/A NTCLUA1  
This system has lost connectivity to node ntclua2. Failover policy will  
be invoked.

6/4/96 12:20:54 PM Failover Manager Information None 517 N/A NTCLUA1  
This system has lost connectivity to node ntclua2. Failover policy will  
be invoked.

6/4/96 12:20:54 PM CluCis Warning None 304 N/A NTCLUA2  
Listener - refusing connection from "NTCLUA1 \xDA".

6/4/96 12:18:53 PM Failover Manager Error None 1024 N/A NTCLUA2  
The Failover Manager encountered an unexpected error: The RPC server is  
not listening.

6/4/96 12:18:52 PM Failover Manager Information None 514 N/A NTCLUA2  
The cluster manager has received a request to stop.

6/4/96 12:17:47 PM Failover Manager Information None 515 N/A NTCLUA2  
The cluster manager has put group group107 ON LINE on this system.  
Reason: System Startup.

6/4/96 12:17:36 PM Failover Manager Information None 515 N/A NTCLUA2  
The cluster manager has put group group106 ON LINE on this system.  
Reason: System Startup.

6/4/96 12:17:20 PM Failover Manager Information None 515 N/A NTCLUA2  
The cluster manager has put group group105 ON LINE on this system.  
Reason: System Startup.

## Application Event Log

6/4/96 12:16:47 PM Failover Manager Information None 515 N/A NTCLUA2  
The cluster manager has put group group104 ON LINE on this system.  
Reason: System Startup.

6/4/96 12:16:43 PM FM Disk DLL Information None 516 N/A NTCLUA2  
Disk partition \Device\Harddisk4\Partition1 was assigned default drive  
letter F:.

6/4/96 12:16:16 PM Failover Manager Information None 515 N/A NTCLUA1  
The cluster manager has put group group102 ON LINE on this system.  
Reason: System Startup.

6/4/96 12:15:47 PM Failover Manager Information None 515 N/A NTCLUA1  
The cluster manager has put group group101 ON LINE on this system.  
Reason: System Startup.

6/4/96 12:15:31 PM Failover Manager Information None 515 N/A NTCLUA1  
The cluster manager has put group group100 ON LINE on this system.  
Reason: System Startup.

6/4/96 12:15:11 PM CluCis Warning None 304 N/A NTCLUA2  
Listener - refusing connection from "NTCLUA1 \xDA".

6/4/96 12:14:34 PM Failover Manager Information None 513 N/A NTCLUA2  
The cluster manager is operational on this system.

6/4/96 12:14:33 PM FM Disk DLL Warning None 790 N/A NTCLUA2  
Disk device PhysicalDrive3 (signature 0x343e0c4a) has no corresponding  
entry in the cluster disk configuration database. It will be ignored and  
will not be eligible to be brought ON LINE.

6/4/96 12:14:31 PM Failover Manager Error None 1024 N/A NTCLUA2  
The Failover Manager encountered an unexpected error: Unable to create--  
already exists.

6/4/96 12:14:25 PM Failover Manager Information None 513 N/A NTCLUA1  
The cluster manager is operational on this system.

6/4/96 12:14:25 PM FM Disk DLL Warning None 790 N/A NTCLUA1  
Disk device PhysicalDrive3 (signature 0x343e0c4a) has no corresponding  
entry in the cluster disk configuration database. It will be ignored and  
will not be eligible to be brought ON LINE.

6/4/96 12:14:22 PM FM Disk DLL Information None 517 N/A NTCLUA1  
Disk configuration database entry \_disk\_f4ab9d20 created.

6/4/96 12:14:19 PM FM Disk DLL Information None 517 N/A NTCLUA1  
Disk configuration database entry \_disk\_3d364641 created.

6/4/96 12:14:17 PM FM Disk DLL Information None 517 N/A NTCLUA1  
Disk configuration database entry \_disk\_49cef72c created.

6/4/96 12:14:14 PM Failover Manager Error None 1024 N/A NTCLUA1  
The Failover Manager encountered an unexpected error: Unable to create--  
already exists.

6/4/96 12:14:13 PM FM Disk DLL Information None 517 N/A NTCLUA1  
Disk configuration database entry \_disk\_343e0c4a created.

## Application Event Log

```
6/4/96 12:14:10 PM FM Disk DLL      Information None 517 N/A NTCLUA1
    Disk configuration database entry _disk_f42cb51b created.
6/4/96 12:14:08 PM FM Disk DLL      Information None 517 N/A NTCLUA1
    Disk configuration database entry _disk_e03e73f1 created.
6/4/96 12:14:06 PM FM Disk DLL      Information None 517 N/A NTCLUA1
    Disk configuration database entry _disk_37a611de created.
```

---

# Glossary

## **adapter**

An integrated circuit expansion board that communicates with and controls a device or system.

## **bus**

A collection of wires in a cable or copper traces on a circuit board used to transmit data, status, and control signals. Examples of PC buses include EISA, ISA, PCI, and SCSI.

## **bus slots**

Connectors inside the computer that are used for attaching add-on cards and devices to a bus. Also known as expansion slots.

## **cluster**

A loosely coupled set of systems that is addressed and managed like a single system, but provides high levels of availability in the event of a failure through redundant CPUs, storage, and data paths.

## **cluster alias**

A common cluster name used by supported end-user clients to access the cluster.

## **cluster IP address**

An Internet Protocol (IP) address that can fail over from one cluster server to the other. To specify a cluster IP address, you create an **IP failover object** in Cluster Administrator. Cluster IP addresses must be unique; you cannot use the IP addresses of the individual cluster servers.

**cluster member**

One of the server systems in a cluster configuration. End-user clients are not members of the cluster.

**cluster share**

A Windows NT file share contained on a shared cluster disk.

**database failover**

Pertains to a database failing over instead of the entire application.

**device driver**

A software module that provides an interface for communication between the operating system and system hardware, for example, a SCSI controller.

**differential**

A SCSI bus transmission method in which each signal is sent on two wires. The signal is derived by taking the difference in voltage between the two wires, effectively eliminating unwanted noise in the wire. *See also* **single-ended**.

**failback**

The automatic migration of failover groups from the alternate server to the primary server after the primary server that caused an initial failover returns to operational status.

**failover**

In a cluster system failure, the relocation of cluster services (such as applications) or resources (such as a cluster share) to end-user clients using backup paths.

**failover group**

A logical group of failover objects. The groups are typically made up of storage devices and applications. For example Microsoft SQL Server and the disks used to store the SQL Server database would be a logical failover group. A group can also be made up of one or more disks without an associated application. *See also* **failover object**.

**failover object**

Any cluster service or resource for which you want to ensure availability in the event of a system failure. Examples of failover objects include disks, database applications such as Microsoft SQL Server and Oracle7 Workgroup Server, and any application that can be launched and shut down in a script. Failover objects can be collected into failover groups. *See also* **failover group**.

**failover policy**

The plan of action the cluster software follows for a failover group. Each failover group is associated with a failover policy that is defined using Cluster Administrator. *See also failover group.*

**failover server**

The cluster server that functions as a backup path for the primary server for a given set of cluster services or resources. In the case of a failure, the failover server assumes responsibility for relocated cluster services or resources based on the failover policy. A failover server is itself a primary server for a given set of cluster services or resources. *See also failover group, failover policy, primary server.*

**fast SCSI**

A SCSI-2 transfer mode that operates at 10 MB/second, twice as fast as regular SCSI. *See also SCSI.*

**fast wide SCSI**

Wide SCSI operating at twice the rate of regular wide SCSI. *See also SCSI.*

**fault tolerance**

A method of ensuring the availability of a computing environment by using a backup system that mirrors the primary system. The backup system is typically called a “hot standby.” The backup system does not provide any additional computing capacity; it is available only for use in the event of a failure of the primary system. For this reason, it is a costly method of ensuring availability.

**IP failover object**

A **cluster IP address**, subnet mask, and a specified set of one or more network adapters on each cluster server that can enable the address. You can use IP failover objects to fail over IP socket-based applications on a cluster.

**IRQ**

Interrupt Request. A signal used by devices to indicate that they need attention from the CPU. Computers have several IRQ channels so that many devices can be attached, each one to its own IRQ, and serviced by the CPU.

**jumper**

A small plastic and metal connector used to bridge the gap between two or more pins. Jumpers are commonly used for configuring devices and adapter cards.

**MIB**

Management information base. SNMP defines a set of variables that the host must keep. Because different network-management services are used for different types of devices or for different network-management protocols, each service has its own set of objects. The entire set of objects that any network-management service or protocol uses is referred to as its MIB. *See also* **SNMP**.

**Microsoft SQL Server database failover**

Refers to an SQL Server database failing over, not the entire application. Upon primary server failure, the failover server will service the databases on the shared disks.

**named pipe**

A network transport used for interprocess communication.

**name server**

Software installed on the cluster servers that works with the client software to create the illusion of a single system through aliases. Using aliases, the client is unaware of the name of each server or how the cluster work load is distributed.

**noise**

Unwanted and usually interfering electrical signals.

**NTFS**

NT File System. The standard file system for the Microsoft Windows NT operating system.

**Oracle Server database failover**

Refers to an Oracle Server instance failing over. Upon primary server failure, the failover server will service the instance database on the shared disks.

**primary server**

The preferred server through which a given set of cluster resources or services are made available.

**RAID**

Redundant Array of Independent Disks. A collection of storage devices configured to provide higher data transfer rates, data recovery capability, or both.

**redundancy**

A method of protecting against failures by building in extra, backup components to a system.

**regular SCSI**

8-bit SCSI. *See also* **SCSI**.

**resource-based failover policy**

A failover policy in which only those resources impacted by a particular failure are failed over. DIGITAL Clusters for Windows NT employs a resource-based failover policy. *See also* **failover policy** and **server-based failover policy**.

**scaleable**

In a system, the ability to add additional capacity as needs require.

**SCSI**

Small Computer System Interface. An intelligent bus protocol for transmitting data and commands between a variety of devices. There are many implementations of SCSI, including fast SCSI, wide SCSI, and fast wide SCSI.

**SCSI-2**

The second generation of SCSI; includes many improvements to SCSI-1, including fast SCSI, wide SCSI, and mandatory parity checking. *See also* **SCSI**.

**SCSI-3**

The third generation of SCSI; introduces improvements to the parallel bus and high-speed bus architectures. *See also* **SCSI**.

**SCSI ID**

A number used on SCSI devices to uniquely identify them among other devices on the bus. *See also* **SCSI**.

**server-based failover policy**

A failover policy in which all system services are failed over in the event of a nonrecoverable system failure. Server-based failover policies are very limiting, because even those resources that were unaffected by the failure are relocated to the backup server. *See also* **failover policy** and **resource-based failover policy**.

**single-ended**

A SCSI bus transmission method in which each signal is carried by a single wire. Single-ended buses are more susceptible to noise than differential buses. *See also* **differential**.

**SMB**

Server Message Block. A protocol that allows a set of computers to access shared resources as if they were local.

**SNMP**

Simple Network Management Protocol. SNMP is a network management protocol widely used in TCP/IP networks, and, more recently, with Internet Packet Exchange (IPX) networks. SNMP transports management information and commands between a management system (such as ServerWORKS™) and an SNMP agent.

**symmetric multiprocessing (SMP)**

A method of adding computing capacity to a system by adding processors.

**terminator**

An electrical circuit attached to each end of a SCSI bus to minimize signal reflections and extraneous noise.

**UNC**

Universal naming convention for the LANMAN (LAN Manager) protocol. This is the traditional `\\server\share` syntax.

**wide SCSI**

A SCSI-2 bus that is 16 or 32 bits wide. Regular SCSI is 8 bits wide. *See also* **SCSI**.

**workload balancing**

The ability to partition the work load by assigning resources (such as disks and databases) to cluster servers and defining failover and failback policies. This offers an efficient and effective use of server resources.

---

# Index

## A

---

### Adapters

- adding, 8–2
- changing, 8–2

### Adapters list for IP failover

- specifying, 9–14, 9–15

### Alias, cluster

- definition, 1–5

### Application event log , DIGITAL Clusters

(*example*), D–2 to D–5

## C

---

### CFMD Registry key, A–8

### CfmdTrace Registry tuning parameter, A–10

### CFS Registry key, A–9

### CisTrace Registry tuning parameter, A–11

### Client tuning parameters

- Windows 95, A–23
- Windows for Workgroups, A–25
- Windows NT, A–24

### CluDisk Registry key, A–9

### CLUIVP utility

- using to verify your cluster installation, A–4

### CluPort Registry key, A–8

### Cluster

- definition, 1–3

### Cluster Administrator

- displaying the cluster topology, 6–2 to 6–7
  - Class View, 6–5

- Cluster View, 6–4

- System View, 6–3

- quitting, 6–2

- starting, 6–2

### Cluster alias, definition, 1–5

### Cluster buses

- displaying, 8–4

### Cluster components that log events, 8–8

### Cluster concepts, 2–1 to 2–5

### Cluster disk alias

- changing, 8–5

- creating, 8–5

- default, 8–5

### Cluster disk names

- displaying, 8–6

### Cluster members, definition, 1–7

### Cluster MIB (management information base), 1–9

### Cluster Monitor utility, A–5

### Cluster topology

- displaying, 6–2 to 6–7

### ClusterFailoverManager Registry key, A–9

### ClusterName Registry tuning parameter, A–16

### ClusterNameService Registry key, A–10

### CLUXFER utility, A–4

### Command execution

- starting with a script failover object, 9–7

- stopping with a script failover object, 9–8

### Command prompt commands

- NET SHARE, A–2

- NET VIEW, A–2

ConnectionTimeout Registry tuning parameter, A-12

Controllers

moving, 8-4

Creating

failover group, 9-17

IP failover object, 9-13

Oracle failover object, 9-2

script failover object, 9-7

## D

---

Database application

starting, 9-6

stopping, 9-6

Database failover (*example*), 2-3

Database software, configuring for failover, 3-1 to 3-25

Deleting

IP failover object, 9-16

Oracle failover object, 9-6

script failover object, 9-9

SQL Server failover object, 9-11

DIGITAL Clusters

benefits, 1-5 to 1-7

high availability, 1-5

scalability, 1-6

supports commodity hardware, 1-6

supports industry standards, 1-6

client tuning parameters

Windows 95, A-23

Windows for Workgroups, A-25

Windows NT, A-24

comparison with fault-tolerant systems, 1-6

configuration (*example*), 1-4

function in a PC LAN, 1-7

physical connections, 1-9

resources, management of, 1-9

services, management of, 1-9

supported clients, 1-8

supported server architectures, 1-8

product definition, 1-3

product overview, 1-1 to 1-5

DIGITAL Clusters Disk Driver (CluDisk) Registry snapshot (*example*), B-8

DIGITAL Clusters event logs, D-1 to D-5

application event log (*example*), D-2 to D-5

system event log (*example*), D-1

DIGITAL Clusters Failover Management Database (CFMD) Registry snapshot (*example*), B-1 to B-6

DIGITAL Clusters Failover Manager

repeated failover threshold, 8-10

DIGITAL Clusters Failover Manager Registry snapshot (*example*), B-10

DIGITAL Clusters Failover Manager trace log, A-26

(*example*), C-1 to C-6

DIGITAL Clusters File System (CFS) Registry snapshot (*example*), B-9

DIGITAL Clusters log disk

choosing a failover group for, 8-8

managing, 8-8

DIGITAL Clusters Name Server Registry snapshot (*example*), B-11

DIGITAL Clusters Port Driver (CluPort) Registry snapshot (*example*), B-7

DisableFailoverNetDelay Registry tuning parameter, A-16

Disabling failover for a group, 9-23

Disk Administrator, Windows NT

using to assign fixed drive letters to partitions, A-1

using to format and partition disks, A-1

DiskArbitrationInterval Registry tuning parameter, A-17

DiskErrorSeparation Registry tuning parameter, A-19

DiskErrorThreshold Registry tuning parameter, A-18

Domain names

for web servers, 5-1

Domino server, 4-2

Control Panel settings for, 4-5

creating failover groups, 4-5

creating failover scripts, 4-19

creating Windows NT services, 4-17

installing partitioned servers, 4-5

IP failover

partitioned servers, 4-3

single server, 4-2

using TCP/IP domain names, 4-5

web server support, 4-20

## E

---

- Event log, Windows NT
  - logging of cluster events, A-27
  - managing, 8-6
- Executing a command using a script failover object, 9-6

## F

---

- Failback
  - definition, 2-3
- Failover
  - definition, 2-1
  - preventing during routine maintenance, 8-9
- Failover group
  - available objects, 9-22
  - contents, 9-18, 9-22
  - creating, 9-17
  - definition, 2-2
  - disabling failover, 9-23
  - modifying behavior, 9-22
  - modifying contents, 9-21
  - name, 9-18, 9-22
  - objects, 9-18
- Failover object
  - definition, 2-2
- Failover policy
  - definition, 2-2
- Failover server
  - specifying for a failover group, 9-19
- Failover threshold, repeated, of DIGITAL Clusters
  - Failover Manager, 8-10
- Failover, database (*example*), 2-3
- FailoverEvaluateDelay Registry tuning parameter, A-13
- Filters for Windows NT event log, 8-7
- FmLogLevel Registry tuning parameter, A-21
- FmTrace Registry tuning parameter, A-20
- FmTraceOutput Registry tuning parameter, A-19
- FmTraceVerbosity Registry tuning parameter, A-22

## I

---

- IIS. *See* Microsoft IIS, Web servers
- IP failover object
  - adapters list

- specifying, 9-14, 9-15
- adding to a failover group, 9-12
- creating, 9-13
- deleting, 9-16
- modifying, 9-14
- name
  - specifying, 9-14
- subnet mask
  - modifying, 9-15
  - specifying, 9-14

## L

---

- Load balancing, 9-20, 9-24
- Local buses
  - displaying, 8-4
- Log disk
  - choosing a failover group for, 8-8
  - managing, 8-8
- Log Watch Registry snapshot (*example*), B-12
- Logical port numbers
  - assigning to adapters, 8-2
- LogLevel Registry tuning parameter, A-15
- LogWatch Registry key, A-9
- Lotus Notes. *See also* Domino server
  - directory pointer files, 4-23
    - creating, 4-30
  - Domino 4.5 server, 4-2
  - Domino server
    - clustering methods, 4-2
    - creating failover scripts, 4-19
    - installing, 4-5
    - web server support, 4-20
  - Notes 4.11 server, 4-21
    - clustering methods, 4-21
    - failover method, 4-23
    - failover scripts, 4-32
    - installing, 4-23
  - overview, 4-1

## M

---

- Managing
  - adapter configuration, 8-2
  - disk aliases, 8-5
  - event log, 8-6
  - manual failover, 8-9

- SQL Server databases, 8–11
- Manual failover, managing, 8–9
- Members, cluster
  - definition, 1–7
- MIB (management information base), cluster, 1–9
- Microsoft IIS, 5–3
  - configuring, 5–4
  - installing, 5–3
  - sample failover scripts, 5–4
- Modifying
  - failover group contents, 9–21
  - IP failover object, 9–14
  - Oracle failover object, 9–4
  - script failover object, 9–9

## N

---

- NET SEND command, 9–6
- NET SHARE command
  - using to verify exported file shares, A–2
- NET VIEW command
  - using to verify file shares, A–2
- NETMON network monitor utility
  - using to troubleshoot cluster problems, A–3
- Netscape Enterprise Server
  - installing, 5–6, 5–7
  - sample scripts, 5–7
- Network monitor utility, NETMON
  - using to troubleshoot cluster problems, A–3

## O

---

- Oracle Server
  - database, specifying, 9–3
  - failover object
    - comment field, specifying, 9–3
    - creating, 9–2
    - deleting, 9–6
    - modifying, 9–4
    - name, specifying, 9–3
    - user name, specifying, 9–3
  - installing and configuring for failover, 3–14 to 3–25
  - listener name, specifying, 9–4
  - password, specifying, 9–3
  - pipe name, specifying, 9–3
  - sid parameter file, specifying, 9–3

- sid, specifying, 9–3
- user type, specifying, 9–3

## P

---

- Password
  - Oracle, specifying, 9–3
  - SQL Server failover object, modifying, 9–11
- Primary server
  - specifying for a failover group, 9–19, 9–22

## R

---

- ReconnectWait Registry tuning parameter, A–14
- regedt32, Windows NT Registry editor, A–1
- Registry keys
  - cluster tuning parameters, A–10–A–23
    - CfmdTrace, A–10
    - CisTrace, A–11
    - ClusterName, A–16
    - ConnectionTimeout, A–12
    - DisableFailoverNetDelay, A–16
    - DiskArbitrationInterval, A–17
    - DiskErrorSeparation, A–19
    - DiskErrorThreshold, A–18
    - FailoverEvaluateDelay, A–13
    - FmLogLevel, A–21
    - FmTrace, A–20
    - FmTraceOutput, A–19
    - FmTraceVerbosity, A–22
    - LogLevel, A–15
    - ReconnectWait, A–14
  - cluster-specific, A–8–A–10
    - CFMD key, A–8
    - CFS, A–9
    - CluDisk, A–9
    - CluPort, A–8
    - ClusterFailoverManager, A–9
    - ClusterNameService, A–10
    - LogWatch, A–9
    - SCSI, A–10
- Registry snapshots, B–1–B–14
  - DIGITAL Clusters Disk Driver (CluDisk) (*example*), B–8
  - DIGITAL Clusters Failover Management Database (CFMD) (*example*), B–1 to B–6

- DIGITAL Clusters Failover Manager  
(*example*), B-10
- DIGITAL Clusters File System (CFS)  
(*example*), B-9
- DIGITAL Clusters Name Server (*example*),  
B-11
- DIGITAL Clusters Port Driver (CluPort)  
(*example*), B-7
- Log Watch (*example*), B-12
- SCSI Device Map (*example*), B-13 to B-14
- Registry, Windows NT
  - using to find cluster information, A-8 to  
A-10
- Restrictions on script failover object command,  
9-8

## S

---

- Script failover object
  - command restrictions, 9-8
  - creating, 9-7
  - deleting, 9-9
  - modifying, 9-9
  - using, 9-6
- SCSI adapter information
  - displaying, 8-4
- SCSI device map Registry key, A-10
- SCSI Device Map Registry snapshot (*example*),  
B-13 to B-14
- ServerWORKS, 1-7
- Services applet, Windows NT
  - using to verify cluster services, A-2
- Simple Network Management Protocol (SNMP)
  - agent, 1-7
  - management tools
    - ServerWORKS, 1-7
- SMS (System Management Server) network
  - monitor program, NETMON, A-3
- SNMP (Simple Network Management Protocol)
  - agent, 1-7
  - management tools
    - ServerWORKS, 1-7
- SQL Server
  - databases, managing, 8-11
  - failover object
    - comment, 9-11
    - deleting, 9-11

- name, 9-11
  - password
  - modifying, 9-11
  - installing and configuring for failover,  
3-1 to 3-14
- Start command execution
  - using a script failover object, 9-7
- Stop command execution
  - using a script failover object, 9-8
- Subnet mask of IP failover address, 9-15
- Symmetric multiprocessing (SMP) processor, 1-3
- Synchronizing of databases by log disk, 8-8
- System event log, DIGITAL Clusters (*example*), D-  
1

## T

---

- Trace log, DIGITAL Clusters Failover Manager,  
A-26
- Trace log, DIGITAL Clusters Failover Manager  
(*example*), C-1
- Troubleshooting, 11-1 to 11-9
  - client problems, 11-5 to 11-7
  - configuration problems, 11-1 to 11-4
  - database problems, 11-7 to 11-9
  - failover problems, 11-4 to 11-5
  - tools and resources, A-1 to A-28
    - blue screen messages, A-27 to A-28
    - DIGITAL Clusters Failover Manager
      - trace log, A-26
    - Registry, A-8 to A-10
    - Registry key tuning parameters,  
A-10 to A-23
    - software utilities, A-1 to A-8

## U

---

- Utilities
  - CLUIVP, A-4
  - Cluster Monitor, A-5
  - CLUXFER, A-4
  - NETMON network monitor, A-3

## V

---

- Views, cluster, displaying
  - Class View, 6-5

Cluster View, 6-4  
System View, 6-3

## **W**

---

### Web servers

configuring for failover, 5-1  
overview, 5-1

Domino server, 4-20

Microsoft IIS, 5-3

Netscape Enterprise Server, 5-5  
using multiple web servers, 5-2

### Windows 95 clients

tuning parameters, A-23

Windows for Workgroups clients  
tuning parameters, A-25

### Windows NT clients

tuning parameters, A-24

### Windows NT event log

filters for, 8-7

logging of cluster events, A-27

managing, 8-6

### Windows NT Services applet

using to verify cluster services, A-2

### Windows NT system device location

displaying, 8-4