Digital Clusters for Windows NT

Administrator's Guide

Order Number: AA–QVUTA–TE

June 1996

This guide provides a conceptual overview of the Digital Clusters for Windows NT™ product, describes key features and benefits of clusters, and presents procedures on using Cluster Administrator to perform routine cluster administration tasks.

Revision/Update Information:	This is a new manual.
Operating System and Version:	Microsoft® Windows NT Version 3.51 with Service Pack 4
Software Version:	Digital Clusters for Windows NT Version 1.0

Digital Equipment Corporation
Maynard, Massachusetts

June 1996

Contents

ABOUT THIS GUIDE

CHAPTER 1 INTRODUCTION

CHAPTER 2 CLUSTER ADMINISTRATOR OVERVIEW

CHAPTER 3 CONFIGURING DATABASE SOFTWARE FOR FAILOVER

CHAPTER 4 GETTING STARTED WITH CLUSTER ADMINISTRATOR

CHAPTER 5 CONFIGURING YOUR CLUSTER

CHAPTER 6 MANAGING A CLUSTER

CHAPTER 7 WORKING WITH FAILOVER OBJECTS AND GROUPS

CHAPTER 8 APPLICATION CONSIDERATIONS

CHAPTER 9 TROUBLESHOOTING

APPENDIX A TROUBLESHOOTING TOOLS AND RESOURCES

APPENDIX B REGISTRY SNAPSHOT

APPENDIX C FAILOVER MANAGER TRACE LOG

APPENDIX D EVENT LOGS

GLOSSARY

FIGURES

About This Guide

This guide provides a conceptual overview of the Digital Clusters for Windows NT product, describes key features and benefits of clusters, and gives step-by-step instructions on how to use Cluster Administrator to perform routine cluster administration tasks.

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 nine 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	Introduces Cluster Administrator and discusses failover objects, failover groups, and the concepts of failover and failback.
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™ Workgroup Server.
Chapter 4	Gives step-by-step instructions on starting and quiting Cluster Administrator and displaying the cluster topology.

Chapter 5	Outlines the steps you need to perform to complete your initial cluster configuration.
Chapter 6	Presents step-by-step procedures on how to manage a cluster using Cluster Administrator.
Chapter 7	Describes how to create, modify, and delete failover groups, Oracle failover objects, script failover objects, and SQL failover objects using Cluster Administrator.
Chapter 8	Discusses client application failover considerations. Specific examples of client applications are included.
Chapter 9	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 for Windows NT Failover Manager trace log.
Appendix D	Contains an example of a typical Digital Clusters for Windows NT event log.

Conventions

The following conventions are used in this guide:

Convention	Meaning
Bold	Bold type indicates the actual commands, words, or characters that you type in a dialog box or at the command prompt.
Italic	Italic type indicates a placeholder for information on parameters that you must provide. For example, if the procedure asks you to type *filename*, 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.
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

Chapter 1
Introduction

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 has begun to deliver 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.

The first version of Digital Clusters for Windows NT focuses on minimizing the down time caused by software, network and sytem failures, and:

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

What Is Digital Clusters for Windows NT?

Digital Clusters for Windows NT Version 1.0 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 for Windows NT is extensible, flexible, and hardware independent. The first release of 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 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

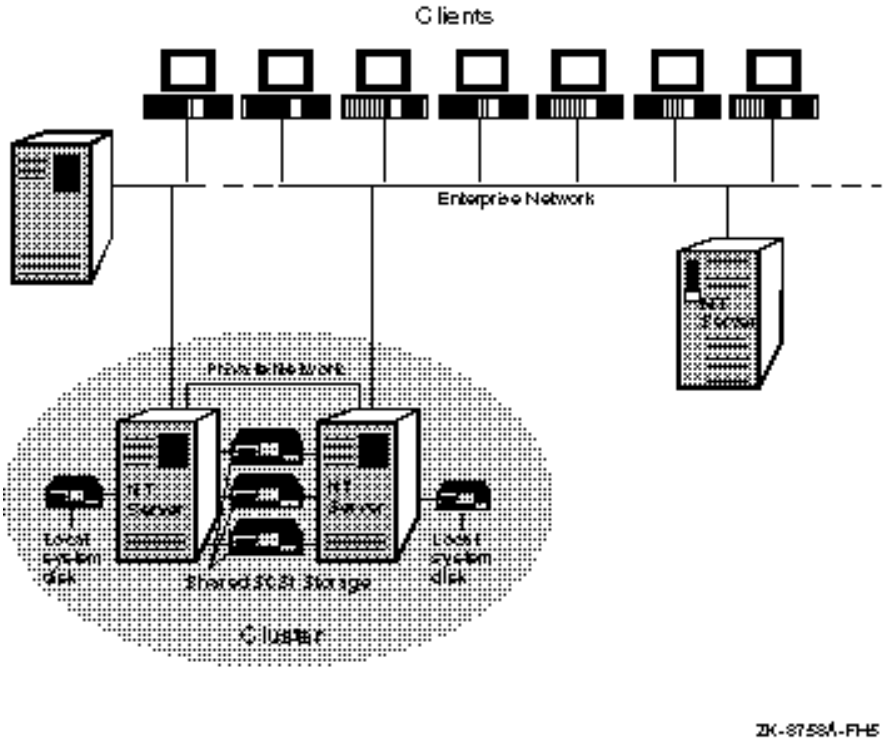

Digital strongly recommends a second, dedicated ethernet connection between the paired clustered servers for the following reasons:

- 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. Digital Clusters for Windows NT uses a common cluster name, or *cluster alias*, which makes the functioning of the cluster transparent to the end user. 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. Existing client applications can make use of the failover function without modification. Depending on the application, the user sees either 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.

The level of availability provided by Digital Clusters for Windows NT compares well with fault-tolerant systems. Digital Clusters for Windows NT 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 for Windows NT, 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. It reduces your investment in new resources because it does not depend on custom hardware or proprietary interconnects.

Scalability

Digital Clusters for Windows NT is a highly scalable 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 scalability of Digital Clusters for Windows NT rests with the partitioned data model of its software architecture. This model delivers numerous benefits, including:

- Improved system performance. Because the work load is divided between the servers in the cluster and processed independently, better *workload balancing* is possible than with other types of solutions. Better workload balancing results in improved system performance.
- Greater flexibility. Dividing the work load into smaller components provides greater flexibility. Digital Clusters for Windows NT 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 for Windows NT enables scalability. In contrast, both the mirrored backup and the high degree of synchronization of fault-tolerant systems undermine scalability.

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 for Windows NT 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.

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

Digital Clusters for Windows NT is fully compatible with Windows NT Server Version 3.51. Unlike other solutions, it 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 for Windows NT 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.

Supported Server Architectures

Digital Clusters for Windows NT 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 and can use the cluster alias to access the cluster. 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.

End-user clients access and manage Digital Clusters for Windows NT 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. The cluster configuration can be changed from any one of the servers in the cluster.

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 for Windows NT allows workload partitioning to the disk level by assigning disks on the shared 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 to 7 for details on using Cluster Administrator.) There is no need to turn off the power or recable hardware components.

Digital Clusters for Windows NT supports the use of standard Windows NT management tools, such as File Manager 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.

Chapter 2
Cluster Administrator Overview

This chapter introduces Cluster Administrator, discusses the concepts of *failover* and *failback*, and defines the terms *failover objects*, *failover groups*, and *failover policy*. It also presents a database failover example.

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.

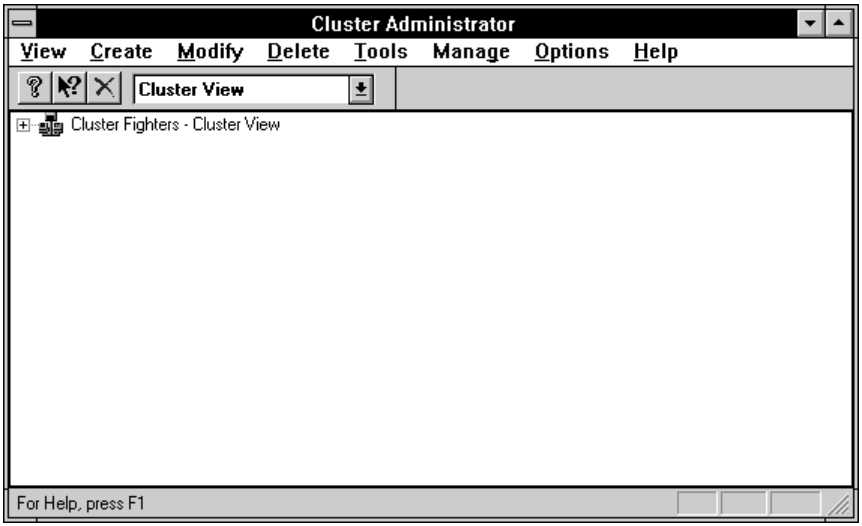

The main window contains a menu bar, toolbar, status bar, and shows the Cluster View by default.

The toolbar provides quick access to commands through shortcut buttons and allows you to change Views by selecting from a list, as described in Chapter 4.

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.

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. This transparent relocation of cluster services and resources is referred to as *failover*.

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, 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*, 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 easily defined using Cluster Administrator. (See Chapter 7 for details.) They often consist of a collection of applications and storage devices that are used together. For example, a database server application such as Microsoft SQL Server and the shared cluster disk 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. 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 describes what happens when the cluster server causing the failover returns to an operational status. If failback is enabled, the failover groups automatically migrate back to the primary server in their original configuration. If failback is not enabled, the failover groups remain on the alternate cluster server. 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.

Database Failover Example—Normal

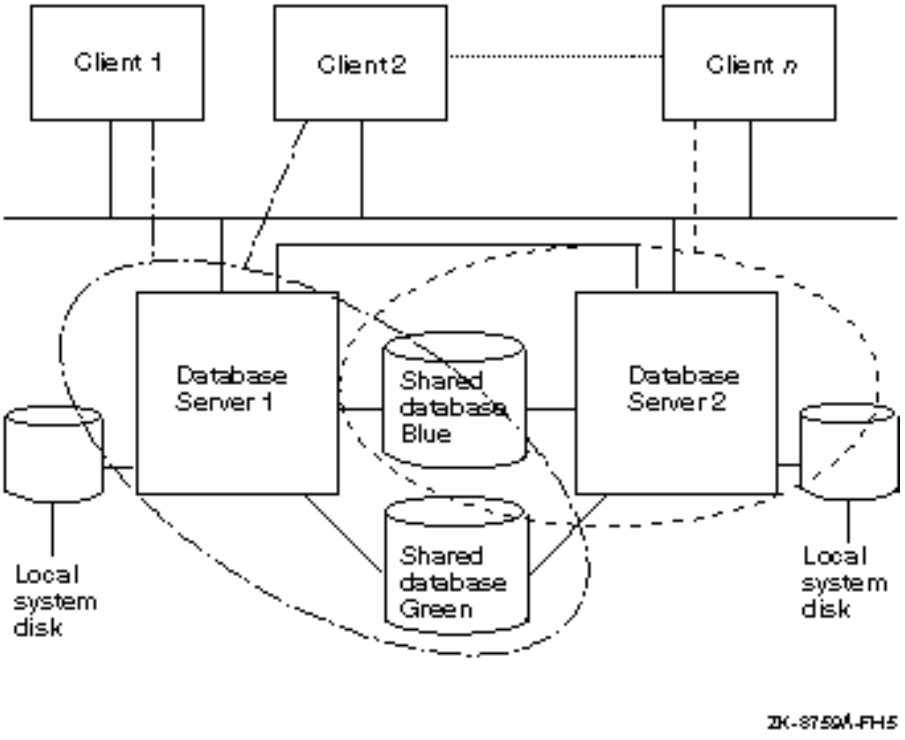

Several end-user clients are accessing the Blue database via Server 1, while other end-user clients are accessing the Green database via Server 2. In this example, we will assume that all the clients have connected to the cluster using the cluster alias. The clients do not need to know which server a 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

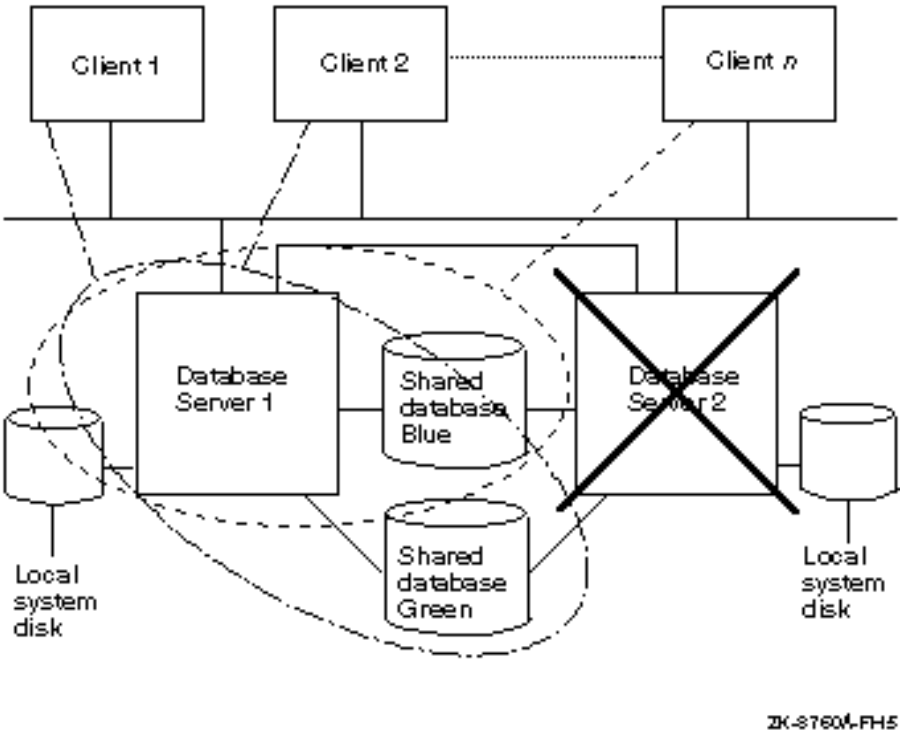

Note that it is necessary for the database application to be installed on both cluster servers for failover to occur. For detailed instructions on database software installation and configuration requirements for failover, see Chapter 3. For a discussion of application considerations, see Chapter 8.

Chapter 3
Configuring Database Software for Failover

Before using Cluster Administrator to configure for Microsoft SQL Server or Oracle7 Workgroup Server failover, you must complete the database software installation and configuration requirements outlined in this chapter.

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 databases in the Digital Clusters for Windows NT environment.

Software Requirements

Software requirements for *each* cluster server are as follows:

- Acquire one or more Microsoft SQL Server licenses in accordance with the Microsoft SQL Server licensing requirements.
- Install Microsoft SQL Server, Version 6.5.
- Configure the SQL master database on a local disk.

Access to Shared Disks

Access to shared disks is not required when installing SQL Server. However, it is required when adding new SQL databases that you would like to designate as highly available. The new databases can be created either before or after cluster software installation.

Designating Primary and Failover Servers

In SQL 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 becomes available again. 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.

To add shareable databases and configure the server systems, you must first designate a primary server and a failover server. The shareable databases must be created from the primary server.

Creating SQL Databases Before Installing the Cluster Software

If you choose to create shared SQL databases prior to installing the cluster software, only the primary server must be allowed access to the shared disks. Digital recommends you do this by physically disconnecting the failover server from the shared bus. Then create the shared databases.

Creating SQL Databases After Installing the Cluster Software

If you choose to create shared SQL databases after installing the cluster software, you must access the shared disks from the primary server. Digital recommends you do this by creating a failover group that includes the shared disks where the databases will reside. At this point in the SQL Server configuration, it is optional that the failover group includes an SQL failover object. For information on creating a failover group, see the section Creating a Failover Group in Chapter 7. For information on creating an SQL failover object, see the section Creating an SQL Failover Object in Chapter 7.

Once you have created the failover group, verify that the primary server has control of the shared disks.

▶ **To verify that the primary server has control of the shared disks:**

1. From the Tools menu, choose Disk Administrator. Disks not controlled by the primary server will be designated as OFF-LINE.
2. If the shared disks are not controlled by the primary server, initiate a manual failover. See the section Managing Manual Failover in Chapter 6 for instructions.

Then create the SQL databases on the shared disks.

SQL Server Database Failover

SQL Server database failover refers to a database failing over, not the entire application. The SQL Server software can be running on both cluster servers and simultaneously accessing locally installed databases. However, only one server system at a time can access databases on a given shared disk. Upon primary server failure, the failover server will service the shared databases.

To ensure failover, the servers and databases must be properly configured by invoking stored procedures supplied by SQL Server. Invoke these procedures using either the ISQL/w application or the Enterprise Manager Query Tool, which are part of the SQL Server product. For further information, see the documentation and online help packaged with the Microsoft SQL Server Client-Server Database Management System for Windows NT product.

Configuring for High Availability of SQL Databases

Once you have defined the shared databases, you need to configure the server systems for high availability. SQL Server must be running on both cluster servers at this time. Digital recommends that the cluster software also be installed.

▶ **To configure the primary server:**

1. Record the failover server:

   ```
   sp_addserver FailoverServerName, fallback
   ```

2. Give the system administrator (sa) account on the failover server permission to log in as the sa on the primary server:

   ```
   sp_addremotelogin FailoverServerName, sa, sa
   ```

3. Enroll all shared databases owned by the primary server for potential failback support:

   ```
   sp_fallback_enroll_svr_db PrimaryServerName, DatabaseName
   ```

To configure the failover server:

1. Record the primary server for which failover support may be provided:

   ```
   sp_addserver PrimaryServerName
   ```

2. Give the sa account on the primary server permission to log in as the sa on the failover server:

   ```
   sp_addremotelogin PrimaryServerName, sa, sa
   ```

3. Verify that the RPC connections are functioning:

   ```
   PrimaryServerName...sp_helplogins
   ```

4. Enroll all shared databases owned by the primary server for potential failover support:

   ```
   sp_fallback_enroll_svr_db PrimaryServerName, DatabaseName
   ```

Oracle7 Workgroup 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. 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 Workgroup Server in the Digital Clusters for Windows NT environment.

Software Requirements

Software requirements for *each* cluster server are as follows:

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

Oracle Instance Failover Prerequisites

For an Oracle instance to failover, the following prerequisites must be met:

- 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 file, control file, log file, and parameter file, must reside on one or more shared disks.
- The shared disks and Oracle failover object must be in the same failover group.

Creating an Oracle Instance for Failover

This section gives instructions on how to create an Oracle instance for failover.

▶ **Before creating an Oracle instance for failover:**

1. If you have not already done so, install the cluster software.
2. Using Cluster Administrator, create a failover group containing all shared disks to be used for Oracle7 Workgroup Server. See the section Creating a Failover Group in Chapter 7 for instructions.

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

Note

Assign identical fixed drive letters on each cluster server for the shared disks used for Oracle instance failover. For instructions on how to assign a fixed drive letter, see the last procedure in this section.

Repeat the steps in the following procedure for each instance that you configure for failover.

▶ **To create an Oracle instance for failover:**

1. On the primary server, create an Oracle instance for failover using an Oracle database administrator tool.
2. Use a shared disk to store all files associated with the instance, including the data file, control file, log file, and parameter file.
3. Create the associated database on one or more shared disks. If you are working with an existing database, move the database data files to the shared disks.

4. *For Version 7.2 only:* If you are running Oracle7 Workgroup Server Version 7.2, perform the following steps for *each additional instance after the first instance* that you configure for failover:

    a. Create a separate network listener for the instance using Oracle Network Manager for Windows. See the Oracle Network Manager for Windows product documentation for instructions.

    b. Specify Named Pipe as the network protocol for each listener.

    c. Specify a unique name for the named pipe associated with each network listener.

    d. Start the network listener using the Services icon on the Control Panel.

5. Verify that you can start the server instance.

6. Verify that you can access the associated database using an Oracle database administrator tool. If you cannot access the associated 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.

7. Shut down the instance using an Oracle database administrator tool. Be sure to leave services for the instance running.

8. Using the Services icon on the Control Panel, verify that services for the instance are set to restart automatically.

9. Manually fail over the disk group to the failover server. See the section Managing Manual Failover in Chapter 6 for instructions.

10. Once the failover server can access the shared disks, create an *identical* Oracle instance on the failover server.

11. *For Version 7.2 only:* Repeat step 4.

12. Repeat steps 5 to 7. When repeating step 7, shut down the instance on the failover server.

13. Manually fail back the disk group to the primary server. See the section Managing Manual Failover in Chapter 6 for instructions.

14. Create an Oracle failover object, supplying information about the instance that you created in step 1. See the section Creating an Oracle Failover Object in Chapter 7 for instructions.

15. Modify the Oracle failover group you created (in the previous procedure) to include the Oracle failover object that you created in step 14. See the section Modifying a Failover Group in Chapter 7 for instructions.

16. Verify that the Oracle Server instance is now running on the primary server. If the instance is not running on the primary server, verify that the information you supplied when creating the Oracle failover object in step 14 is correct.

▶ **To assign a fixed drive letter to a shared disk:**

1. On the primary server for the shared disks used for Oracle instance failover, choose the Administrative Tools program group in Program Manager.

2. Choose the Disk Administrator icon.

3. Select a shared disk for which you want to modify the drive-letter assignment.

4. From the Tools menu, choose Drive Letter.

   The Assign Drive Letter dialog box is displayed.

5. 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.

6. Repeat steps 3 to 5 for each shared disk used for Oracle instance failover.

7. Fail over the disk group to the failover server. See the section Managing Manual Failover in Chapter 6 for instructions.

8. Once failover has completed, repeat steps 1 to 6 on the failover server.

9. On the failover server, fail over the disk group to the primary server. See the section Managing Manual Failover in Chapter 6 for instructions.

10. Using Disk Administrator, verify that the shared disks have failed over.

Oracle Server Database Failover

Throughout the Digital Clusters for Windows NT documentation, *Oracle Server database failover* refers to an Oracle7 Workgroup Server instance failing over. Oracle7 Workgroup 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.

If an Oracle instance fails while the server system on which it is running continues to function properly, database administrator intervention may be required. Once the problem has been remedied, the database administrator can restart the instance.

Initiating Manual Failover of an Oracle Instance

If you need to manually fail over an instance, Digital strongly recommends that you first shut down the instance using the database administrator tools provided with the Oracle7 Workgroup Server software. This guarantees orderly shutdown of the instance. Then use Cluster Administrator to initiate a manual failover.

▶ **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 6 for instructions.

Chapter 4
Getting Started with Cluster Administrator

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

Starting Cluster Administrator

▶ **To start Cluster Administrator:**

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:**

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

Any of these views can be the starting point for a cluster administration operation.

Displaying the System View

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

The System View can be 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 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.

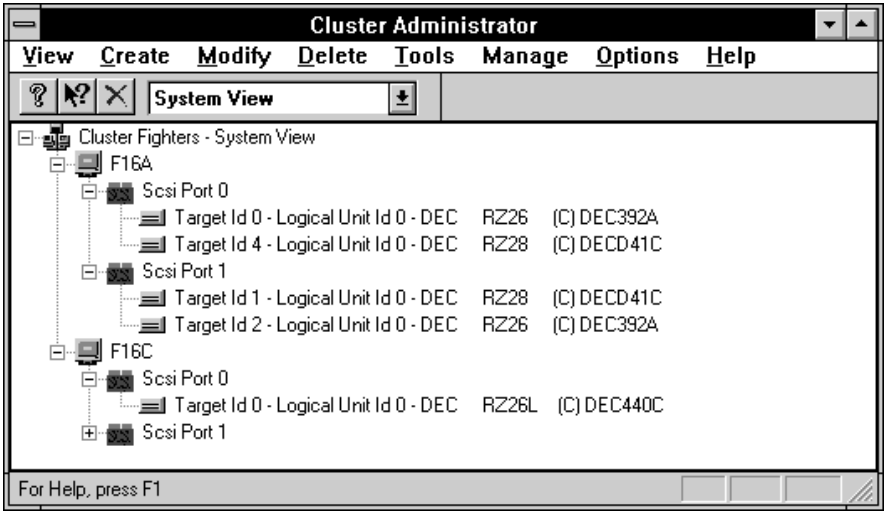

Displaying the Cluster View

The Cluster View allows you to examine the cluster from a failover group perspective. 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 a recognized cluster applications such as Oracle7 Workgroup Server and SQL Server, and scripts.

The Cluster View can be 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 Cluster View:

Object Type	Default Action
Group	Modify Failover Group
Disk	Manage Disk Alias
Oracle	Modify Oracle Failover Object
SQL	Modify SQL Failover Object
Script	Modify Script Failover Object

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.

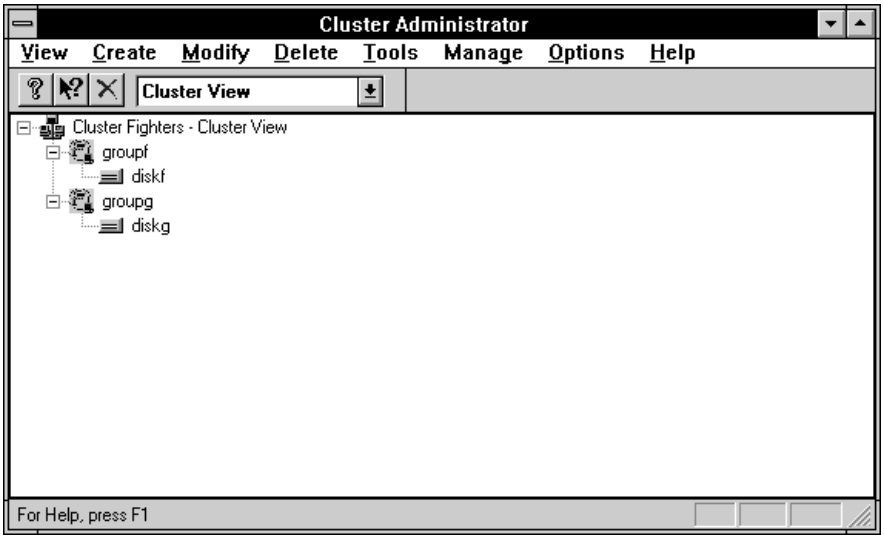

Displaying the Class View

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 groups, disks, shares, servers, Oracle objects, SQL objects, and script objects. Oracle objects, SQL objects, and script objects are classified as Applications in the Class View.

The Class View can be 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
Group	Modify Failover Group
Disk	Manage Disk Alias
Server	None
Oracle	Modify Oracle Failover Object
SQL	Modify SQL Failover Object
Script	Modify Script Failover Object

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.

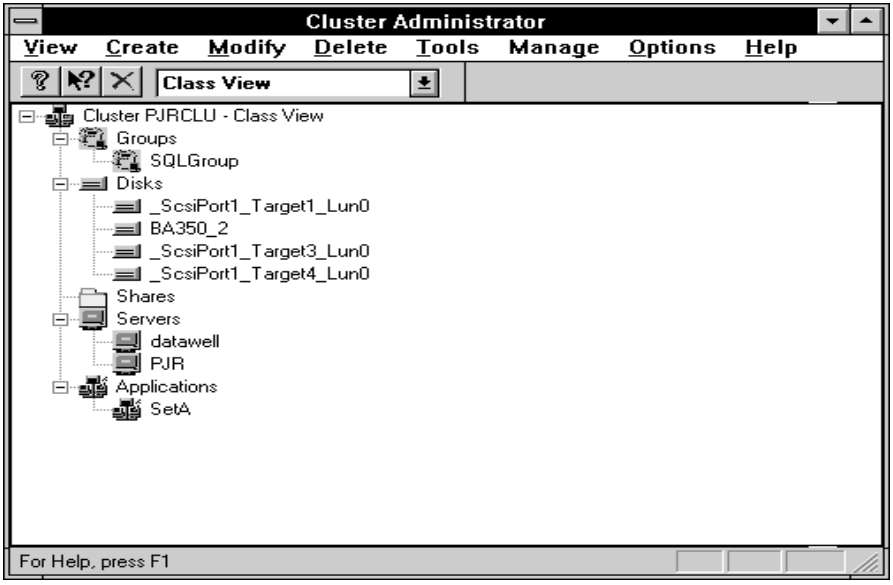

Chapter 5
Configuring Your Cluster

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 *Digital Clusters for Windows NT 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. See the section Starting Cluster Administrator in Chapter 4 for instructions.

   The Cluster View is shown by default.

2. Display the Class View. See the section Displaying the Class View in Chapter 4 for instructions.

   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. The information displayed is useful to know before creating a failover group.

3. If you plan to take advantage of the Microsoft SQL Server or Oracle7 Workgroup Server database failover features offered by the Digital Clusters for Windows NT software, you must first:

   a. Complete the appropriate database software installation requirements outlined in Chapter 3.

   b. Create an SQL failover object or an Oracle failover object as instructed in the sections Creating an SQL Failover Object and Creating an Oracle Failover Object in Chapter 7.

4. 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 7 for instructions.

5. Display the Cluster View. See the section Displaying the Cluster View in Chapter 4 for instructions.

   The Cluster View allows you to examine the cluster from a failover group perspective. It includes all defined failover groups and members of each group.

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 groups you created in step 4. Disks not controlled by the server on which you run Disk Administrator will be designated as OFF-LINE.

   - Determine which drive letters are associated with cluster disks.

8. Optionally, you may want to create one or more network file shares on a cluster disk:

Note

You can create network file shares only from the server system that has the disk on line.

a. From the Tools menu, choose File Manager.

b. Select a cluster drive.

c. From the Disk menu, choose Share As.

d. Enter a share name in the Share Name box.

e. Enter a path specification in the Path box.

Chapter 6
Managing a Cluster

This chapter presents step-by-step procedures on how to manage a cluster using the Cluster Administrator. The topics covered in this chapter include:

- Managing an Adapter Configuration
- Managing Disk Aliases
- Managing an Event Log
- Managing the Log Disk
- Managing Manual Failover

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 7 for information about creating, modifying, or deleting failover groups.

Managing an 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 the 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, the 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), and then reboot the system. Or, if the wrong adapter is selected, select the correct adapter and then reboot the system.

If you remove an 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.

To manage an adapter configuration:

1. From the Manage menu, choose Adapter Configuration.

   The Manage Adapter Configuration dialog box is displayed.

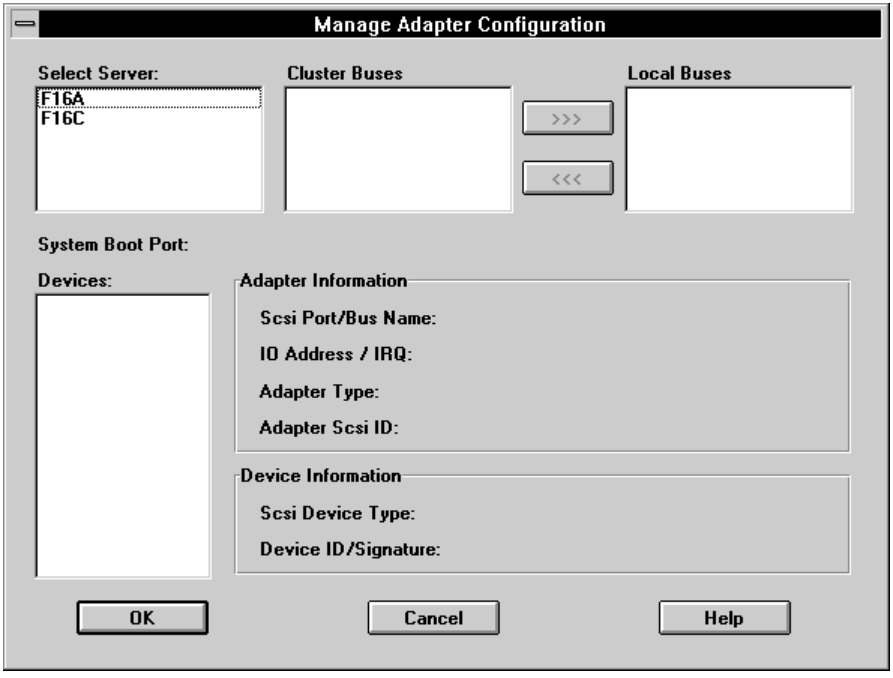

   In the dialog box, the Select Server list box displays the computer names of the server systems in the cluster.

2. In the Select Server list box, select a server for which you want to change the SCSI bus configuration.

The 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, the Cluster Administrator displays 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. The 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 type string, the target ID number, 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 does not automatically update the 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.

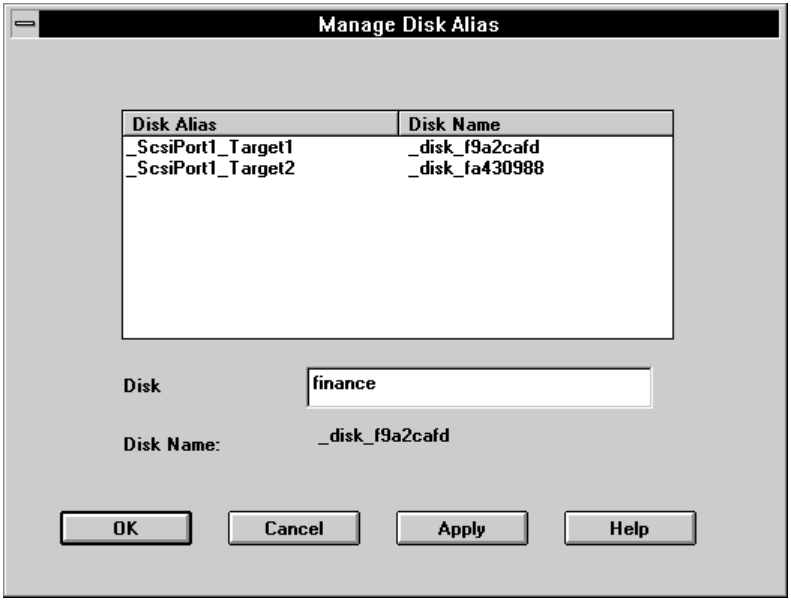

The dialog box displays all disk names that are cluster-available (that is, are configured as part of a cluster instead of being configured as local). For each disk, the Disk Alias list box first displays the disk name automatically, and then as the disk comes on line, changes to the disk volume name.

2. Select one of the displayed disk names.
3. In the Disk 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 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.

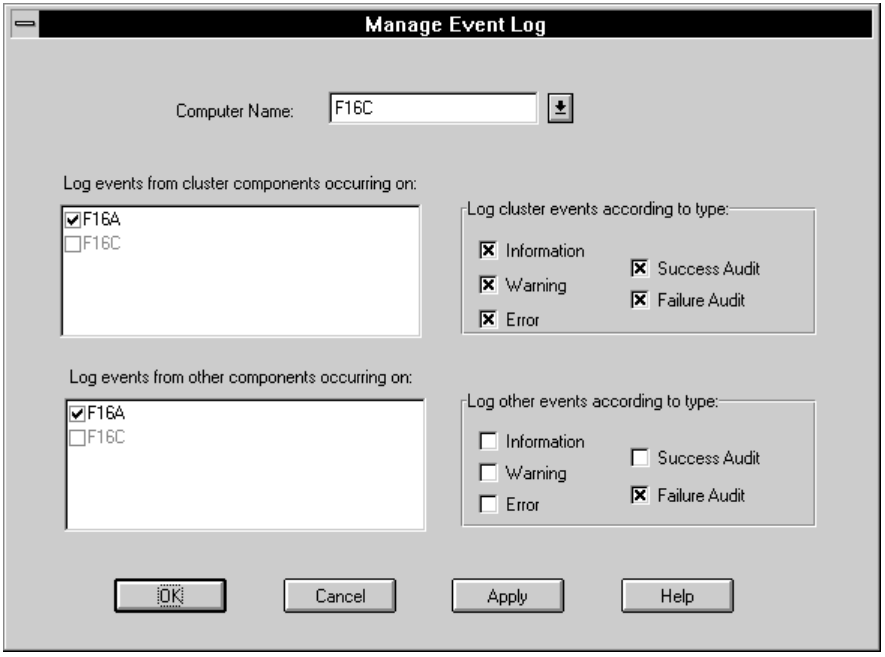

The selected computer name of the current server is displayed automatically in the Computer Name list box. You can manage event logging for that server, or select the other server to manage the event logging on the failover server.

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 checked (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 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.
4. Choose either:
    - 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

The following cluster components of Digital Clusters for Windows NT log events:

- Device drivers (including cluster disk driver, cluster file driver, and cluster port driver)
- Cluster Configuration Database server
- Cluster Failover Manager
- Cluster Name Server
- Cluster Event Log Service

Managing the Log Disk

The Cluster Configuration Database 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 again becomes available, 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 the 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.

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 all shared disks.

   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.

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 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.

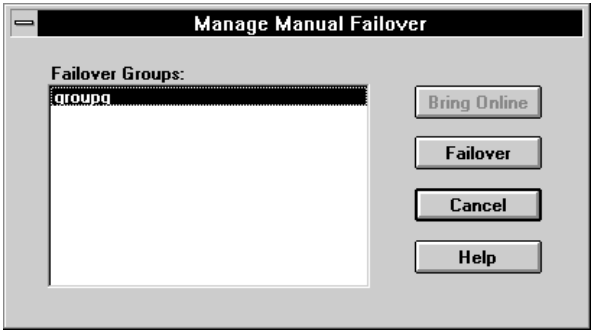

2. In the Failover Groups list box, select a group and click Failover.

   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.

Chapter 7
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 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 the 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 the 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 failover, refer to Chapter 3, Configuring Database Software for Failover.

Creating an Oracle Failover Object

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

Note

Once 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.

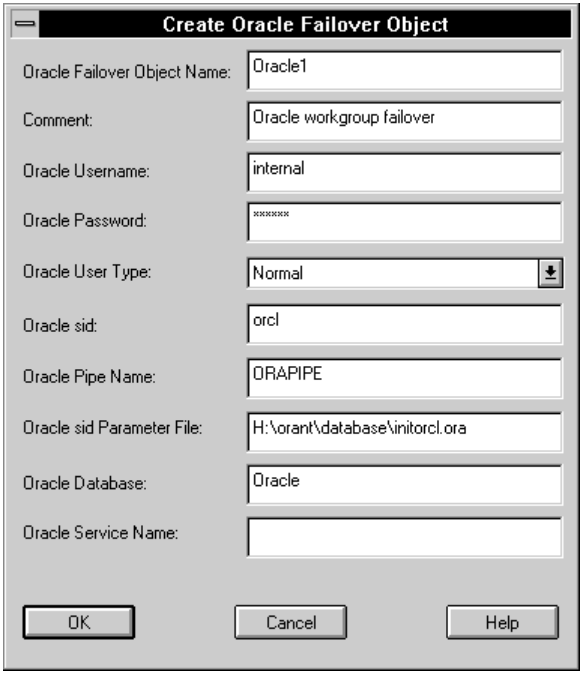

2. Supply the Cluster Administrator with information about your Oracle database in the following text boxes:

   - 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 Workgroup Server Version 7.1, you need to use the default name "internal" because this is the only name that can start the database. For Oracle7 Workgroup Server Version 7.2, 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 the 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 — Specifies:
     - For Oracle7 Workgroup Server Version 7.1, the Oracle system identifier as the pipe name.
     - For Oracle7 Workgroup Server Version 7.2, for each instance, the pipe name of the network listener associated with that instance. For example, use ORAPIPE for the default network listener.

     Note that you can check your Oracle7 Workgroup Server Version 7.2 pipe name by using the SQL*NET™ configuration tool. Refer to 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 Service Name — Specifies a service name alias. This text box is not used for Oracle7 Workgroup Server Version 7.1 or 7.2 but might be required by Version 7.3.

  This text box is reserved for future use. You do not need to supply an alias name for this release of the Cluster Administrator.

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 the 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.

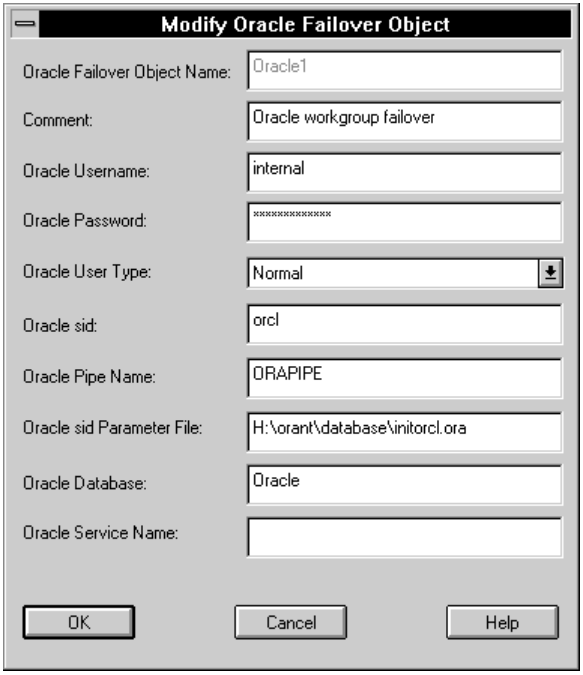

The selected Oracle failover object is displayed in the Oracle Failover Object Name text box.

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

3. Update the 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.

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.

   When you refresh the view, the Class view shows the new configuration without the Oracle failover object.

Working with a Script Failover Object

The Cluster Administrator can use a script failover object to execute a command 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 might 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 use it to notify a system administrator of a failover event. You can set up such notification with the `net send` command.

To use a script failover object, create a script failover object with the Cluster Administrator, as described in the next section. Then add the object to a failover group by using 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

Creating a script failover object to a group allows you to use a generic failover object to apply an NT script startup command when a group is coming on line or going off line.

Note

Once you create a failover object, as described in this section, it becomes an available resource that you can add to a failover group. Remember to add the resource to a group by selecting Failover Group from the Create menu 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.

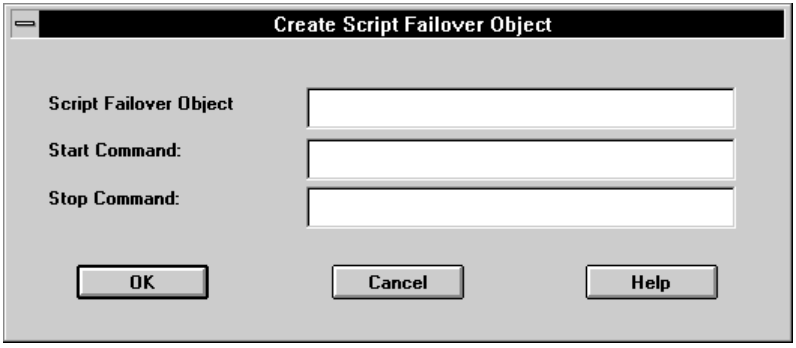

2. In the Script Failover Object text box, enter the object name associated with the specific script. The script object name can be any name that is different from other script objects.

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

   The start command is executed when the script failover object is put 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`). It can be 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 be run when the Failover Manager takes a group off line.

   The stop command is executed when the script failover object is put off line as part of a failover group. It follows the same rules as the start command.

   Both the start and stop commands are optional, in which case no command will be executed during the corresponding transition.

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:

- Access rights for the commands and the account must match — The start or stop command must be executed with the same access rights as the Cluster Failover Manager account. The commands also require the same path that is associated with the Cluster Failover Manager account. If you provide no path for the commands, the Cluster Administrator uses its own default path.

  If you want to use the same script file on both servers, make sure to locate the script on shared storage, preferably on a disk that is a member of the same group as the script.

  Also, the script must be on storage that is accessible to the failover manager. Therefore, if the script is on local storage, manually place the script on both servers of the cluster. In general, do not locate the script on a network share, as network shares are unavailable to the failover manager. If the script is on a cluster disk, place the disk in the same failover group as the script.

- 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 Cluster Administrator information about 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 the Modify menu, choose Script Failover Object, 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.

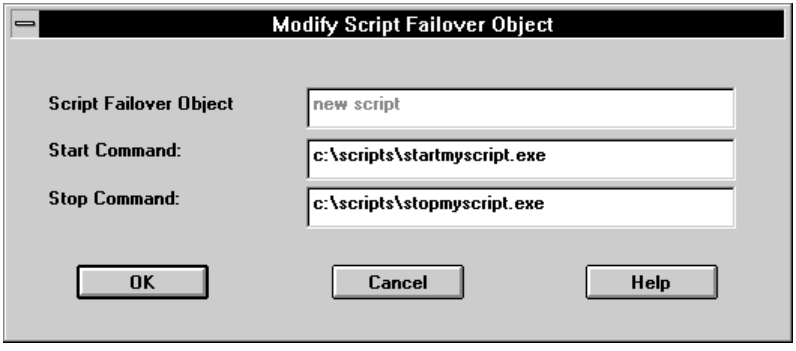

2. Edit any of the information in the text boxes to reflect the changes made to the script failover object.
3. Click OK to accept the changes, or Cancel to cancel the changes.

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.

   The Cluster Administrator deletes the selected script from the cluster database.

Working with an SQL Failover Object

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

To ensure that the Cluster Administrator properly fails over a disk containing a Microsoft SQL database server that resides on a shared disk:

1. Create a failover object for the SQL database server.

   For each distinct set of disks containing SQL databases that you want the Cluster Administrator to fail over, create a failover object. Doing this defines an instance of a failover object.

2. Create a failover group containing that SQL object.

   When you create the failover group and associate the SQL failover object with it, a binding is defined between the disks in the group and the SQL databases that are both located on that disk and enrolled as failover databases. (For information about enrolling databases, refer to Chapter 3.)

   In general, for each failover group, you need to create only one SQL object.

For example, you might want to create a cluster structure similar to the one shown in this Cluster View:

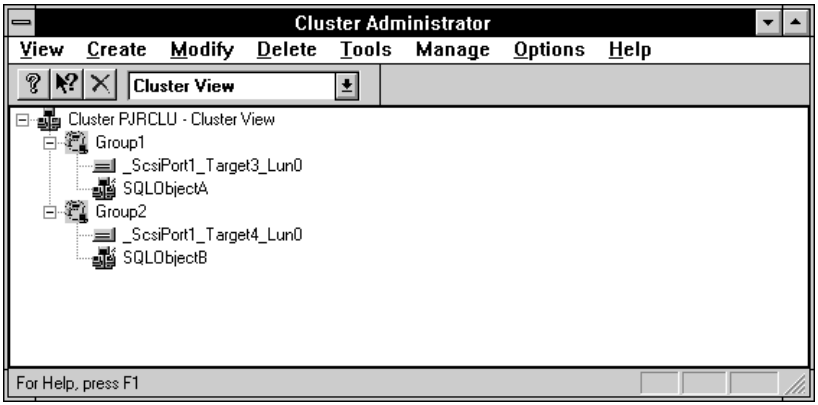

In this example, a system administrator created two SQL failover objects, SQLObjectA and SQLObjectB. For each of these objects, the system administrator then created a failover group and associated a specific cluster disk with that group. (Group1 contains one disk and the failover object associated with that disk. Group2 contains the other disk and the other failover object.)

Creating an SQL Failover Object

Use the Create SQL Failover Object dialog box to create an SQL failover object. You can then create a failover group and add the SQL failover object to the group. For information about creating a failover group, see the section Working with a Failover Group in this chapter.

Once you add an object to a cluster, as described in this section, it is an available resource that you can add to a failover group. Remember to add the resource to a group by selecting Failover Group from the Create menu and moving the object from Available Failover Objects into the Group Contents. An object also can be added to an existing group from the Modify menu.

To create an SQL object:

1. From the toolbar, select Create, SQL Failover Object.

The Create SQL Failover Object dialog box is displayed.

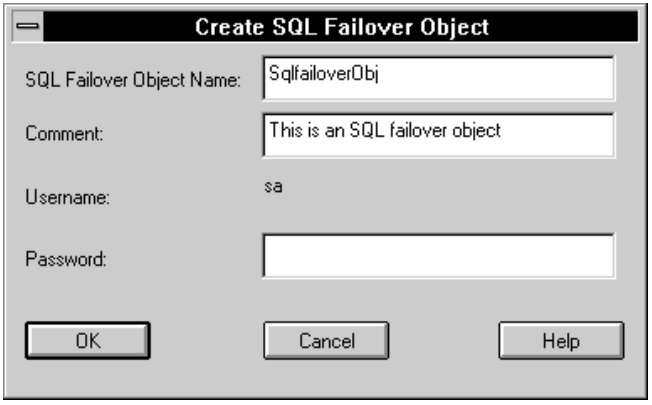

2. In the SQL Failover Object Name text box, supply the name of the failover object. Use this name when adding this SQL object to a failover group.
3. In the Comment text box, optionally supply a comment that is displayed when this SQL failover object is advertised via the cluster name service software.
4. In the Username text box (read only), the user name is specified by default for the SQL account associated with this instance of the database. The user name always defaults to sa.
5. In the Password text box, supply the password for the SQL sa account, which is used to protect access to an instance of the database. Note that the SQL password has a length limit of 30 characters.
6. Click OK to save the modifications to the settings, or click Cancel to cancel the modifications.

Modifying an SQL Failover Object

Modify the characteristics of an SQL failover object whenever these characteristics change. For example, you can update the information if the system administrator password is changed.

To modify an SQL object:

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

The Modify SQL Failover Object dialog box is displayed.

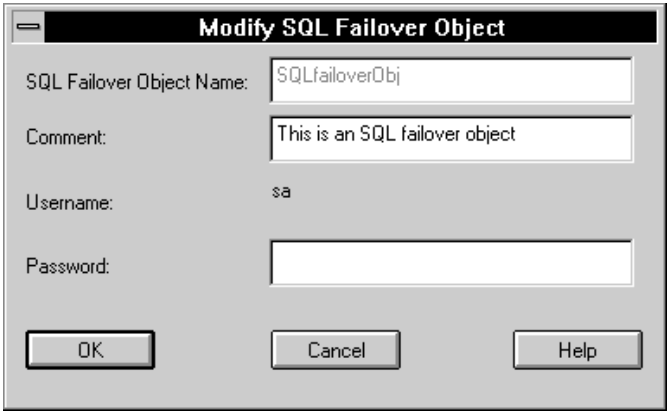

In the SQL Failover Object Name text box**,** the name of the selected SQL failover object is displayed. If you supplied a comment when first creating the SQL failover object, the comment is displayed also.

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 failover object is advertised via the cluster name service software.
4. In the Password text box, you can change the password for the SQL 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 Failover Object

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

To delete an SQL object:

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

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

3. From the View dialog box, click Refresh to see the new configuration (without the SQL object).

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 a Failover Group

Using a failover group ensures that the 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 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.

Creating a Failover Group

Create a generic failover group to ensure that the 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.

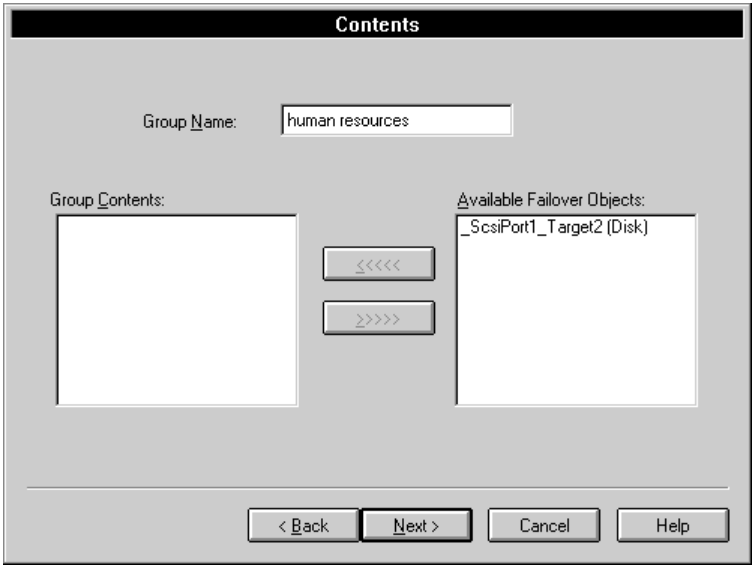

The Contents page of the Create Failover Group contains:

- 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, the 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. Refer to the section Deleting a Failover Group for information about deleting a group.

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

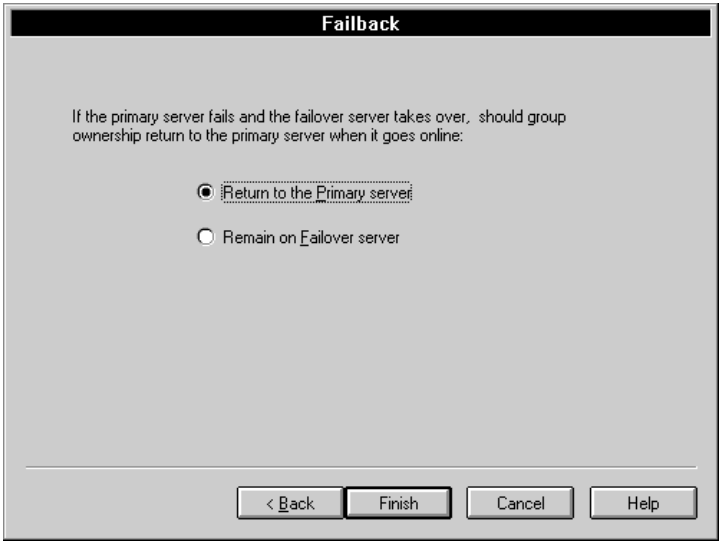

6. On the Failback dialog box, specify which server should regain control of the group after a failure. Specify that the group should either:

   - 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.

- 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.

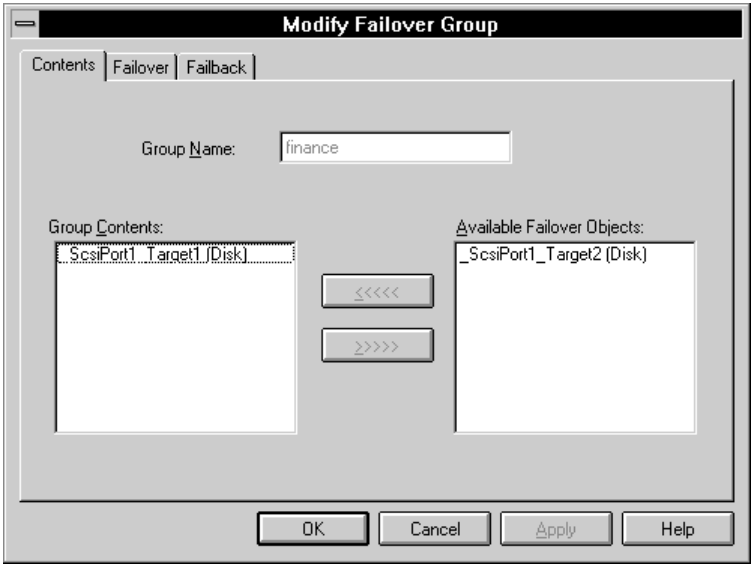

The Contents page of the Modify Failover Group dialog box contains:

- 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.

Use this page to specify which system should be the primary server, that is, where the group will be available normally:

- Primary Server field — Specify the computer name of the system that should be the server under normal circumstances.
- Failover Server field (read only) — Displays the computer name of the system that you want to serve the selected group if the primary server computer fails. Note that, because there are only two servers for any failover group, the failover server defaults to the computer not specified as the primary server.

- Disable Failover for this group checkbox — Disables failover for a group until you enable it again (by removing the check from this checkbox). Disabling failover allows the group to remain on line but not move to the other server upon a cluster failure. 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.

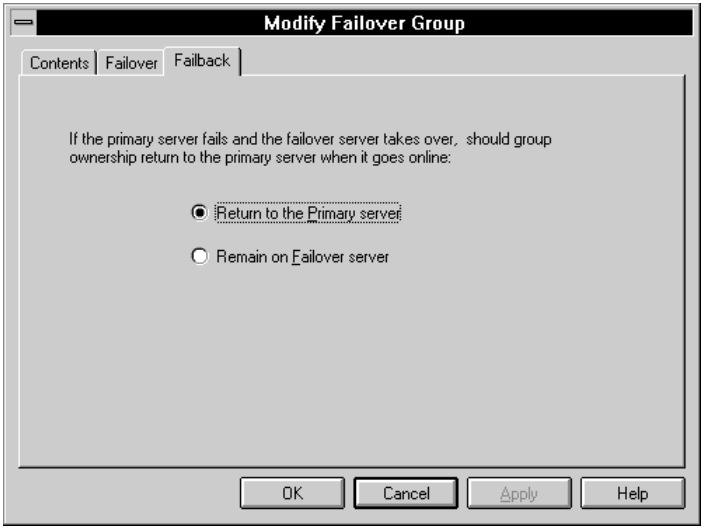

Use this page to modify the failback behavior of a group. The page includes two radio buttons:

- 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 event log 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

Should you decide to reconfigure your failover group configuration and need to delete one or more failover groups, you can do so easily**.**

▶ **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.

   The 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.

Chapter 8
Application Considerations

The Digital Clusters for Windows NT software furnishes high server availability to cluster resources. Version 1.0 provides the following failover capabilities:

- NTFS file services and network shares (NetBEUI, TCP/IP, and IPX/SPX)
- Microsoft SQL Server, Version 6.5
- Oracle7 Workgroup Server, Versions 7.1 and 7.2
- 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 for Windows NT 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 used to send notification messages to the system administrator in the case of a failover event.

Application Handling During a Failover

Like any high-availability product, Digital Clusters for Windows NT 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 client applications, Oracle 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 client cluster 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

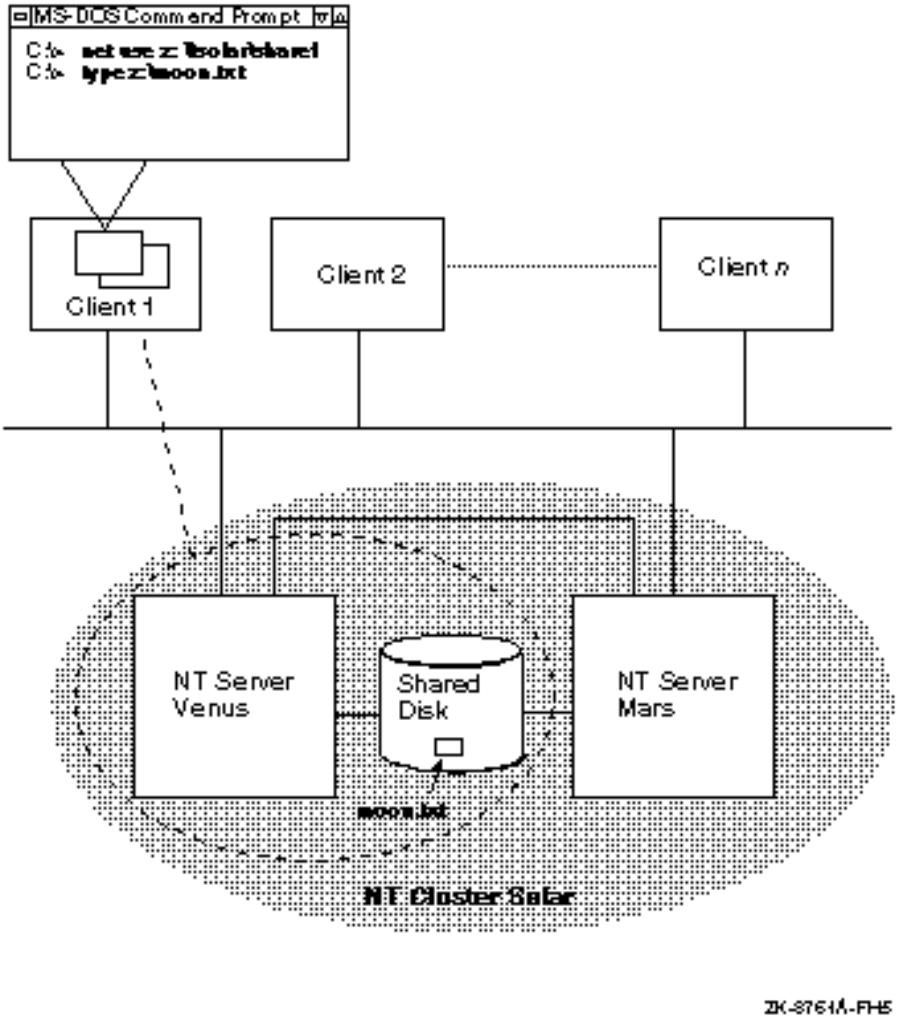

The following command establishes a connection to `\\venus\share1`:

```
net use z: \\solar\share1
```

The next command opens the file `star.txt` on `\\venus\share1`:

```
type z:\star.txt
```

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

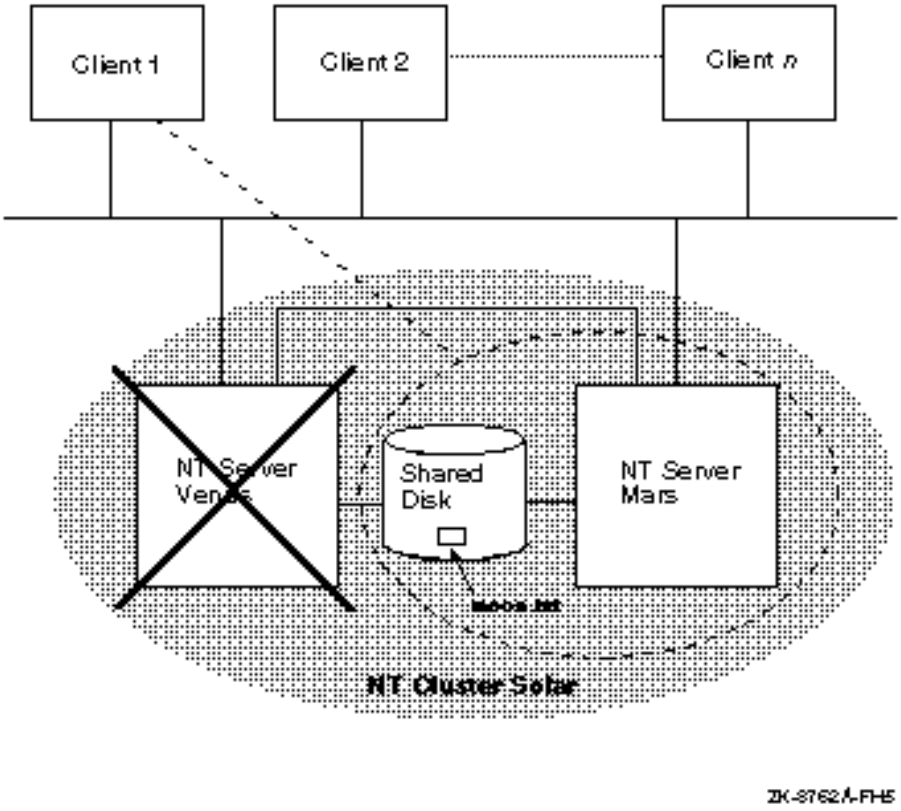

Database Application Failover

Digital Clusters for Windows NT 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 Workgroup 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 Version 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 and Named Pipes
- Failover Times

Open Files and Named Pipes

If a cluster resource that contains an open file or named pipe 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 or named pipe, establishing 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.

When a file or named pipe 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 for Windows NT. 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 cluster server to continue access to the cluster service.

Chapter 9
Troubleshooting

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 `\winnt35\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 Software Utilities in Appendix A for more information about the CLUIVP utility.)

Where are my disks? I can't create a group.

If you run Class View from the Cluster Administrator and there appear to be no shared disks with which to create your cluster groups, check the following:

1. Did you reboot both servers after installing the 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 software will result in no signatures being read.

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 on 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 from Control Panel, verify that the Cfmd Server, the Cluster Failover Manager, and the Cluster 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 *Digital Clusters for Windows NT 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:

   ```
   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 the 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 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 *second* server, examine the 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 second server has control of the group.

   (Note that you also can use FMSTAT and CLUSTAT to view graphically how the cluster resources are allocated. These unsupported cluster utilities are supplied on the Digital Clusters for Windows NT distribution CD–ROM.)

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 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 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 group.

3. Start the Cluster Administrator on the server controlling the group. Select the group and then use the Manual Failover option to move control to the other server.
4. Within a minute or less, the `ConnectionPoint` parameter should move to the second server, demonstrating that the cluster is failing over properly.

Note

Instead of opening the Registry and looking for the `ConnectionPoint` parameter, you also can use FMSTAT and CLUSTAT to view how the cluster resources are allocated and observe dynamically the group failover. These unsupported cluster utilities are supplied on the Digital Clusters for Windows NT distribution CD–ROM.

If this procedure demonstrates that the cluster is operating normally, the failover problem you are seeing is at the client system. See the section My client hangs after a failover, later in this chapter. If you are running a database server such as Microsoft SQL Server or Oracle7 Workgroup Server, you also should refer to the section Database Problems, later in this chapter.

If you did not see the `ConnectionPoint` parameter move to the second server, your cluster is not operating properly. Look in the Failover Manager trace log for error messages that might indicate what the problem is. The trace log file is found in the `temp` subdirectory of the cluster destination directory, as specified when the server software was installed. If you accepted the default directory during installation, the trace log file path is as follows on your system disk:

```
\Program Files\Digital\Clusters\temp\fm#.log
```

where # is the sequential number of the trace log. (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 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 NetBEUI, IPX/SPX, and TCP/IP protocols.) If the client is using only the TCP/IP transport protocol, the client and server systems also must be on the same subnet.

My client doesn't see the cluster it needs.

The Cluster 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 group?

   Again, the cluster software follows the same rules and restrictions as Windows NT. You must create a network share from your cluster 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 Network Share 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 a defined `ConnectionPoint` parameter for the resource, then the cluster Failover Manager has not put the resource on line.

My client hangs after failover.

There are several reasons that a client application might hang or appear to hang after failover, including the following:

- The resource may be taking a long time to failover. 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 name), 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 Network Share key on both servers:

```
HKEY_LOCAL_MACHINE\System\CurrentControlSet\
   Services\Cfmd\Database\NetworkShare
```

2. In the list of resources, double-click on the resource in question and look for the `ConnectionPoint` parameter.
3. If neither server has a defined `ConnectionPoint` parameter for the resource, then the cluster Failover Manager has not put the resource on line. See the section My group won't come on line, earlier in this chapter, for more information.

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 chose 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 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 Workgroup Server database software.

My database isn't available.

If you get an error message telling you that your database is not available, you first should determine whether all the necessary services for your database server are running. The

database administrators for both the Microsoft SQL Server and Oracle7 Workgroup Server products provide methods for determining 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.

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. You must create a database failover object and add it to a cluster failover group. (See Chapter 3, Configuring Database Software for Failover, for more information on preconfiguring your database software.)

   If you are running Oracle7 Workgroup Server, you also must make sure 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 Microsoft SQL Server or Oracle failover object dialog box is correct. Also make sure that you can start and stop the database server manually with the database administrator using the information in the failover object dialog box.

   Note also that if you are using Oracle7 Workgroup 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 Failover Manager trace log for login error messages or indications of problems during cluster startup. If you are running Microsoft SQL Server, you also should 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. If you do not see the shared disks under the Cluster Administrator Class View on the remote system, the problem is *not* a database problem. Refer to the section, Configuration Problems, earlier in this chapter, to troubleshoot the problem.

If the cluster group fails over properly without the database failover object, run the database administrator for your software and enumerate your databases.

- If your database is listed as “offline,” then the database disk was probably 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. Refer to the documentation that came with the Microsoft SQL Server software 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 more information.

Note

If you are using Microsoft SQL Server, but are not running Version 6.5 with Service Pack 1, you might need to run patch script `fallbac3.sql` after starting the MSSQLServer service to enable your databases to be enrolled properly. See Chapter 3 for information on how to run this script.

My Oracle7 Workgroup Server won’t fail over to the other server system.

If you are running Oracle7 Workgroup Server, and your database server will not fail over to the other cluster server, make sure that all the data files for the database server are located on shared disks and all the services for the database server are set to auto start on *both* cluster servers.

My Oracle7 Workgroup Server is running but the client can’t access it.

If you are running Oracle7 Workgroup 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. Digital Clusters for Windows NT supports only the named pipe protocol for automatic failover.

Appendix 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

Windows NT and Digital Clusters for Windows NT provide a variety of software utilities that are useful in troubleshooting your system.

Regedt32

The cluster software stores configuration data in the Windows NT Registry. You can use the Registry editor (located at `\winnt35\system32\regedt32.exe`) to examine, or modify, or both, 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 that came with your Windows NT operating system.

See the section Registry, later in this appendix, 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 given 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 to verify that the cluster services are running. The required cluster services are the following:

- Cfmd Server
- Cluster Failover Manager
- Cluster Name Service

NET SHARE Command

You can use the NET SHARE command to verify that a particular server has exported the proper shares. To display the names of all file shares currently being exported by your server, do the following:

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. To display the shares that are being exported by a particular server, do the following:

1. Open an MS–DOS command prompt window.
2. Type the following command, where *server* is the name of the server you are interested in:

```
net view \\server
```

You also can get the same information by selecting Disk | Connect Network Drive from Windows NT File Manager and browsing the Shared Directories box for the server.

NETMON Network Monitor

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 the following information:

- The cluster announcement is being sent from the server to the client. (A cluster client cannot "see" the Cluster Name Server if it is unable to receive the cluster announcement.)
- Requests and responses are being passed between the client and the Cluster Name Server.

Digital Clusters for Windows NT provides an extension DLL for the System Management Server network monitor program, NETMON, that formats the Cluster Name Service protocol. To enable this extension DLL, do the following on one or both of your cluster servers:

1. Put the file `cnsmon.dll` in the `parsers` subdirectory of NETMON.
2. Using a text editor, edit file `parser.ini` as follows:
   - Add `cnsmon.dll` to the list of DLLs at the beginning of the file by adding the following line:

     ```
     CNSMON.DLL  = 0: CNS
     ```

   - 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
     ```

   - Add the following section to the end of the file:

     ```
     [CNS]
        Comment = "Cluster Name Service protocol"
        FollowSet    =
        HelpFile     =
     ```

3. Use the Network Control Panel to verify that a matching protocol is enabled on the client system.
4. If the client is running TCP/IP, ensure that it is in the same subnet as the server system.

CLUIVP

The cluster distribution CD–ROM contains an unsupported utility (CLUIVP) that you can use to verify your cluster installation. CLUIVP 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 servers. After rebooting both servers, you must stop the cluster Failover Manager on both servers. This puts all shared disks off line.

CLUIVP is 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 cluster failover management database (Cfmd)
- Disk response and conformance to SCSI protocols required by the cluster software

CLUSTAT and FMSTAT

The cluster distribution CD–ROM includes two utility programs (CLUSTAT and FMSTAT) that you can use to determine how your cluster is allocating resources. These programs are dynamic status monitors that show where the cluster has relocated its resources. Note that these utilities are unsupported and undocumented.

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.

Registry Keys

The registry keys used by the cluster software are described briefly in the following sections. Refer to Appendix B for an example of the contents of each key.

Cluster Configuration Management Database (Cfmd) Key

The cluster configuration management database (`Cfmd`) key in the Registry 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
```

Cluster Port Driver (CluPort) Key

The cluster port driver (`CluPort`) key in the Registry 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
```

Cluster Disk Driver (CluDisk) Key

The cluster disk driver (`CluDisk`) key in the Registry contains values used by the Windows NT operating system to start the cluster SCSI disk driver when the system is booted.

The `CluDisk` key is found at:

```
\Registry\Machine\System\CurrentControlSet\Services\CluDisk
```

Cluster File System (Cfs) Key

The cluster file system (`Cfs`) key in the Registry contains values used by the Windows NT operating system to start the cluster file system driver when the system is booted.

The `Cfs` key is found at:

```
\Registry\Machine\System\CurrentControlSet\Services\Cfs
```

Log Watch Key

The log watch (`LogWatch`) key in the Registry 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
```

Cluster Failover Manager Key

The cluster failover manager (`ClusterFailoverManager`) key in the Registry contains entries that are used to control the timing of certain cluster operations, as well as entries specifying the location of the cluster trace log files. (Some of these entries are explained in the section Tuning Parameters, later in this appendix.)

The `ClusterFailoverManager` key also contains values used by the Windows NT operating system to start the cluster Failover Manager when the system is booted.

The `ClusterFailoverManager` key is found at:

```
\Registry\Machine\System\CurrentControlSet\Services
   \ClusterFailoverManager
```

Cluster Name Service Key

The cluster name service (`ClusterNameService`) key in the Registry contains the cluster name cache—a list of the names of all the clusters in the domain to which the local cluster belongs.

The `ClusterNameService` key is found at:

```
\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
```

Tuning Parameters

The Registry contains several parameters that affect the behavior of your cluster. Some of these parameters are described in the following sections.

Cluster parameters are stored in a subkey (`Parameters`) of the keys for their respective services Cfmd and Failover Manager. To examine and modify these parameters, you use the Registry editor, located at `\winnt35\system\32\regedt32.exe`.

Note that any modifications you make to the Registry entries do not take effect until you reboot your system.

CfmdTrace

The `CfmdTrace` parameter is used to control the location of the Cfmd trace file. The trace file contains information used by Digital support personnel.

Initially this parameter is not defined and tracing is not enabled. To enable tracing, do the following:

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 using Edit | Add Key. The Class Name field should be left blank.

2. Create the `CfmdTrace` parameter using Edit | Add Value. 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 so, Cfmd 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 Cfmd server to create log files named `cfmd001.log`, `cfmd002.log`, and so on.

To disable tracing, use the Registry editor to access the key, then select and delete the parameter using Edit | Delete Value.

CisTrace

The `CisTrace` parameter is used to control the location of the cluster infrastructure communication trace file. The trace file contains information used by Digital support personnel.

Initially, this parameter is not defined and tracing is not enable. To enable tracing, do the following:

1. Using the Registry editor, access the following key:

```
\Registry\Machine\System\CurrentControlSet\Services\
   ClusterFailoverManager\Parameters
```

2. Create the `CisTrace` parameter using Edit | Add Value. 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 so, the 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 Failover Manager to create log files named `cis001.log`, `cis002.log`, and so on.

To disable tracing, use the Registry editor to access the key, then select and delete the parameter using Edit | Delete Value.

ConnectionTimeout

The `ConnectionTimeout` parameter is used to control how long (in seconds) 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 seconds is used. To change the timeout value, do the following:

1. Using the Registry editor, access the following key:

   ```
   \Registry\Machine\System\CurrentControlSet\Services\
      ClusterFailoverManager\Parameters
   ```

2. If no `ConnectionTimeout` parameter exists, create it using Edit | Add Value. Set its data type to REG_DWORD.
3. Set the parameter value to the desired timeout, in seconds.
4. To revert to the default of 30 seconds, select and delete the parameter using Edit | Delete Value.

FailoverEvaluateDelay

The `FailoverEvaluateDelay` parameter 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. To change the timeout value, do the following:

1. Using the Registry editor, access the following key:

   ```
   \Registry\Machine\System\CurrentControlSet\Services\
      ClusterFailoverManager\Parameters
   ```

2. If no `FailoverEvaluateDelay` parameter exists, create it using Edit | Add Value. Set its data type to REG_DWORD.
3. Set the parameter value to the desired timeout, in seconds.
4. To revert to the default, select and delete the parameter using Edit | Delete Value.

ReconnectWait

The `ReconnectWait` parameter is used to control how long (in seconds) the communications infrastructure waits before trying to recontact 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. To change the timeout value, do the following:

1. Using the Registry editor, access the following key:

   ```
   \Registry\Machine\System\CurrentControlSet\Services\
      ClusterFailoverManager\Parameters
   ```

2. If no `ReconnectWait` parameter exists, create it using Edit | Add Value. Set its data type to REG_DWORD.
3. Set the parameter value to the desired delay, in seconds.
4. To revert to the default value, select and delete the parameter using Edit | Delete Value.

LogLevel

The `LogLevel` parameter 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 this 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. To change the types of messages logged, do the following:

1. Using the Registry editor, access the following key:

   ```
   \Registry\Machine\System\CurrentControlSet\Services\
      ClusterFailoverManager\Parameters
   ```

2. If no `LogLevel` parameter exists, create it using Edit | Add Value. Set its data type to REG_DWORD.
3. Set the parameter value to the desired level (4, 3, 2, or 1).
4. To revert to the default, select and delete the parameter using Edit | Delete Value.

ClusterName

The `ClusterName` parameter 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 is used to turn off the stabilization delay imposed by the Failover Manager after a server is declared down before failing over to the other server. This delay allows the 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, this parameter is not defined, thereby enabling the stabilization delay. To disable the delay, do the following:

1. Using the Registry editor, access the following key:

   ```
   \Registry\Machine\System\CurrentControlSet\Services\
      ClusterFailoverManager\Parameters
   ```

2. If no `DisableFailoverNetDelay` parameter exists, create it using Edit | Add Value. Set its data type to REG_DWORD.
3. Set the parameter value to any nonzero number.

To reenable the delay, select and delete the parameter using Edit | Delete Value or set the value to 0.

Note that the length of the stabilization delay is fixed and cannot be changed.

DiskArbitrationInterval

The `DiskArbitrationInterval` parameter is used to specify (in seconds) how frequently the server polls the disk to verify its ownership and how long the secondary 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. To change the arbitration interval, do the following:

1. Using the Registry editor, access the following key:

   ```
   \Registry\Machine\System\CurrentControlSet\Services\
      ClusterFailoverManager\Parameters
   ```

2. If no `DiskArbitrationInterval` parameter exists, create it using Edit | Add Value. Set its data type to REG_DWORD.
3. Set the parameter value to the desired interval, in seconds.
4. To revert to the default of 30 seconds, select and delete the parameter using Edit | Delete Value.

DiskErrorThreshold

The `DiskErrorThreshold` parameter 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. To change the threshold value, do the following:

1. Using the Registry editor, access the following key:

   ```
   \Registry\Machine\System\CurrentControlSet\Services\
      ClusterFailoverManager\Parameters
   ```

2. If no `DiskErrorThreshold` parameter exists, create it using Edit | Add Value. Set its data type to REG_DWORD.
3. Set the parameter value to the desired threshold.
4. To revert to the default of 3, select and delete the parameter using Edit | Delete Value.

DiskErrorSeparation

The `DiskErrorSeparation` parameter 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. To change the error separation interval, do the following:

1. Using the Registry editor, access the following key:

   ```
   \Registry\Machine\System\CurrentControlSet\Services\
      ClusterFailoverManager\Parameters
   ```

2. If no `DiskErrorSeparation` parameter exists, create it using Edit | Add Value. Set its data type to REG_DWORD.
3. Set the parameter value to the desired interval, in seconds.
4. To revert to the default value, select and delete the parameter using Edit | Delete Value.

FmTraceOutput

The `FmTraceOutput` parameter is used to control the location of the Failover Manager trace log. The parameter is a bit mask with the following value definitions:

1 Log to file (as named by parameter `FmTrace`)
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 FmTrace. To change the trace log output, do the following:

1. Using the Registry editor, access the following key:

   ```
   \Registry\Machine\System\CurrentControlSet\Services\
      ClusterFailoverManager\Parameters
   ```

2. If no FmTraceOutput parameter exists, create it using Edit | Add Value. Set its data type to REG_DWORD.
3. Set the parameter value to the desired number (1, 2, 4, or 8).
4. To revert to the default, select and delete the parameter using Edit | Delete Value.

FmTrace

The FmTrace parameter is used to specify the full path of the Failover Manager trace log, if trace log output is being directed to a file.

Initially, the 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 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 Failover Manager to create trace log files named fm001.log, fm002.log, and so on.

To change the name of the trace log file, do the following:

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 is used to control what types of messages are to be logged by the 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 this 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. To change the types of messages logged, do the following:

1. Using the Registry editor, access the following key:

   ```
   \Registry\Machine\System\CurrentControlSet\Services\
      ClusterFailoverManager\Parameters
   ```

2. If no `FmLogLevel` parameter exists, create it using Edit | Add Value. Set its data type to REG_DWORD.
3. Set the parameter value to the desired level (4, 3, 2, or 1).
4. To revert to the default, select and delete the parameter using Edit | Delete Value.

FmTraceVerbosity

The `FmTraceVerbosity` parameter is used to control the priority of messages logged to the Failover Manager trace log. The value indicates the level, as follows:

1	Errors and major events
2, 3	Minor events
4, 5	Individual communication protocol messages
6	Data structure dumps

Initially, the parameter is undefined and a default of 3 is used, enabling the logging of all major and minor events. To change the level of detail of messages logged, do the following:

1. Using the Registry editor, access the following key:

   ```
   \Registry\Machine\System\CurrentControlSet\Services\
      ClusterFailoverManager\Parameters
   ```

2. If no `FmTraceVerbosity` parameter exists, create it using Edit | Add Value. Set its data type to REG_DWORD.
3. Set the parameter value to the desired level (1 to 6).
4. To revert to the default, select and delete the parameter using Edit | Delete Value.

Trace Log

The 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:

- Cluster software version number and start time
- Names of shared disks discovered by the cluster 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 Failover Manager trace log file is specified by the `FmTrace` parameter. (See the section Tuning Parameters, earlier in this appendix, for more information.) By default, the file is found in the `temp` subdirectory of the cluster destination directory, as specified when the server software was installed. If you accepted the default directory during installation, the default 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. (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 cluster software maintains an event log that tracks 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.

Cluster components that run as services put their error messages in the Application portion of the event log. These components include the following:

- Failover Manager
- CluCis
- Cfmd
- FM Disk DLL

Cluster components that run as drivers put their error messages in the System portion of the event log. These components include the following:

- CluPort
- 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`*x* `Bus 0`

  This message indicates that the shared SCSI bus adapter is located on bus 0 or port *x*. It is a normal startup message. You can use it to verify that the cluster is using the proper adapter.

- `Warning! No Cluster Adapters found`

  This message indicates that no shared SCSI bus adapters have been detected during the scan of the hardware. This 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 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. Use the Manage Adapter Configuration function of the 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 ScsiPort`*`X`*

  This message indicates that the cluster port driver encountered an error while trying to take control of the specified SCSI bus adapter. Consult the event log for more information.

Appendix B
Registry Snapshot

This appendix contains snapshots of those parts of the Windows NT Registry that pertain to Digital Clusters for Windows NT.

Cluster Configuration 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
```

```
        _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)
                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
```

```
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"  \
                                "ntclua1"
        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"  \
                                "ntclua1"
        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"  \
                                "ntclua1"
        ObjectList = REG_MULTI_SZ "FMDisk\_disk_37a611de"
        Reevaluate = REG_DWORD 0x00000001
        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
    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
```

Cluster 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"  \
                                              "ntclua1"
                    Name = NTCLUA2tontclua1bus0
                    Checksum = REG_DWORD 0x00011395
            Scsi Port 2
                Scsi Bus 0
                    Initiators = REG_MULTI_SZ "NTCLUA2"  \
                                              "ntclua1"
                    Name = NTCLUA2tontclua1bus1
                    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
```

Cluster 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
```

Cluster 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
```

Cluster 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 = ntsgwest\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
```

Cluster 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

```
\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 = ncrc810
        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
                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
```

Appendix C
Failover Manager Trace Log

This appendix contains an annotated example of a typical Digital Clusters for Windows NT Failover Manager trace log.

The trace log header contains useful information such as the product version, build number, cluster name, nodename, 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 Failover Manager. Note that 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...
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 groups in the database are loaded. (In this case, there are seven groups.) Note the `tid` field. It can be used to follow threads of activity. This is very important when mutliple 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 this next section, the disk DLL is initialized. Note the phases marked step 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'
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
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\fmdisk\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\fmdisk\device.c    Line: 1328
```

In the next section, the Failover Manager finishes intialization. In this example, neither SQL nor Oracle is installed. If they were, 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
```

In the next section, the Failover Manager puts the groups on line. In this case, four of the groups are primary on this machine—groups 104, 106, 105, and 107. Each group has one disk. The Failover Manager tells the disk DLL to put the disk on line, and because each disk is in a separate 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

Note that in this example, group 104 contains the log disk. That group is brought on line first, and the 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 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\fmdisk\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\fmcore\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
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
```

Appendix D
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
   succcessfully 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

```
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".
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.
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.
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 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.

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**.

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.

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.

digital™

Digital Clusters for Windows NT

Administrator's Guide Addendum

Order Number: AV–R171A–TE

October 1996

This addendum provides updated information on installing and configuring Microsoft® SQL Server™ and Oracle7™ Workgroup Server with Digital Clusters for Windows NT.

Revision/Update Information:	This addendum supercedes Chapter 3 of the *Digital Clusters for Windows NT Administrator's Guide*, First Edition, AA–QVUTA–TE.
Operating System and Version:	Microsoft Windows NT Version 3.51 with Service Pack 5
Software Version:	Digital Clusters for Windows NT Version 1.0 with Service Pack 1

Digital Equipment Corporation
Maynard, Massachusetts

October 1996

Contents

ABOUT THIS ADDENDUM

CHAPTER 3 CONFIGURING DATABASE SOFTWARE FOR FAILOVER

About This Addendum

This addendum supplies updated information on installing and configuring Microsoft SQL Server and Oracle7 Workgroup Server with Digital Clusters for Windows NT Version 1.0 with Service Pack 1. Also included are recommendations, restrictions, and verification procedures. The information supercedes Chapter 3, Configuring Database Software for Failover of the *Digital Clusters for Windows NT Administrator's Guide*, First Edition.

Audience

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

Organization

This addendum consists of one chapter, as follows:

Chapter 3	Presents Microsoft SQL Server and Oracle7 Workgroup Server installation and configuration requirements in the Digital Clusters for Windows NT environment, recommendations, restrictions, and verification procedures.

Conventions

The following conventions are used in this addendum:

Convention	Meaning
Bold	Bold type indicates the actual commands, words, or characters that you type in a dialog box or at the command prompt.
Italic	Italic type indicates a placeholder for information on parameters that you must provide. For example, if the procedure asks you to type *filename*, 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*
- *Digital Clusters for Windows NT Administrator's Guide*
- Online help

Chapter 3
Configuring Database Software for Failover

This chapter presents Microsoft SQL Server and Oracle7 Workgroup 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.

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 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.

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 SQL Server with Digital Cluster for Windows NT:

- The cluster server systems 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 cluster client systems must be configured with the same network transports as the server systems. For example, if you use the TCP/IP network protocol, clients must be in the same IP subnet as the server systems and NetBIOS must be installed on each client. You can bypass the client subnet restriction if you have NetBEUI installed on both the clients and the servers.
- Both cluster servers must reside in the same Windows NT domain.

- Neither cluster server can be a member of any other cluster.
- The domain account ClusterAdmin must have the proper rights (that is, Domain Administrator, Administrator, and Login as Service).

Configuration and Run-Time Recommendations

Before you configure Microsoft SQL Server to run with Digital Clusters for Windows NT, 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.
- If you have not already done so, Digital strongly recommends that you install Service Pack 1 for SQL Server Version 6.5, which is located on the Digital Clusters for Windows NT installation CD–ROM in the directory `\SQL65sp1`. By installing Service Pack 1, you will not later have to run the script `fallbac3.sql`, discussed in the sections Running the SQL Server Patch Script, and Configuring SQL Server for Failover.
- 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.

Restrictions

The following Microsoft SQL Server restrictions have been identified in Digital Clusters for Windows NT Version 1.0 with Service Pack 1:

- You cannot install SQL Server while the Cluster Failover Manager and Cluster Failover Manager 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.
- You must configure the MSSQLServer service to start manually.

- Due to an SQL Server software limitation, you can run SQL Server on only one cluster server at a time. This cluster server is the one servicing the SQL Server databases on the shared disks. The alternate cluster server cannot be processing another SQL Server database, whether it be on a local or shared disk.

  However, you can configure the cluster software to run Oracle Server on one cluster server while SQL Server is running on the other, each accessing databases that reside on the shared disks. Similarly, you can run custom server applications on the other cluster server using the scripting capability offered by Digital Clusters for Windows NT. The custom server application can access data on the shared disks. See the section Working with a Script Failover Object in Chapter 7 for details.

- Do not use the ISQL/w application or Enterprise Manager Query Tool to register SQL Server using the cluster alias. SQL Server management functions should be done using individual server names. This restriction does not apply to database queries to the cluster alias.
- You can have only one failover group containing an SQL failover object per cluster.
- You can fail over multiple SQL Server databases, but these must be in the same failover group.
- When creating a failover group for an SQL failover object, you must specify only the disks that contain the SQL Server databases and optionally, the log disk for the cluster. The SQL failover object must be placed *last* in the failover group.

Running the SQL Server Patch Script

To use Microsoft SQL Server Version 6.5 with Digital Clusters for Windows NT Version 1.0 with Service Pack 1, you may need to run the script `fallbac3.sql` after starting the MSSQLServer service. This script, a patch from Microsoft to correct potential SQL Server failback problems, is provided on the Digital Clusters for Windows NT distribution CD–ROM.

You will need to run this patch script if you are not running Microsoft SQL Server Version 6.5 with Service Pack 1 *and* either of the following conditions exist:

- Microsoft SQL Server is installed on your server but the MSSQLServer service was not running at the time that you installed the cluster server software.
- Microsoft SQL Server was installed after the cluster server software.

If Microsoft SQL Server was installed and the MSSQLServer service was running when you installed the cluster server software, the Setup program will have run the script automatically and you need do nothing further.

Configuring SQL Server for Failover

This section presents the steps you must perform to configure Microsoft SQL Server for failover.

To configure the SQL Server software for failover:

1. Install Microsoft SQL Server on a local disk on each cluster server, configuring the MSSQLServer service to start manually.

Caution

For Digital Clusters for Windows NT Version 1.0 with Service Pack 1, do not change the default named pipe name for SQL 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.

   c. Go to step 9.

3. Install the cluster software on each server system. See the *Digital Clusters for Windows NT Configuration and Installation Guide* for instructions.

4. Install a new SQL Server stored procedure:

   a. Using the Services applet of the Windows NT Control Panel, check if the MSSQLServer service is running. If not, start the service.

   b. Return to Program Manager and open an MS–DOS window.

   c. At the MS–DOS command prompt, type:

      ```
      isql -Usa -P -S -n -i
      "cluster-directory\sp_fallback_DEC_perm_svr_db.sql"
      ```

      where *cluster-directory* is the drive and path name of the destination directory that you specified during the cluster server software installation. If you used the default directory, this drive and path name will be:

      ```
      c:\Program Files\Digital\Cluster
      ```

      Note that the quotes in the command string are required.

5. Run the SQL Server Version 6.5 patch script, as outlined here, if your configuration falls into either of the categories described in the previous section, Running the SQL Server Patch Script:

   a. Using the Services applet of the Windows NT Control Panel, check if the MSSQLServer service is running. If not, start the service.

   b. Return to Program Manager and open an MS–DOS window.

   c. At the MS–DOS command prompt, type:

      ```
      isql -Usa -P -S -n -i "cluster-directory\fallbac3.sql"
      ```

      where *cluster-directory* is the drive and path name of the destination directory that you specified during the cluster server software installation. If you used the default directory, this drive and path name will be:

      ```
      c:\Program Files\Digital\Cluster
      ```

      Note that the quotes in the command string are required.

6. On the primary server, use Cluster Administrator to create a failover group containing all the shared disks to be used for Microsoft SQL Server. See the section Creating a Failover Group in Chapter 7 for instructions.

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

Caution

When creating a failover group for an SQL failover object, you must specify only the disks that contain the SQL Server databases and optionally, the log disk for the cluster. The SQL failover object must be placed *last* in the failover group.

7. On each cluster server, assign identical fixed drive letters for each shared disk used for SQL Server database failover:

   On the Primary Server:

   a. On the primary server for the shared disks used for SQL Server database failover, choose the Administrative Tools program group in Program Manager.

   b. Choose the Disk Administrator icon.

   c. Select a shared disk for which you want to modify the drive-letter assignment.

   d. From the Tools menu, choose Drive Letter.

      The Assign Drive Letter dialog box is displayed.

   e. 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.

   f. Repeat steps c to e for each shared disk used for SQL Server database failover.

   g. Manually fail over the disk group to the failover server. See the section Managing Manual Failover in Chapter 6 for instructions.

   On the Failover Server:

   h. Once failover has completed, repeat steps a to f on the failover server. Digital recommends that you assign the same fixed drive letters on both server systems.

   i. Manually fail back the disk group to the primary server. See the section Managing Manual Failover in Chapter 6 for instructions.

On the Primary Server:

j. Using Disk Administrator, verify that the shared disks have failed over. Disks not controlled by the primary server will be designated as OFF-LINE.

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

9. Configure the server systems and databases for high availability:

   a. Start the MSSQLServer service on each cluster server.

   b. *On the Primary Server:*

      Invoke the next stored procedures supplied with the Microsoft SQL Server Client-Server Database Management System for Windows NT product. You can invoke these procedures using either the ISQL/w application or the Enterprise Manager Query Tool:

      - Record the failover server:

        ```
        sp_addserver FailoverServerName, fallback
        ```

      - Give the system administrator (sa) account on the failover server permission to log in as the sa on the primary server:

        ```
        sp_addremotelogin FailoverServerName, sa, sa
        ```

      - Enroll all databases on the shared disks owned by the primary server for potential failback support:

        ```
        sp_fallback_enroll_svr_db PrimaryServerName,
        DatabaseName
        ```

   c. *On the Failover Server:*

      Invoke the next stored procedures:

      - Record the primary server:

        ```
        sp_addserver PrimaryServerName
        ```

      - Give the sa account on the primary server permission to log in as the sa on the failover server:

        ```
        sp_addremotelogin PrimaryServerName, sa, sa
        ```

      - Verify that the RPC connections are functioning:

        ```
        PrimaryServerName...sp_helplogins
        ```

- Enroll all databases on the shared disks owned by the primary server for potential failback support:

  ```
  sp_fallback_enroll_svr_db PrimaryServerName,
  DatabaseName
  ```

10. If you have not already done so, install the cluster software on each server system. See the *Digital Clusters for Windows NT Configuration and Installation Guide* for instructions.
11. Install a new SQL Server stored procedure:
    a. Using the Services applet of the Windows NT Control Panel, check if the MSSQLServer service is running. If not, start the service.
    b. Return to Program Manager and open an MS–DOS window.
    c. At the MS–DOS command prompt, type:

       ```
       isql -Usa -P -S -n -i
       "cluster-directory\sp_fallback_DEC_perm_svr_db.sql"
       ```

       where *cluster-directory* is the drive and path name of the destination directory that you specified during the cluster server software installation. If you used the default directory, this drive and path name will be:

       ```
       c:\Program Files\Digital\Cluster
       ```

       Note that the quotes in the command string are required.
12. Run the SQL Server Version 6.5 patch script, as outlined here, if your configuration falls into either of the categories described in the previous section, Running the SQL Server Patch Script:
    a. Using the Services applet of the Windows NT Control Panel, check if the MSSQLServer service is running. If not, start the service.
    b. Return to Program Manager and open an MS–DOS window.
    c. At the MS–DOS command prompt, type:

       ```
       isql -Usa -P -S -n -i "cluster-directory\fallbac3.sql"
       ```

       where *cluster-directory* is the drive and path name of the destination directory that you specified during the cluster server software installation. If you used the default directory, this drive and path name will be:

       ```
       c:\Program Files\Digital\Cluster
       ```

       Note that the quotes in the command string are required.
13. Manually stop the MSSQLServer service on each server system.

14. Using Cluster Administrator, create an SQL failover object. See the section Creating an SQL Failover Object in Chapter 7 for instructions.

15. If you have not already done so, use Cluster Administrator on the primary server to create a failover group containing all the shared disks to be used by SQL Server and the SQL Server failover object that you created in step 14. The SQL failover object must be placed *last* in the failover group. See the section Creating a Failover Group in Chapter 7 for instructions.

    Otherwise, modify the failover group that you created in step 6 to include the SQL failover object that you created in step 14. The SQL failover object must be placed *last* in the failover group. See the section Modifying a Failover Group in Chapter 7 for instructions.

Caution

When creating or modifying a failover group for an SQL failover object, you must specify only the disks that contain the SQL Server databases and optionally, the log disk for the cluster.

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 key:

      ```
      HKEY_LOCAL_MACHINE/System/CurrentControlSet/
      Services/Cfmd/Database/Pipes/SQL
      ```

   b. Expand the key 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.

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. Using either the ISQL/w application or the Enterprise Manager Query Tool, use the cluster alias, 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 the Digital Customer Support Center closest to you. You will find a list of worldwide Digital Customer Support Centers and associated telephone numbers in the cover letter of your Digital Clusters for Windows NT kit.

On the Primary Server:

5. Bring the primary server on line again.
6. 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:

7. Repeat step 4.

Modifying an SQL Failover Group

If, after creating or modifying a failover group that contains an SQL failover object, you want to add another shared disk or database to the failover group, you must follow the instructions in the next procedure.

▶ **To modify a failover group that contains an SQL failover object:**

1. Use Cluster Administrator to check if the failover group is on line on the primary server. If not, manually fail over the failover group. See the section Managing Manual Failover in Chapter 6 for instructions.
2. Add the new shared disk to the failover group. See the section Modifying a Failover Group in Chapter 7 for instructions.
3. Use a database administrator tool to create additional SQL Server databases (and optionally, the transaction log) on the shared disk.
4. Configure the new databases for high availability:
   a. Start the MSSQLServer service on each cluster server.
   b. Enroll the new databases on the shared disk for potential failback support by running the following stored procedure on each cluster server. You can run this procedure using either the ISQL/w application or the Enterprise Manager Query Tool:

      ```
      sp_fallback_enroll_svr_db PrimaryServerName,
      DatabaseName
      ```

   c. Manually stop the MSSQLServer service on each server system.
5. Remove the SQL failover object (created in step 14 of the procedure in the section Configuring SQL Server for Failover) from the failover group. See the section Modifying a Failover Group in Chapter 7 for instructions.

Caution

Removing the SQL failover object from the failover group will place all shared SQL Server databases off line. Users will not be able to access SQL Server using the cluster alias.

6. Delete the SQL failover object. See the section Deleting an SQL Failover Object in Chapter 7 for instructions.
7. Create a new SQL failover object. See the section Creating an SQL Failover Object in Chapter 7 for instructions.

8. Add the new SQL failover object to the failover group. See the section Modifying a Failover Group in Chapter 7 for instructions.

Caution

When creating or modifying a failover group for an SQL failover object, you must specify only the disks that contain the SQL Server databases and optionally, the log disk for the cluster. The SQL failover object must be placed *last* in the failover group.

Withdrawing an SQL Server Database from Failover Support

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

▶ **To remove an SQL Server database from cluster failover support:**

1. Check if the SQL Server database is on line on the primary server. If not, use Cluster Administrator to manually fail over the disk. See the section Managing Manual Failover in Chapter 6 for instructions.
2. 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.
3. On the failover server, execute the following stored procedure:

   ```
   sp_fallback_withdraw_svr_db PrimaryServerName,
   DatabaseName
   ```

4. On the primary server, execute the next stored procedure:

   ```
   sp_fallback_DEC_perm_svr_db PrimaryServerName,
   DatabaseName
   ```

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 withdrawing an SQL Server database does not affect primary server access.

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, Removing 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 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 7 for instructions.
3. Use a database administrator tool to expand or shrink the SQL Server database.
4. Reenroll the database for failover support by executing the following stored procedure on *each* cluster server:

   ```
   sp_fallback_enroll_svr_db PrimaryServerName, DatabaseName
   ```

5. If you added or deleted any shared disks from the failover group referenced in step 2, you must re-create the SQL failover object using Cluster Administrator:
   a. Remove the SQL failover object from the failover group. See the section Modifying a Failover Group in Chapter 7 for instructions.

   Caution

   Removing the SQL failover object from the failover group will place all shared SQL Server databases off line. Users will not be able to access SQL Server using the cluster alias.

   b. Delete the SQL failover object. See the section Deleting an SQL Failover Object in Chapter 7 for instructions.
   c. Create a new SQL failover object. See the section Creating an SQL Failover Object in Chapter 7 for instructions.

d. Add the SQL failover object to the failover group. See the section Modifying a Failover Group in Chapter 7 for instructions.

Caution

The SQL failover object must be placed *last* in the failover group.

Oracle7 Workgroup 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. 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 Workgroup Server in the Digital Clusters for Windows NT environment.

Prerequisite Information

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

Software Requirements

Software requirements are as follows:

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

Oracle Server Database Failover

Throughout the Digital Clusters for Windows NT documentation, *Oracle Server database failover* refers to an Oracle7 Workgroup Server instance failing over. Oracle7 Workgroup 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.

If an Oracle instance fails while the server system on which it is running continues to function properly, database administrator intervention may be required. Once the problem has been remedied, the database administrator can restart the instance.

Configuration Requirements

For an Oracle instance to fail over, the following prerequisites must be met:

- 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 file, control file, log file, and parameter file, must reside on one or more shared disks.
- The shared disks and Oracle failover object must be in the same failover group.

Creating an Oracle Instance for Failover

This section presents instructions on how to create an Oracle instance for failover.

▶ **Before creating an Oracle instance for failover:**

1. If you have not already done so, install the cluster software.
2. 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, choose the Administrative Tools program group in Program Manager.

   b. Choose the Disk Administrator icon.

   c. Select a shared disk for which you want to modify the drive-letter assignment.

   d. From the Tools menu, choose Drive Letter.

      The Assign Drive Letter dialog box is displayed.

   e. 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.

   f. Repeat steps c to e for each shared disk used for Oracle instance failover.

g. Manually fail over the disk group to the failover server. See the section Managing Manual Failover in Chapter 6 for instructions.

On the Failover Server:

h. Once failover has completed, repeat steps a to f on the failover server. Digital recommends that you assign the same fixed drive letters on both server systems.

i. Manually fail back the disk group to the primary server. See the section Managing Manual Failover in Chapter 6 for instructions.

On the Primary Server:

j. Using Disk Administrator, verify that the shared disks have failed over.

3. Using Cluster Administrator, create a failover group containing all shared disks to be used for Oracle7 Workgroup Server. See the section Creating a Failover Group in Chapter 7 for instructions.

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

Repeat the following procedure for *each* instance that you configure for failover.

▶ **To create an Oracle instance for failover:**

On the Primary Server:

1. Create an Oracle instance for failover using an Oracle database administrator tool.
2. Use a shared disk to store all files associated with the instance, including the data file, control file, log file, and parameter file.
3. Create an associated database on one or more shared disks. If you are working with an existing database, move the database data files to the shared disks.
4. *For Version 7.2 only:*

   If you are running Oracle7 Workgroup Server Version 7.2, perform the following steps for *each additional instance after the first instance* that you configure for failover:

   a. Create a separate network listener for the instance using Oracle Network Manager for Windows. See the Oracle Network Manager for Windows product documentation for instructions.

   b. Specify Named Pipe as the network protocol for each listener.

   c. Specify a unique name for the named pipe associated with each network listener.

   d. Start the network listener using the Services icon on the Control Panel.

5. Verify that you can start the server instance.

6. Verify that you can access the associated database using an Oracle database administrator tool. If you cannot access the associated 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.

7. Shut down the instance using an Oracle database administrator tool. Be sure to leave services for the instance running.

8. Verify that services for the instance are set to restart automatically using the Services icon on the Control Panel.

9. Manually fail over the disk group to the failover server. See the section Managing Manual Failover in Chapter 6 for instructions.

On the Failover Server:

10. Once the failover server can access the shared disks, repeat steps 1 to 8 to create an *identical* Oracle instance to the one you created on the primary server.

11. Manually fail back the disk group to the primary server. See the section Managing Manual Failover in Chapter 6 for instructions.

On the Primary Server:

12. Create an Oracle failover object, supplying information about the instance that you created in step 1. See the section Creating an Oracle Failover Object in Chapter 7 for instructions.

13. Modify the disk failover group that you created in the previous procedure to include the Oracle failover object that you created in step 12. See the section Modifying a Failover Group in Chapter 7 for instructions.

14. 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 12 is correct.

15. Verify that the Oracle Server instance can fail over properly:

    a. Initiate a manual failover of the Oracle Server instance. See the section Managing Manual Failover in Chapter 6 for instructions.

    b. Verify that the Oracle Server instance is running on 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 Workgroup Server software. This guarantees orderly shutdown of the instance. Then use Cluster Administrator to initiate a manual failover.

▶ **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 6 for instructions.