



File Director™

Administrator Guide

NeoPath Networks, Inc.
2171 Landings Drive
Mountain View, CA 94043
+1 650.691.7700 (Voice)
+1 650.961.6108 (Fax)
<http://www.neopathnetworks.com>

support@neopathnetworks.com
+1 866-NEOPATH (+1 866-636-7284)

Important Notice

NeoPath™ Networks File Director™

© 2005 NeoPath Networks. All rights reserved.

NeoPath and File Director are trademarks of NeoPath Networks in the U.S.A. and/or other countries.

All other company, brand, or product names used herein may be trademarks or registered trademarks of their respective companies or organizations.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of NeoPath Networks.

NeoPath Networks may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from NeoPath Networks, the furnishing of this document does not give you any license to these patents, trademarks copyrights, or other intellectual property.

The File Director is delivered with certain independent code that is licensed under the GNU General Public License (“GPL”), the GNU Library General Public License (“LGPL”), and/or other open source licenses (“Open Source Code”), and Open Source Code is licensed to Customer in accordance with the applicable open source license. You may obtain a copy of the GPL at <http://www.gnu.org/licenses/gpl.html>, and of the LGPL at <http://www.gnu.org/licenses/lgpl.html>. For a period of three years from the date of your purchase of the File Director, NeoPath Networks will at your request provide a copy of the source code for the code licensed under the GPL and LGPL, including any modifications made by NeoPath Networks. Notwithstanding anything to the contrary, to the extent that any of the terms and conditions of this Agreement conflict with any such open source licenses, the conflicting terms and conditions shall not apply to the corresponding Open Source Code.

For license and copyright information for third party software included with the File Director, see Appendix B, “Third-Party Software License and Copyright Information.”

The information in this document is subject to change without notice. NeoPath Networks shall not be liable for any damages resulting from technical errors or omissions which may be present in this document, or from use of this document.

Version: 1.0.7

Updated: April 2005

CONTENTS

Audience.....	ix
Contacting NeoPath Networks Technical Support	ix
How to Use This Guide.....	ix
Conventions.....	x
Getting Online Help	xi

Chapter 1: Introducing the File Director

Understanding the Problems of Network Attached Storage Servers	1-1
What is the File Director?	1-2
File Director Functionality and Benefits	1-4
Examples of Scenarios for Deploying the File Director	1-5
Scenario 1: Balance Storage Capacity Among Legacy File Servers	1-5
Scenario 2: Transparently Add Storage and Present a Single Hierarchy.....	1-9
Scenario 3: Deploy the File Director and New File Servers and Clients into a New Environment	1-13
What are the Tasks for Configuring and Using the File Director?	1-16
Required Configuration Tasks.....	1-16
Ongoing Operational Tasks	1-16

Chapter 2: Configuring the File Director for Your Environment

Using the Management Console to Configure the File Director	2-1
Logging Into the File Director Management Console	2-1
Setting the Amount of Information Displayed in the Management Console.....	2-3
Sorting Information in the Management Console	2-3
Filtering Information in the Management Console.....	2-4
Adding Administrator Accounts	2-4
Using the Command Line Interface to Configure the File Director	2-5
Configuring the File Director for Your Network	2-5
Configuring the File Director IP Addresses.....	2-6
Configuring the File Director Network Interfaces	2-8
Setting the Default Gateway for the File Director	2-8
Setting a Proxy IP Address for the File Director.....	2-9
Specifying a DNS Server and DNS Search Domain.....	2-10
Specifying a WINS Server	2-10
Changing the File Director Host Name.....	2-11
Setting the File Director Date, Time, and Time Zone	2-11
Specifying an NTP Server	2-12

Setting Up the File Director to Access File Servers	2-13
Overview of Accessing Shares and Exports with the File Director	2-13
Specifying a File Server	2-16
Displaying the File Server Access Configuration	2-18
Displaying and Matching Shares and Exports for a File Server	2-18
Displaying Shares and Exports for a File Server	2-18
Matching or Unmatching Physical Shares and Exports	2-19
Specifying Virtual Servers and Virtual Shares.....	2-20
Specifying a Virtual Server.....	2-20
Specifying a Virtual Share or Virtual Export	2-21
Configuring ACL Entries for an NFS Virtual Export	2-23
Configuring Clients	2-26

Chapter 3: Working with Synthetic Directories

About Synthetic Directories, Synthetic Links, and Union Directories	3-1
Scenarios for Setting Up and Using Synthetic Directories	3-2
Synthetic Directory Scenario 1: Export Aggregation	3-2
Synthetic Directory Scenario 2: Server Aggregation.....	3-4
Synthetic Directory Scenario 3: Union Directory	3-5
Creating and Using Synthetic Directories, Synthetic Links, and Union Directories	3-8
Creating a Synthetic Directory.....	3-8
Creating Synthetic Links for a Synthetic Directory.....	3-9
Creating Union Directories.....	3-11
Creating a Virtual Share or Export for a Synthetic Directory	3-13
Deleting a Synthetic Directory	3-15
Using the CLI to Work with Synthetic Directories.....	3-16
About Directory Storage	3-16
Procedure: Adding a Synthetic Directory	3-17

Chapter 4: Using the File Director to Migrate Data

About Migrating Data	4-1
Configuring File Server Security Before Migrating Data	4-2
General CIFS Migration Guideline	4-2
Network Appliance CIFS Migration Guidelines.....	4-2
Samba CIFS Migration Guidelines	4-3
General Multi-protocol Migration Guidelines	4-4
Windows NFS Multi-protocol Migration Guidelines	4-5
Network Appliance Multi-protocol Migration Guideline	4-5
Samba and NFS Multi-protocol Migration Guidelines	4-6
Creating a Migration Job	4-6

Checking the Status of a Migration Job that is Running	4-9
Displaying a List of Migration Jobs	4-10
Displaying the File Migration Log	4-11
Retrying a Failed Migration Job	4-11
Displaying Migration Statistics for a Physical Share or Export	4-12
Finding the Current Location of Files or Directories	4-13
Finding the Location of a Physical File or Directory	4-13
Finding the Location of a Virtual File or Directory	4-14

Chapter 5: Using Policies for Automatic File Migration

About Policies	5-1
Summary of Steps for Setting Up and Using a Policy	5-1
Policy Execution	5-3
Using and Creating Schedules	5-4
Monitoring Traffic for Source Shares or Exports	5-6
Working with Policy Rules	5-7
Creating a Policy Rule	5-8
Inserting a Condition into a Rule	5-9
Saving a Policy Rule	5-10
Reverting a Workpad Rule to a Saved Rule	5-11
Combining Conditions in a Rule	5-11
Nesting Conditions in a Rule	5-12
Deleting a Condition from a Rule	5-16
Deleting a Saved Rule	5-16
Editing a Saved Rule	5-17
Working with Policies	5-18
Creating and Running Policy	5-18
Rehearsing a Policy	5-19
Editing a Policy	5-20
Deleting a Policy	5-20
Displaying the Policy Execution Log	5-20

Chapter 6: Monitoring and Troubleshooting the File Director

Configuring SNMP on the File Director	6-1
Configuring an SNMP Trap Community	6-2
About the SNMP Traps for the File Director	6-3
Deleting a Trap Community	6-4
Configuring an SNMP Access Community	6-4
Deleting an Access Community	6-5
Downloading and Saving the MIB File from the File Director	6-5

Displaying Log Files for the File Director	6-6
Displaying the System Log File.....	6-6
Displaying the Error Log	6-6
Displaying the Accumulator Log	6-7
Displaying the System Debug Log	6-7
Displaying the Switching Service Log.....	6-7
Configuring Log File Rotation	6-8
Displaying Traffic Statistics for the File Director	6-9
Displaying NFS and NLM Traffic Statistics.....	6-9
Displaying CIFS Traffic Statistics	6-10
Displaying Traffic Statistics for the File Director Network Interfaces.....	6-10
Troubleshooting the File Director.....	6-11
Displaying the Status of the File Director Network Interfaces.....	6-11
Performing a Ping Test from the File Director	6-12
Displaying a Routing Table from the File Director	6-12
Restarting or Stopping System Services on the File Director	6-13
Removing a Server	6-14
Restarting and Shutting Down the File Director	6-15
Backing Up and Restoring Configuration Information for the File Director	6-15
Backing Up the File Director Configuration.....	6-15
Restoring the File Director Configuration.....	6-17
Updating and Rolling Back the Software on the File Director	6-19
Updating the Software on the File Director	6-19
Rolling Back the Software on the File Director.....	6-20

Chapter 7: Setting Up a High Availability Cluster of File Directors

Overview of High Availability Clusters.....	7-1
Primary and Backup Nodes.....	7-1
Active and Standby Status	7-1
Heartbeat Communications and Transfer of Service.....	7-2
Configuration Changes	7-2
Setting Local or Cluster Scope	7-2
Configuration Guidelines for a Cluster.....	7-4
Summary of Configuration Settings	7-5
Summary of Steps for Setting Up a Cluster.....	7-6

Scenario A: Configuring Two New, Unconfigured File Directors as a Cluster	7-9
Task A-1: Run the Initial Configuration Script	7-9
Task A-2: Connect the Heartbeat Interfaces	7-9
Task A-3: Configure IP Addresses on the Primary	7-10
Task A-4: Configure the Multicast Base Address and Set a Backup Node for the Cluster	7-12
Task A-5: Configure Cluster-Wide Settings	7-14
Task A-6: Configure the Backup Node	7-14
Task A-7: Complete User Account and Network Configuration	7-16
Task A-8: Verify the Cluster Configuration	7-16
Scenario B: Configuring an Existing File Director into a Cluster with Another New or Existing File Director	7-17
Task B-1: Verify Stand-Alone Configuration	7-17
Task B-2: Connect the Heartbeat Interfaces	7-18
Task B-3: Configure IP Addresses	7-19
Task B-4: Configure Node Roles and the Multicast Base Address of the Cluster	7-20
Task B-5: Configure the Remaining Node	7-22
Task B-6: Confirm User Account and Network Configuration	7-24
Task B-7: Verify the Cluster Configuration	7-24
Checking the Status of Each Cluster Node	7-25
Troubleshooting Problems with Setting Up a Cluster	7-26
Disbanding a Cluster	7-26
Disbanding a Cluster If Both Nodes are Operating	7-26
Disbanding a Cluster If Only One Node is Operating	7-28
Changing the Multicast Base Address	7-29

Appendix A: Specifications for the File Director

File Director Hardware Specifications	A-1
File Director Compatibility	A-1
Compatibility with NFS Clients	A-1
Compatibility with NFS Servers	A-2
Compatibility with CIFS Clients	A-2
Compatibility with CIFS Servers	A-2
TCP and UDP Port Specifications	A-2

Appendix B: Third-Party Software License and Copyright Information

Index

About This Guide

This *File Director Administrator Guide* provides overview information and step-by-step instructions for configuring and maintaining the File Director™ from NeoPath™ Networks.

For information on installing and initializing the File Director, see the *File Director Quick Start Guide*.

For reference information on using the Command Line Interface (CLI) commands for configuring and maintaining the File Director, see the *File Director CLI Reference Guide*.

All of these guides are included with the File Director as Adobe Acrobat files (.pdf). (To install a free version of Adobe Acrobat Reader, see <http://www.adobe.com>.)

Audience

This guide is for Information Technology (IT) administrators who have experience installing and configuring network servers and equipment.

Contacting NeoPath Networks Technical Support

If you have any questions or you need technical assistance with the File Director hardware or software, call the NeoPath Networks Technical Support Department at 408-970-0300. You can also visit the Service & Support area of the NeoPath Networks Web site at www.neopathnetworks.com.

Before you call NeoPath Networks, make sure you know the Service Tag and the Express Service Code that uniquely identify your File Director hardware. Both of these identifiers are located on top in the front right corner of the File Director chassis. The Service Tag is also visible on a vertical bar code label above the purple mouse port on the rear panel of the File Director.

How to Use This Guide

Here is a summary of the chapters in this guide:

This chapter or appendix	Describes
Chapter 1, “Introducing the File Director”	An overview of the File Director
Chapter 2, “Configuring the File Director for Your Environment”	How to configure the File Director for your network, and share or export data from file servers to clients
Chapter 3, “Working with Synthetic Directories”	How to set up synthetic directories on the File Director

This chapter or appendix	Describes
Chapter 4, “Using the File Director to Migrate Data”	How to use the File Director to migrate data between file servers
Chapter 5, “Using Policies for Automatic File Migration”	How to create and use a policy to automate file migration
Chapter 6, “Monitoring and Troubleshooting the File Director”	How to use tools to monitor and troubleshoot the File Director
Chapter 7, “Setting Up a High Availability Cluster of File Directors”	How to configure a cluster of File Directors
Appendix A, “Specifications for the File Director”	Hardware and compatibility specifications for the File Director
Appendix B, “Third-Party Software License and Copyright Information”	Third-party software license and copyright information
“Index”	The Index for this guide

Conventions

The following conventions are used in this guide:

- On-screen text appears in bold font. For example:
Click **Add New User** under the **List of Users** section.
- Information that you type appears in a bold, monospace font. For example:
admin
- A **Tip** suggests ways to make a task easier or faster. For example:
Tip: This example uses four virtual shared directory names, but you can also choose to set up a *synthetic directory*.
- A **Note** or **Important** contains important information. For example:
Note: The File Director supports CIFS and NFS (v2 and v3) NAS protocols.
Important: You must also configure the NFS file server to allow access from this File Director IP address.
- A **Warning** describes actions that can cause data loss or problems. For example:
Warning: Be careful when changing the IP configuration settings. If you enter incorrect settings, the File Director may become inaccessible from the network.

Getting Online Help

In addition to this guide, you can get context-sensitive online Help for each set of configuration settings in a section of the Management Console.

To get online Help:

- Click the Help link in the upper right corner of each configuration section.

The screenshot displays the File Director Management Console interface. On the left is a sidebar with a list of configuration categories, each with a dropdown arrow. The main content area is titled "HOST AND DNS CONFIGURATION" and shows the user "admin" is logged in. It contains three main configuration sections, each with a "HELP" link in its top right corner:

- SETTING HOST NAME:** Includes a "Host Name" field with the value "server1" and an "Edit" button.
- DNS SERVER LIST:** Includes a table with one row showing the IP address "172.22.1.10" and a "Delete" button. Below the table is an "Add" button.
- DNS SEARCH DOMAIN:** Includes a "SEARCH DOMAIN" field with the value "neopathnetworks.com" and "Assign" and "Remove" buttons.

Annotations with arrows point to the "HELP" links in the top right of each of these three configuration sections.

CHAPTER 1

Introducing the File Director

This chapter contains the following topics that introduce you to the File Director:

- Understanding the Problems of Network Attached Storage Servers
- What is the File Director?
- File Director Functionality and Benefits
- Examples of Scenarios for Deploying the File Director
- What are the Tasks for Configuring and Using the File Director?

Understanding the Problems of Network Attached Storage Servers

Network Attached Storage (NAS) file servers provide storage for large quantities of data on a network and are critical for the business activity of many enterprises. But NAS file servers have some disadvantages and problems that you, as an IT administrator, must solve:

- **Data storage requirements are growing but budgets are constrained.**

There is a large growth in the amount of data that many enterprises must store, but you have an increasingly constrained budget. You are forced to use existing file servers more efficiently, or else purchase less expensive file servers to minimize expenses.

- **To increase data storage capacity, downtime is required.**

To increase data storage, either by increasing the capacity of existing file servers or by adding new file servers, you must disable client access. When you add a new export, shared directory, or file server, you must also reconfigure clients to make use of it.

- **Migration of data between file servers disrupts client access.**

While migrating data between file servers, you must disable client access. The migration must be scheduled carefully to minimize interruption to company productivity. Furthermore, migrating data between file servers is conceptually simple, but the process has several time-consuming steps:

- Disable client access to the data on the source file server.
- Copy the data to the destination file server.
- Perform a verification that the copy succeeded.
- Reconfigure clients with the new paths to the migrated data.
- Enable client access to the migrated data.
- Remove original data on the source file server.

- **Capacity and throughput utilization is not balanced across file servers.**

The differences in demands for data can result in some file servers being over used while other file servers are idle. File servers can be better utilized if frequently accessed data is spread across file servers, but you must disrupt client access to migrate data.

- **Managing multiple file servers requires significant resources.**

You must manage file servers individually because they are typically standalone “silos” (isolated islands of storage), which create an increasingly significant workload as the number of file servers increases. To control management costs, you must consolidate file servers.

- **Data is stored on file servers without consideration for its characteristics.**

Application designers and client users often haphazardly determine which data to store on the various file servers. A better approach is to store very important data on the most reliable servers, and store frequently accessed data on high performance servers. Lower priority data can then be stored on lower cost servers.

- **As the number of file servers increases, client users can have trouble finding data.**

With standalone file servers, you cannot combine and present the data distributed on multiple file servers into a unified view for client users. Client users must navigate to multiple shares or mount multiple directories to access data that is stored on multiple file servers.

What is the File Director?

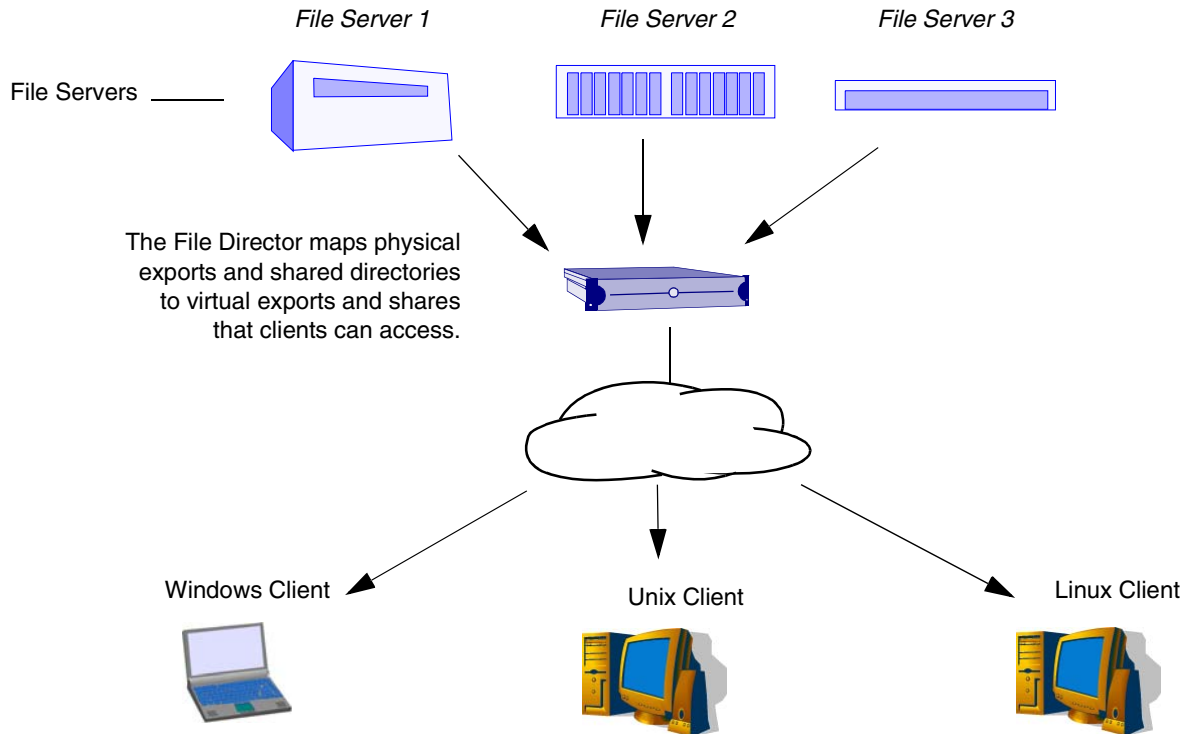
The File Director solves the problems of standalone file server silos by aggregating file servers into a unified view of NAS resources and decoupling clients from the servers.

The File Director is deployed between clients and file servers and transparently redirects communications between them on an IP network. No installation of client or server software is required. The File Director uses the native file access capabilities built into the operating systems of servers and clients. The File Director stores its own configuration data but it does not store any enterprise data. Enterprise data is stored on the file servers.

The File Director introduces a virtual layer between clients and file servers. Instead of connecting to physical shares and file servers directly, clients connect to virtual shares and virtual servers that the File Director presents.

When a client wants to perform operations, such as read and write on directories or files of a virtual share, the client sends requests to a virtual server on the File Director. The File Director modifies the requests and redirects them to the appropriate physical share or export where the actual data is located. Then it forwards the replies from the file server to the client.

As a result, the client is decoupled from the file server storing the data.



The File Director can provide clients with a single virtual export or share name to access data on multiple file servers regardless of the physical locations of the data.

With the File Director, you can migrate data between file servers at any time with no disruption to clients who may be accessing the data at the same time. (Clients may experience slower access due to migration traffic.) You can also easily change the mapping of exports and paths to their real storage locations.

The File Director provides additional features in the NAS system that are not available with file servers alone. For example, you can present client users with a single virtual directory (called a *synthetic directory*) that consists of multiple subdirectories physically located on multiple file servers. Clients can use the synthetic directory to access data stored on multiple file servers as if the data were stored on a single file server.

Note: The File Director supports CIFS and NFS (v2 and v3) network file access protocols for redirecting client to server communications on an IP network. For more information on supported protocols, servers and clients, see “File Director Compatibility” on page A-1.

File Director Functionality and Benefits

With the File Director, you can:

- **Increase headroom capacity by consolidating multiple file servers and eliminating data silos.**

The File Director can logically combine multiple servers and present them as a single file server to clients. You can use the File Director to perform all of the time-consuming steps involved in migrating data to balance storage utilization across servers. The benefit is that the File Director allows leveraging of current investments in data storage servers.

- **Transparently consolidate and add file servers on the fly.**

The File Director allows you to add file servers without any downtime for client users.

- **Balance storage use across file servers.**

By using the File Director, you can migrate data stored on highly loaded file servers to less loaded servers, which allows higher utilization of existing servers and postpones the purchase of new storage. By using *policies* (pre-defined selection criteria), you can have the File Director automatically migrate data based on particular file attributes at the time you specify.

- **Centralize management of access to NAS shares and exports, and placement of data in those NAS shares and exports.**

The File Director includes both a web-based management console and a text-based command line interface for simplifying the management of access to NAS shares and exports and placement of data on those NAS shares and exports. This is especially useful for common day-to-day data management tasks across different servers.

- **Transparently expand capacity and add functionality to less expensive file servers.**

The File Director can reduce expenses by adding new functionality such as data migration to low-cost and legacy file servers.

- **Remove the dependency between managing clients and managing file servers.**

The client views the File Director as a virtual server for data. You can use the File Director to change the storage locations of data without affecting clients.

- **Provide high availability of file server access by setting up a cluster of File Directors.**

In a cluster, a standby File Director automatically takes over if the active File Director in the cluster fails for any reason.

Examples of Scenarios for Deploying the File Director

There are many possible scenarios for deploying the File Director in various environments. This section contains some examples of scenarios that illustrate typical usage.

If your objective is to:	See this scenario example:
Balance storage capacity among legacy file servers via file migration without disrupting client access and without requiring client reconfiguration	“Scenario 1: Balance Storage Capacity Among Legacy File Servers” on page 1-5
Transparently add new storage in a legacy environment, and optionally present clients with a single hierarchy for accessing multiple shared directories on multiple file servers	“Scenario 2: Transparently Add Storage and Present a Single Hierarchy” on page 1-9
Deploy the File Director and all new file servers and clients into a new environment at the same time	“Scenario 3: Deploy the File Director and New File Servers and Clients into a New Environment” on page 1-13

Note: The following examples use CIFS terminology and syntax, but the same concepts and File Director functionality apply to NFS file servers. Only the NFS configuration details are different.

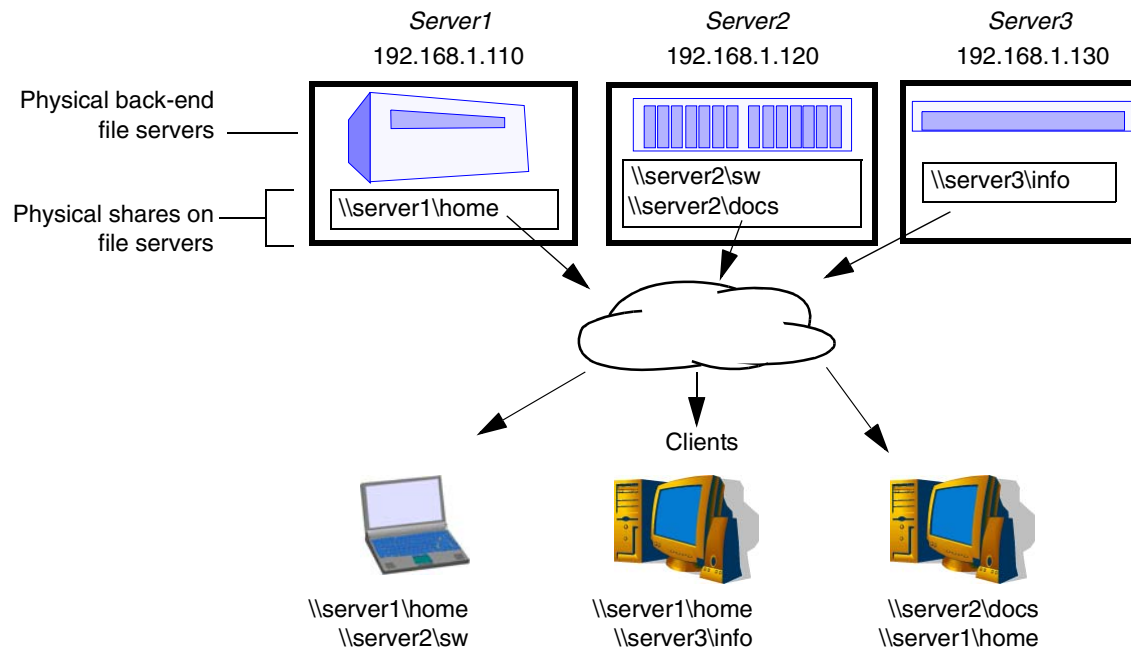
Scenario 1: Balance Storage Capacity Among Legacy File Servers

This scenario illustrates how to balance data storage among legacy file servers without disrupting client access and without requiring client reconfiguration by using the File Director to migrate the data.

In this scenario, the File Director is deployed into an existing environment where file servers and clients are already configured and being used. You can avoid changing the existing client configurations and change the smaller number of file servers instead. This is accomplished by reassigning the file servers' existing IP addresses to the File Director virtual servers, and assigning new system host names and IP addresses to the file servers. The File Director is configured with virtual server names and virtual share names that are identical to the existing file server names and physical share names. All client requests to the previous server names and shares are then directed to the File Director, which maps the requests to the file servers.

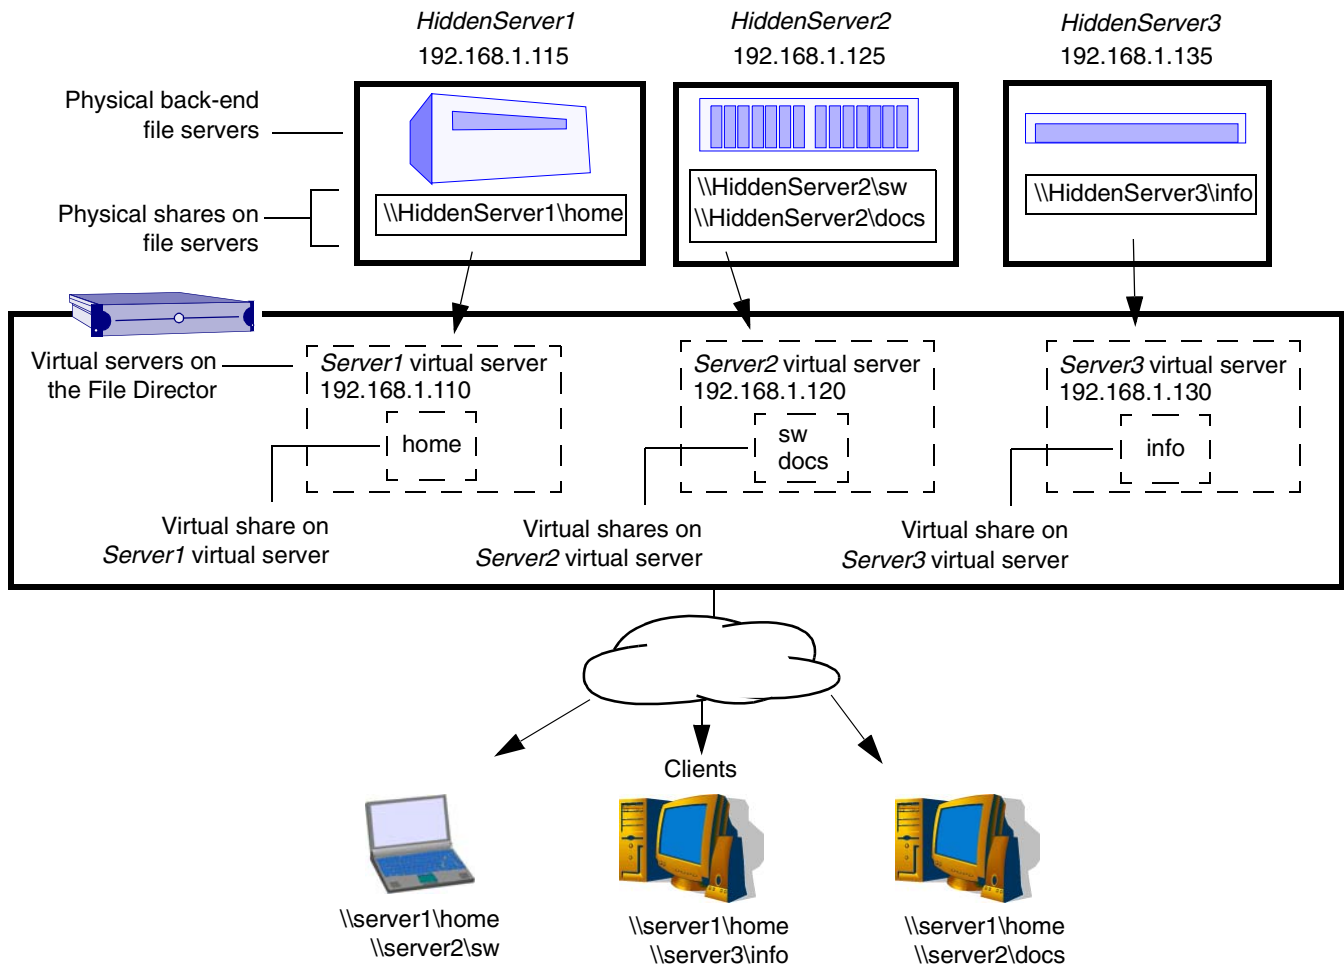
After the deployment and reconfiguration is complete, you can balance data storage among the file servers by using the File Director for transparent online migration of the data. Client access to the migrated data is not affected because clients are accessing the same server and share names through the File Director, which then redirects the requests to the new location.

Here is an example of a legacy environment before the File Director is deployed:



Example of a legacy environment before deploying the File Director

Here is an example of a legacy environment after the names and IP addresses of the file servers have been changed, and the File Director is configured with the old IP addresses, server names, and share names:



Example of a legacy environment after deploying the File Director

SUMMARY OF TASKS FOR THIS DEPLOYMENT SCENARIO

Here is a summary of the steps for this deployment scenario:

- Change IP addresses and names of existing file servers. (Update the DNS server if the file servers are used for other purposes, such as application servers.)
- Configure the File Director with virtual server names that are exactly the same as the old file server names, and assign the old file server IP addresses to the virtual servers.
- For each virtual server name, configure the File Director with virtual share names that are exactly the same as the physical share names on the file server.
- Map the virtual server and share names to the physical shares of file servers.

CONFIGURATION OF FILE SERVERS FOR THIS SCENARIO

Each of the file servers must be reconfigured as shown in the following table:

File Server Configuration

Old File Server Name	Old IP Address	New File Server Name	New IP Address	Physical Shared Directories
Server1	192.168.1.110	HiddenServer1	192.168.1.115	home
Server2	192.168.1.120	HiddenServer2	192.168.1.125	sw docs
Server3	192.168.1.130	HiddenServer3	192.168.1.135	info

CONFIGURATION OF THE FILE DIRECTOR FOR THIS SCENARIO

The File Director is configured with three virtual server names for clients to access instead of the file servers. Each virtual server is configured to provide virtual shares that have the same names as, and are mapped to, the physical shared directories on the file servers.

The following table summarizes the configuration of the virtual server names and virtual shared directories on the File Director:

File Director Configuration

Virtual Server Name	IP Address	Virtual Shared Directories	Physical Shared Directories
Server1	192.168.1.110	home	\\HiddenServer1\home
Server2	192.168.1.120	sw docs	\\HiddenServer2\sw \\HiddenServer2\docs
Server3	192.168.1.130	info	\\HiddenServer3\info

CONFIGURATION OF CLIENTS FOR THIS SCENARIO

No client reconfiguration is required.

Scenario 2: Transparently Add Storage and Present a Single Hierarchy

This scenario illustrates how to accomplish two independent objectives:

- Transparently add storage to a legacy environment without disrupting client access and without requiring client reconfiguration. You can implement this objective with or without implementing the next objective.
- Present clients with a single hierarchy for accessing multiple shared directories on multiple file servers.

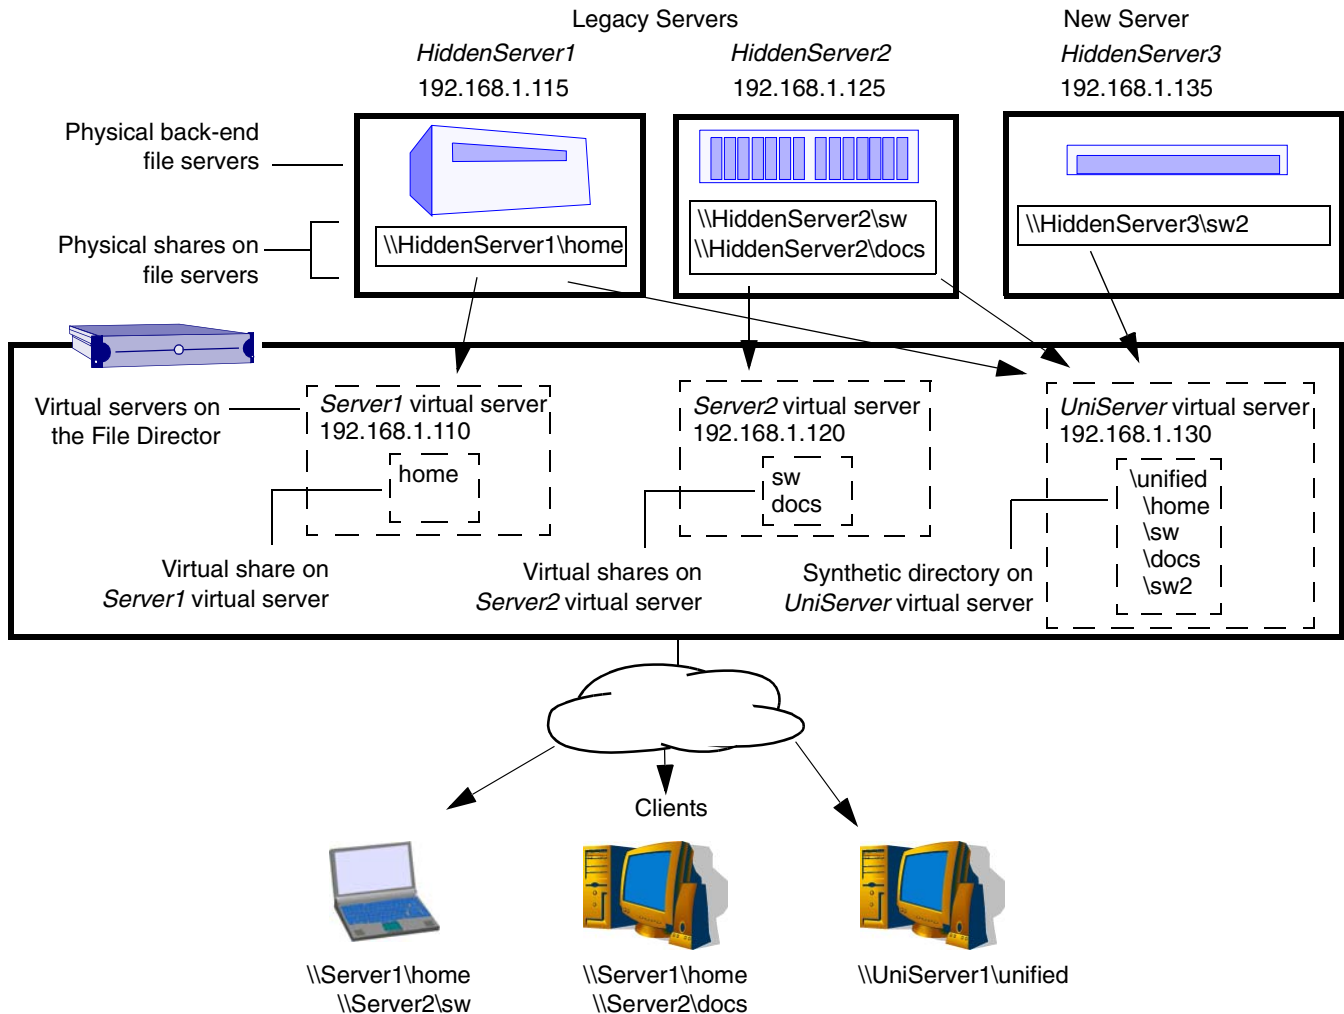
The environment for this scenario is similar to the environment in the previous legacy scenario, where you assign the file servers' existing IP addresses to the File Director virtual servers, and new system host names and IP addresses to the file servers.

In this scenario:

- You transparently add a new (and potentially inexpensive) file server to increase storage, but the new server's shares are mapped to virtual servers that have the same names as the legacy file servers. No client reconfiguration is required to access the new storage.
- You can transparently migrate files to the new file server without affecting clients.
- Optionally, you can also create a synthetic directory that presents clients with a single directory hierarchy that consists of all shares on all file servers.

In the following example, the new file server is called *HiddenServer3* and its share, *sw*, is mapped on the File Director to the *Server2* virtual server.

A synthetic directory called *\unified* on the *UniServer* virtual server presents clients with a hierarchy that contains the other shares on the *Server1* and *Server2* virtual servers. Synthetic directories allow clients access to all shares by using a single namespace for a virtual server and share.



Example of transparently adding storage to a legacy environment and using a File Director synthetic directory

SUMMARY OF TASKS FOR THIS DEPLOYMENT SCENARIO

Here is a summary of the steps for this deployment scenario:

- Change IP addresses and names of existing file servers. (Update the DNS server if the file servers are used for other purposes, such as application servers.)
- Configure the File Director with virtual server names that are exactly the same as the old file server names, and assign the old file server IP addresses to the virtual servers.
- Map the virtual server and share names to the physical shares of file servers.
- For each virtual server name, configure the File Director with virtual share names that are exactly the same as the physical share names on the file server.
- Add a new file server and map its shares to the virtual server and share names for the legacy servers.
- Create a synthetic directory that contains the physical shares of file servers under a single directory hierarchy.

CONFIGURATION OF FILE SERVERS FOR THIS SCENARIO

Each of the file servers must be reconfigured as shown in the following table:

File Server Configuration

Old File Server Name	Old IP Address	New File Server Name	New IP Address	Physical Shared Directories
Server1	192.168.1.110	HiddenServer1	192.168.1.115	home
Server2	192.168.1.120	HiddenServer2	192.168.1.125	sw docs
<i>none</i> (new file server)	<i>none</i> (new file server)	HiddenServer3	192.168.1.135	sw2

CONFIGURATION OF THE FILE DIRECTOR FOR THIS SCENARIO

The File Director is configured with three virtual server names for clients to access instead of the file servers. Each virtual server is configured to provide virtual shares that have the same names as, and are mapped to, the physical shared directories on the file servers. A new virtual server called *UniServer* is configured to provide a synthetic directory that contains a hierarchy of the shares on the other virtual servers.

The following table summarizes the configuration of the virtual server names and virtual shared directories on the File Director:

File Director Configuration

Virtual Server Name	IP Address	Virtual Shared Directories	Physical Shared Directories
Server1	192.168.1.110	home	\\HiddenServer1\home
Server2	192.168.1.120	sw docs	\\HiddenServer2\sw \\HiddenServer2\docs

The following table summarizes the configuration of the synthetic directory and synthetic links on the File Director:

File Director Configuration

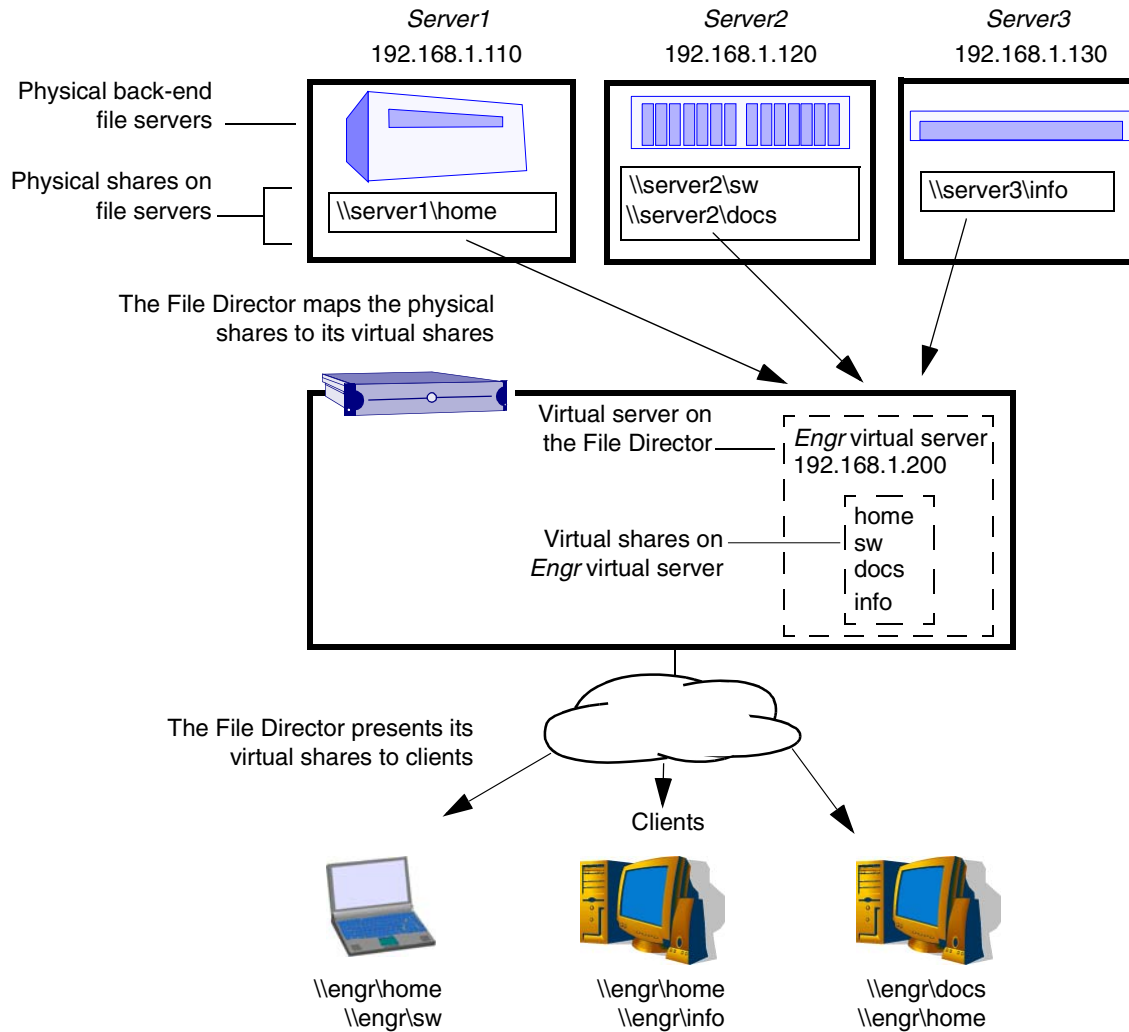
Virtual Server Name	IP Address	Synthetic directory	Use these synthetic links on the File Director	To map to these physical servers and exports
UniServer	192.168.1.130	\unified	\unified\home \unified\sw \unified\docs \unified\sw2	\\HiddenServer1\home \\HiddenServer2\sw \\HiddenServer2\docs \\HiddenServer3\sw2

CONFIGURATION OF CLIENTS FOR THIS SCENARIO

No client reconfiguration is required to access the File Director virtual shares because they have the same names and IP addresses as the old file servers. Clients must be configured to access the new virtual server with the synthetic directory.

Scenario 3: Deploy the File Director and New File Servers and Clients into a New Environment

This scenario illustrates how to deploy and configure the File Director and all new file servers and clients at the same time.



Example of a File Director deployment in a new environment with new clients and file servers

SUMMARY OF TASKS FOR THIS DEPLOYMENT SCENARIO

Here is a summary of the steps for this deployment scenario:

- Install new file servers and assign IP addresses to each server.
- Install the File Director, assign one or more IP Addresses to the File Director, and configure the DNS server with the File Director IP address (or addresses) and host system name.
- Configure the File Director with a virtual server name and map virtual share name(s) to physical shares of file servers.
- Configure the DNS server with virtual share host name and the File Director IP address.
- Configure all clients to use the File Director virtual share names and virtual shared directories. No clients are configured to access the file servers directly.

CONFIGURATION OF FILE SERVERS FOR THIS SCENARIO

Each of the file servers is configured to provide the shared directories shown in the following table:

File Server Configuration

File Server Name	IP Address	Physical Shared Directories
Server1	192.168.1.110	home
Server2	192.168.1.120	sw docs
Server3	192.168.1.130	info

CONFIGURATION OF THE FILE DIRECTOR FOR THIS SCENARIO

The File Director is configured with a virtual server name called *Engr* for clients to access instead of the file servers. The virtual server is configured to provide four virtual share names, which are mapped to the four physical shared directories on the file servers.

Note: You can also create multiple virtual servers on the File Director and map any number of virtual shares within each virtual server. For an example of multiple virtual servers, see “Scenario 1: Balance Storage Capacity Among Legacy File Servers” on page 1-5.

The following table summarizes the configuration of the virtual server name and virtual shared directories on the File Director:

File Director Configuration

Virtual Server Name	IP Address	Virtual Shared Directories	Physical Shared Directories
Engr	192.168.1.200	home	\\Server1\home
		sw	\\Server2\sw
		docs	\\Server2\docs
		info	\\Server3\info

Note: In this example, the virtual share names are exactly the same as the physical share names, but that is not a requirement. For example, a physical share name *hr_info* can be mapped on the File Director to *info*.

Tip: This example uses four virtual shared directory names, but you can also choose to set up a *synthetic directory*, which presents clients with a single directory hierarchy that consists of the union of multiple shares or exports on multiple file servers. For example, a synthetic directory called *\\MyRemoteFiles* could be mapped to the four physical directories in this example and presented as one unified hierarchy to clients. For more information on creating a synthetic directory, see Chapter 3, “Working with Synthetic Directories.”

CONFIGURATION OF CLIENTS FOR THIS SCENARIO

Configure the clients to access the virtual share names on the File Director by using these paths:

- \\engr\home
- \\engr\sw
- \\engr\docs
- \\engr\info

For example, to access the physical share called *\\server1\home*, a client would access the virtual share called *\\engr\home*, which is mapped in the File Director to *\\server1\home*.

What are the Tasks for Configuring and Using the File Director?

After initializing the File Director (see the *File Director Quick Start Guide*), you're ready to use the File Director Management Console or the Command Line Interface to finish configuring the File Director for your environment. (See Chapter 2, "Configuring the File Director for Your Environment.")

Some configuration tasks are required for the File Director to operate properly. Other tasks are ongoing operational tasks for using the File Director.

Required Configuration Tasks

Here are the required tasks for configuring the File Director:

Required Configuration Tasks	For more information, see
<input type="checkbox"/> If your clients and servers are on different subnets, or you connected the <i>eth1</i> through <i>eth5</i> interface ports, assign IP addresses to the <i>eth1</i> through <i>eth5</i> ports.	"Configuring the File Director IP Addresses" on page 2-6
<input type="checkbox"/> Specify information about the file servers that have shares or exports that you want the File Director to present to clients.	"Setting Up the File Director to Access File Servers" on page 2-13
<input type="checkbox"/> Specify the names of the virtual servers and virtual shares or exports that you want the File Director to present to clients.	"Specifying Virtual Servers and Virtual Shares" on page 2-20
<input type="checkbox"/> If you want to provide clients with a single virtual directory that consists of multiple physical subdirectories, set up a synthetic directory.	Chapter 3, "Working with Synthetic Directories"
<input type="checkbox"/> If you want to provide high availability, set up a cluster of File Directors.	Chapter 7, "Setting Up a High Availability Cluster of File Directors"

Ongoing Operational Tasks

Here are the ongoing operational tasks for using the File Director:

Operational Task	For more information, see
<input type="checkbox"/> Migrate data between file servers	Chapter 4, "Using the File Director to Migrate Data"
<input type="checkbox"/> Use policies to configure automatic file migrations	Chapter 5, "Using Policies for Automatic File Migration"
<input type="checkbox"/> Monitor the status of the File Director	Chapter 6, "Monitoring and Troubleshooting the File Director"

CHAPTER 2

Configuring the File Director for Your Environment

This chapter contains the following topics that explain how to configure the File Director for your network, and share or export data from file servers to clients:

- Using the Management Console to Configure the File Director
- Using the Command Line Interface to Configure the File Director
- Adding Administrator Accounts
- Configuring the File Director for Your Network
- Changing the File Director Host Name
- Setting the File Director Date, Time, and Time Zone
- Specifying an NTP Server
- Setting Up the File Director to Access File Servers
- Displaying and Matching Shares and Exports for a File Server
- Specifying Virtual Servers and Virtual Shares

Using the Management Console to Configure the File Director

After installing the File Director on your network, you can use the File Director Management Console in a web browser to configure the File Director for your environment. You can run the Management Console on any computer that can access the File Director over the network.

Important: Internet Explorer 6 and Mozilla 1.0 to 1.6 are supported for use with the File Director Management Console.

Logging Into the File Director Management Console

To log into the File Director Management Console:

1. Enter one of the following URLs in a Web browser on a computer that can access the File Director over a network or the Internet:
 - Enter the File Director host name. For example:
`https://server_name.company.com`
 - Enter the File Director IP address:
`https://<IP address>`
where <IP address> is any IP address assigned to the File Director, such as the IP address specified during initialization.
2. When the security certificate alert appears, accept the certificate.

- When the File Director login dialog box appears, enter an administrator user name, such as the default name:

admin

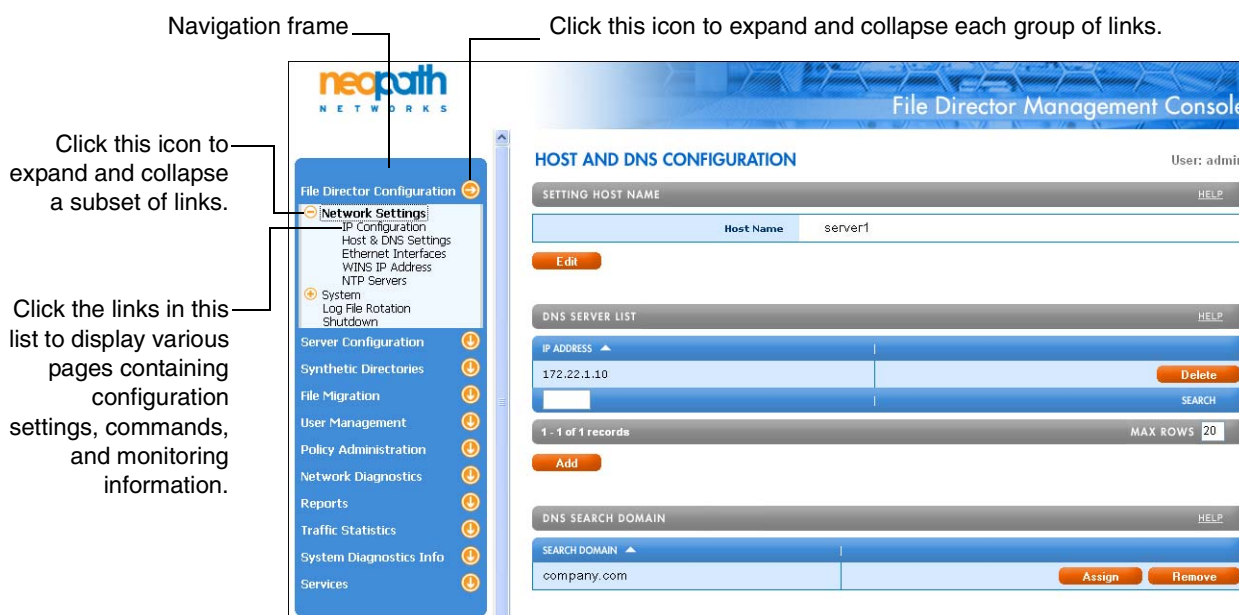
- Enter the administrator password.

Note: The user name and password are case sensitive. You can use the user name and password you specified during initialization, or you can create additional accounts later. (See “Adding Administrator Accounts,” next.)

- Click **OK**.

After you log in, the File Director Management Console appears.

The Management Console is composed of several pages where you select options, enter configuration information, and choose commands to configure and administer the File Director. Some pages contain status information and reports that you can use to monitor the File Director.



Note: The File Director automatically times out a session after 60 minutes of no activity.

Setting the Amount of Information Displayed in the Management Console

You can set the maximum number of rows to display in each section of the Management Console. You can also navigate to other pages of rows.

The screenshot shows the 'FILE SYSTEM MONITOR' interface. At the top, it says 'LIST OF FILE SYSTEMS' and 'Page: 1 2 3 4 5 > >> 17'. Below this is a table with columns: FILE SERVER, FILE SERVER SHARE OR EXPORT, PROTOCOL, MONITORED, RESYNC SCHEDULE, and a 'Display' button. The table contains five rows of data. At the bottom, there is a 'MAX ROWS' field set to 5 and a 'SEARCH' button. A callout points to the 'MAX ROWS' field with the text: 'To set the number of rows to display in each section, enter a number here, and then click MAX ROWS or press the Enter key.' Another callout points to the page navigation links with the text: 'Click here to navigate to other pages in this section.' A third callout points to the '1 - 5 of 85 records' text with the text: 'Current set of records'.

Current set of records

Click here to navigate to other pages in this section.

To set the number of rows to display in each section, enter a number here, and then click MAX ROWS or press the Enter key.

Notes:

- To go to the next page in a section, click on the single right arrow next to the page numbers.
- To go to the next set of 5 pages, click on the double right arrow.
- To go to the previous page, click on the single left arrow.
- To go to the previous set of 5 pages, click on the double left arrow.
- If the number of rows is less than the size of the page, the page numbers will not be displayed.

Sorting Information in the Management Console

You can sort the order of information displayed in the columns of a section in ascending or descending order. The default sort order is in ascending order of the first column.

The screenshot shows the 'FILE SYSTEM MONITOR' interface. At the top, it says 'LIST OF FILE SYSTEMS' and 'Page: 1 2 3 4 5 > >> 17'. Below this is a table with columns: FILE SERVER, FILE SERVER SHARE OR EXPORT, PROTOCOL, MONITORED, RESYNC SCHEDULE, and a 'Display' button. The table contains five rows of data. At the bottom, there is a 'MAX ROWS' field set to 5 and a 'SEARCH' button. A callout points to the 'FILE SERVER' column header with the text: 'An up arrow indicates this column is sorted in ascending order. A down arrow indicates descending order.' Another callout points to the 'FILE SERVER' column header with the text: 'To sort a column in ascending order, click the column label. To sort in descending order, click the label again.'

An up arrow indicates this column is sorted in ascending order. A down arrow indicates descending order.

To sort a column in ascending order, click the column label. To sort in descending order, click the label again.

Filtering Information in the Management Console

You can filter the information displayed in a section. For example, you can display only rows that contain CIFS in the Protocol column.

FILE SYSTEM MONITOR

LIST OF FILE SYSTEMS Page: 1 2 3 4 5 17 HELP

FILE SERVER	FILE SERVER SHARE OR EXPORT	PROTOCOL	MONITORED	RESYNC SCHEDULE	
server1	mktg	CIFS	No		Display
server2	sales	CIFS	No		Display
server2	eng	CIFS	No		Display
server3	hr	CIFS	No		Display
server4	finance	CIFS	No		Display

1 - 5 of 85 records Page: 1 2 3 4 5 17 MAX ROWS 5

To apply the filter text, click Search or press the Enter key.

Enter filter text in the field at the bottom of a column.

Adding Administrator Accounts

During initialization, you used the default administrator account for the File Director and specified a password for it. You can specify additional administrator accounts that have the same permissions as the default administrator account.

To add an administrator account for the File Director:

1. In the Navigation frame on the left side of the Management Console, expand **User Management** and then click the **User Management** link.

The **User Management** page appears.

USER MANAGEMENT User: admin

LIST OF USERS HELP

USER NAME	USER DESCRIPTION	USER STATUS	
admin	File Director setup and administration account	enabled	Delete Display

1 - 1 of 1 records MAX ROWS 20

Add new User

2. Click **Add New User** under the **List of Users** section.
3. Enter a user name for logging in the administrator account.
4. Enter a description for the account.
5. Enter a password for the administrator account.

6. From the **User Status** drop-down menu, select a status (**Enabled** or **Disabled**) for the administrator account.
7. Click **Submit** to commit the new settings.

Using the Command Line Interface to Configure the File Director

You can also use the Command Line Interface (CLI) commands to configure and maintain the File Director. All of the functionality in the File Director Management Console is available through CLI commands.

For more information on using the CLI commands, see the *File Director CLI Reference Guide*.

Configuring the File Director for Your Network

If your data storage requirements are medium to large and you require maximum bandwidth and throughput, it is recommended that you connect all of the File Director interface ports to your network and assign at least one IP address to each port. This recommended installation, which is shown in the *File Director Quick Start Guide*, dedicates separate port(s) for traffic to and from clients, and for traffic to and from file servers. During initialization, you specified an IP address and it was automatically assigned to the *eth0* interface port on the motherboard. If you connected the other ports, you must now add IP addresses for those ports.

Important: You must also add one IP address for each virtual server you want to use on the File Director. Virtual servers cannot share IP addresses.

Note: If your data storage requirements are small, you can use a single interface port and IP address for configuration, clients, and servers.

You can configure the following network settings for the File Director:

- IP Address and netmask for each File Director interface. You can specify one or more IP addresses or subnets for each interface. For example, you might need to offer service from the File Director to multiple subnets. You can also set the scope of an IP address to be local for a single File Director, or as part of a cluster of two File Directors. (See “Configuring the File Director IP Addresses,” next.)
- File Director network interface settings, such as speed or duplex mode. (See “Configuring the File Director Network Interfaces” on page 2-8.)
- Default gateway for the File Director to use. (See “Setting the Default Gateway for the File Director” on page 2-8.)
- Proxy IP address for the File Director to use. (See “Setting a Proxy IP Address for the File Director” on page 2-9.)
- DNS servers for the File Director to use. (See “Specifying a DNS Server and DNS Search Domain” on page 2-10.)
- A domain for the File Director to search for names of file servers. (See “Specifying a DNS Server and DNS Search Domain” on page 2-10.)
- WINS server for the File Director to use. (See “Specifying a WINS Server” on page 2-10.)

Warning: Be careful when changing the server's IP configuration settings. If you enter incorrect settings, the server may become inaccessible from the network. If that happens, you must connect a monitor and keyboard directly to the File Director, or use a serial terminal connection. For more information, see the *File Director Quick Start Guide*.

Configuring the File Director IP Addresses

To configure the File Director IP addresses:

1. In the Navigation frame, expand **File Director Configuration**, expand **Network Settings**, and then click the **IP Configuration** link.

The **Network Settings** page appears.

The screenshot shows the File Director web interface. On the left is a navigation pane with a tree structure. Under 'File Director Configuration', 'Network Settings' is expanded, and 'IP Configuration' is selected. The main area displays the 'NETWORK SETTINGS' page for user 'admin'. It features a section titled 'IP ADDRESS LIST' with a table containing two entries. Each entry has buttons for 'Delete' and 'Display'. Below the table is an 'Add' button. Other sections include 'CLUSTER' with 'PRIMARY NODE', 'BACKUP NODE', and 'MULTICAST IP ADDRESS' fields, and 'DEFAULT GATEWAY' with 'GATEWAY IP ADDRESS' and 'INTERFACE' fields.

IP ADDRESS	SUBNET MASK	INTERFACE	SCOPE	ACTIVE	
172.22.1.242	255.255.255.0	eth0	local	Yes	Delete Display
172.22.1.244	255.255.255.0	eth0	local	Yes	Delete Display

1 - 2 of 2 records MAX ROWS 20

Add

PRIMARY NODE	BACKUP NODE	MULTICAST IP ADDRESS	
(none assigned)	(none assigned)		Assign


GATEWAY IP ADDRESS	INTERFACE	
172.22.1.1	eth0	Assign

2. Click **Add** under the **IP Address List** section.

The **IP Address Configuration** page appears.

3. Enter an IP address.
4. Enter a subnet mask.

5. From the **Interface** drop-down menu, choose which File Director network interface port you want to assign the IP address to:
 - **eth0** or **eth1** (on motherboard)
 - **eth2** or **eth3** (on upper physical interface)
 - **eth4** or **eth5** (on lower physical interface)



The interface ports on the File Director

6. From the **Scope** drop-down menu, choose a scope:
 - Choose **Local** if you want the scope of the IP address to be limited to the File Director you are configuring.
 - If you are setting up a cluster of File Director nodes, choose **Cluster** to allow the IP address to be transferred to a Standby File Director node if the Active (current) File Director fails. For more information on clusters, see Chapter 7, “Setting Up a High Availability Cluster of File Directors.”
7. Click **Submit** to commit the new settings.

The new IP address is listed in the **IP Address List** section of the **Network Settings** page. An *Active* status means the IP address is currently assigned to an interface on the File Director.

To delete an IP address:

- In the **Network Settings** page, click **Delete** next to the IP address you want to delete.

Important: You cannot delete an IP address that is being used by a virtual server or proxy IP address until you first delete that virtual server or proxy IP address.

Configuring the File Director Network Interfaces

You can change the default settings for the File Director network interface settings, such as speed, duplex mode, and auto-negotiation.

To configure the File Director network interfaces:

1. In the Navigation frame, expand **File Director Configuration**, expand **Network Settings**, and then click the **Ethernet Interfaces** link.

The **Ethernet Interface Configuration** page appears.

ETHERNET INTERFACE CONFIGURATION					
LIST OF ETHERNET INTERFACES					
INTERFACE NAME	INTERFACE SPEED	INTERFACE MODE	AUTONEGOTIATE	LINK STATUS	
eth0	1000 Mbps	Full Duplex	yes	connected	Display
eth1	1000 Mbps	Full Duplex	yes	connected	Display
eth2	1000 Mbps	Full Duplex	yes	connected	Display
eth3	1000 Mbps	Full Duplex	yes	connected	Display

2. Locate the interface you want to configure and click **Display**.

The **Ethernet Interface** page appears.

3. Click **Edit**.
4. Choose a speed from the **Interface Speed** drop-down menu.
5. Choose a duplex mode from the **Interface Mode** drop-down menu.
6. Choose whether you want auto-negotiation enabled from the **Autonegotiate** drop-down menu.

Note: The **Auto-negotiation** setting is incompatible with the speed and duplex mode parameters. If you set **Auto-negotiation** to **Yes**, you cannot set the speed or duplex mode parameters. Conversely, if you set **Auto-negotiation** to **No**, you *must* specify the speed and duplex mode parameters.

7. Click **Submit**.

Setting the Default Gateway for the File Director

During initialization, you specified a default gateway for the File Director. If necessary, you can change the default gateway you specified.

To set the default gateway for the File Director:

1. In the Navigation frame, expand **File Director Configuration**, expand **Network Settings**, and then click the **IP Configuration** link.

The **Network Settings** page appears.

2. Click **Assign** in the **Default Gateway** section.

The **Default Gateway Assignment** page appears.

3. For **Gateway IP Address**, enter the IP address for the default gateway.

4. Choose one of the following from the **Interface** drop-down menu:
 - Choose **[Auto-select]** to have the File Director automatically select an interface port based on the default gateway IP address you specified.
 - Choose the interface port you want to associate with the default gateway IP address.
5. Click **Submit**.

After you click **Submit**, the File Director displays the port under **Interface** in the **Default Gateway** section.

Note: If you change the default gateway and the computer you are using to access the Management Console is not routable through the gateway specified, you may lose your network connection to the File Director.

Setting a Proxy IP Address for the File Director

The proxy IP address is used to allow CIFS clients to authenticate to the file server. The proxy address must be an IP address already assigned to one of the interfaces on the File Director.

You can use the File Director Proxy IP address as the default proxy IP address for all CIFS file servers. If necessary, you can also override the File Director Proxy IP address and specify a particular proxy IP address for each CIFS file server. For information, see “Specifying a File Server” on page 2-16. The File Director Proxy IP address is optional, but if you don’t specify it you must specify a proxy IP address for each CIFS file server.

In most cases, using the File Director Proxy IP address as the default proxy IP address for all CIFS file servers is sufficient.

Note: The proxy IP address cannot be used as a virtual server IP address.

Important: Before specifying a CIFS file server, you must specify a default File Director proxy IP address or a proxy IP address for the CIFS server.

Note: A single proxy IP address cannot be used for a CIFS file server that does not support the NetBIOS protocol (TCP port 139). Instead, configure a separate proxy IP address that is used exclusively for this server. For more information, see “Specifying a File Server” on page 2-16.

To set a proxy IP address for the File Director:

1. In the Navigation frame, expand **File Director Configuration**, expand **Network Settings**, and then click the **IP Configuration** link.

The **Network Settings** page appears.
2. Click **Assign** in the **Proxy IP Address** section.

The **Proxy IP Address Assignment** page appears.
3. From the **IP Address** drop-down menu, choose an address to use as the proxy IP address.

The IP addresses listed on the menu are already assigned to one of the interfaces on the File Director.

Important: If you want the File Director to be part of a cluster, you must first set the scope to **Cluster** for the IP address you want to use as the proxy IP address. For more information, see “Configuring the File Director IP Addresses” on page 2-6.

4. Click **Submit**.

Specifying a DNS Server and DNS Search Domain

During initialization, you specified a DNS server for the File Director to use. If necessary, you can change the DNS server you specified, or add other DNS servers.

You can also specify a domain for the File Director to search for names of file servers, such as *department.company.com*.

To specify a DNS server for the File Director:

1. In the Navigation frame, expand **File Director Configuration**, expand **Network Settings**, and then click the **Host and DNS Settings** link.

The **Host and DNS Configuration** page appears.

2. Click **Add** under the **DNS Server List** section.

The **DNS Configuration** page appears.

3. Enter an IP address for the DNS Server and then click **Submit**.

To specify a domain to search:

1. In the Navigation frame, expand **File Director Configuration**, expand **Network Settings**, and then click the **Host and DNS Settings** link.

The **Host and DNS Configuration** page appears.

2. Click **Assign** in the **DNS Search Domain** section.

The **DNS Search Domain** page appears.

3. Enter a domain name and then click **Submit**.

Specifying a WINS Server

You can specify a WINS server for the File Director to use for WINS name lookups and to register the names of CIFS virtual servers.

To specify a WINS server for the File Director:

1. In the Navigation frame, expand **File Director Configuration**, expand **Network Settings**, and then click the **WINS IP Address** link.

The **WINS IP Address Configuration** page appears.

2. Click **Assign**.

3. Enter an IP address for the WINS Server and then click **Submit**.

Changing the File Director Host Name

If necessary, you can change the File Director host name specified during initialization.

Important: If you want to change the host name of a File Director node in a cluster, you must first disband the cluster. After you change the host name, you can then recreate the cluster by specifying the roles of each node and the multicast base address. For more information, see “Disbanding a Cluster” on page 7-26.

To change the File Director host name:

1. In the Navigation frame, expand **File Director Configuration**, expand **Network Settings**, and then click the **Host and DNS Settings** link.

The **Host and DNS Configuration** page appears.

2. Click **Edit** under the **Setting Host Name** section.
3. Enter a new host name, and then click **Submit**.

Setting the File Director Date, Time, and Time Zone

If necessary, you can change the File Director date, time, and time zone.

To change the File Director date, time, and time zone:

1. In the Navigation frame, expand **File Director Configuration**, expand **System**, and then click the **Clock** link.

The **System Date and Time** page appears.

2. Click **Edit** under the **Setting Date and Time** section.
3. For **Date**, enter a new date in the following format: YYYY/MM/DD

4. For **Time**, enter a new time in the following format: HH:MM:SS
5. From the **Time Zone** drop-down menu, choose a time zone.
6. Click **Submit**.


Specifying an NTP Server

If necessary, you can specify an Network Time Protocol (NTP) server for the File Director to use.

To specify an NTP server:

1. In the Navigation frame, expand **File Director Configuration**, expand **Network Settings**, and then click the **NTP Servers** link.

The **NTP Servers** page appears.



2. Click **Add New NTP Server** under the **NTP Servers** section.
3. Enter the IP address or host name of the NTP server, and then click **Submit**.

Setting Up the File Director to Access File Servers

The Common Internet File System (CIFS) protocol and the Network File System (NFS) protocol enable clients on a network to access shared files stored on file servers.

CIFS, which is based on the commonly used Server Message Block (SMB) protocol, is integrated in a large variety of operating systems and file servers.

NFS uses Remote Procedure Call (RPC) primitives and an eXternal Data Representation (XDR) to provide functionality on different types of computers, operating systems, network architectures and transport protocols. By mounting NFS exports, clients can access remote directories located on file servers as part of their local file system.

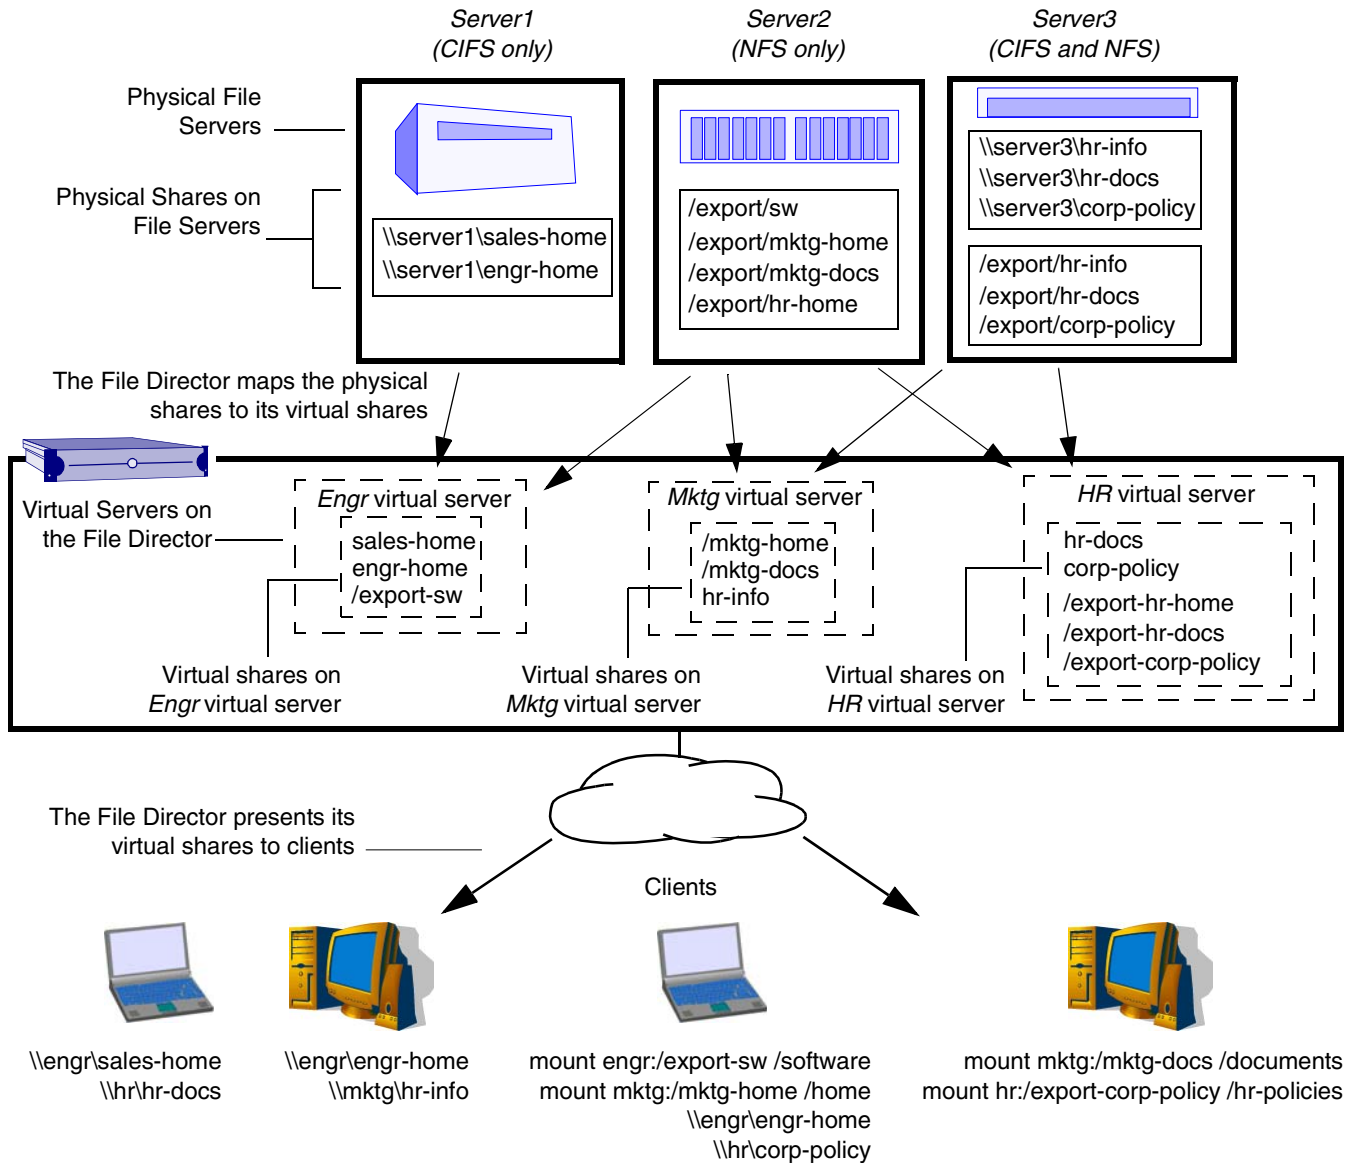
The File Director allows clients to share and access files on file servers that support CIFS, NFS or both protocols. Clients can also use either or both protocols with the File Director.

Overview of Accessing Shares and Exports with the File Director

When using the File Director to share and access files on file servers, clients do not access the file servers directly. Instead, clients access virtual servers and shares and exports on the File Director, which are mapped to the physical file servers and shares and exports you specify. You can also combine physical shares and exports from multiple file servers and present them as virtual shares and exports under a single virtual server, simplifying the access process for clients.

In the following illustration, for example, the File Director is configured to access three file servers. *Server1* supports CIFS only, *Server2* supports NFS only and *Server3* supports both CIFS and NFS. CIFS clients can access the *sales-home* virtual share on the *Engr* virtual server, and NFS clients can access the */export-sw* virtual export on the *Engr* virtual server. The *sales-home* virtual share is mapped to the *sales-home* physical share on the file server named *Server1*, and the */export-sw* virtual export is mapped to the */export/sw* physical export on *Server2*. The *Mktg* and the *HR* virtual servers also have virtual shares and exports mapped to *Server2* and *Server3* physical shares and exports. Note that the names of the virtual shares and

exports *do not* have to be the same as the physical shares and exports. The names in this example are the same for the sake of clarity in the illustration.



The File Director can provide clients with virtual shares and exports for accessing data on multiple file servers regardless of the physical locations of the data.

Here is a summary of the File Director configuration for the virtual servers and virtual shares and exports in the preceding illustration:

File Director Configuration

These virtual servers	And these virtual shares and exports	Are mapped to these physical servers and shares and exports
Engr	sales-home engr-home /export-sw	\\Server1\sales-home \\Server1\engr-home Server2:/export/sw
Mktg	/mktg-home /mktg-docs hr-info	Server2:/export/mktg-home Server2:/export/mktg-docs \\Server3\hr-info
HR	hr-docs corp-policy /export-hr-home /export-hr-docs /export-corp-policy	\\Server3\hr-docs \\Server3\corp-policy Server2:/export/hr-home Server3:/export/hr-docs Server3:/export/corp-policy

There are three basic steps to allow clients access to the file servers through the File Director:

1. Specify the file servers which have shares and/or exports that clients access. See the next section, “Specifying a File Server.”
2. Specify one or more virtual servers that you want to present to clients. See “Specifying Virtual Servers and Virtual Shares” on page 2-20.
3. Specify one or more virtual shares or exports that you want to map to the physical shares or exports on each file server. See “Specifying a Virtual Share or Virtual Export” on page 2-21.

Important: The File Director creates a *.neopath* directory in each physical share or export that you configure the File Director to access. This special directory contains File Director internal information that specifies the location of files that have been migrated away from the share or export, as well as files migrated to the share or export. The File Director prevents clients from accessing this directory, and administrators must not modify this directory when logged in directly to the file server. Any manual changes made to this directory may result in data loss. (*Note:* This does not preclude mapping of read-only physical shares. In this case, creation of a *.neopath* directory fails harmlessly and migration is not possible, but other functionality is available.)

Specifying a File Server

Specify each file server that has physical shares and/or exports you want the File Director to present as virtual shares or exports to clients.

To specify a file server:

1. In the Navigation frame, expand **Server Configuration** and then click the **File Servers** link.

The **File Server Configuration** page appears.

The screenshot shows the 'FILE SERVER CONFIGURATION' page. On the left is a navigation pane with 'File Servers' selected. The main area has a 'LIST OF FILE SERVERS' table with columns for 'SERVER NAME', 'PROTOCOL', and actions. It lists three servers: server1, server2, and server3. Each server has checkboxes for NFS and CIFS protocols and 'Delete' and 'Display' buttons. At the bottom, there is a search bar, a status bar showing '1 - 3 of 3 records' and 'MAX ROWS 20', and an 'Add new File Server' button.

SERVER NAME	PROTOCOL	
server1	<input checked="" type="checkbox"/> NFS <input checked="" type="checkbox"/> CIFS	Delete Display
server2	<input checked="" type="checkbox"/> NFS <input type="checkbox"/> CIFS	Delete Display
server3	<input type="checkbox"/> NFS <input checked="" type="checkbox"/> CIFS	Delete Display

2. Click **Add New File Server** under the **List of File Servers** section.

The **File Server Configuration** page appears.

The screenshot shows the 'FILE SERVER CONFIGURATION' page with the 'CONFIGURE A FILE SERVER' form. The form is divided into three steps. Step 1: Enter the server name, IP address and type of the file server. It includes fields for 'Server Name', '*Server IP Address' (with a hint), '*Server Type' (a dropdown menu showing '[None Selected]'), 'File Director IP Address for Server Access' (a dropdown menu showing '[Auto-select]'), and 'Protocol Types' (checkboxes for NFS and CIFS). Step 2: Enter the CIFS attributes of the file server. It includes fields for 'Admin Name', 'Admin Password', 'CIFS Domain Name', '*CIFS Proxy IP Address' (with a hint), and '*Supported CIFS Protocols' (checkboxes for NetBIOS and SMB Direct Host). Step 3: Submit the request or cancel and try again later. It includes 'Submit' and 'Cancel' buttons.

3. For **Server Name**, specify the host name of the file server.
4. (Optional) For **Server IP Address**, enter the IP address of the file server.

You can leave the IP Address blank and the File Director will attempt to automatically discover the IP address.

5. From the **Server Type** drop-down menu, choose the type of file server.

If you choose **[None-Selected]**, the File Director will attempt to automatically discover the correct type. Neopath recommends selecting a server type because not all server types are automatically discovered.

6. From the **File Director IP Address for Server Access** drop-down menu, choose a File Director IP address that you want the File Director to use when accessing this file server.

If you choose **[Auto-select]**, the File Director will automatically use one of the IP addresses that you have assigned to the File Director. (For information on how to determine which address it is using to access this file server, see the next section, “Displaying the File Server Access Configuration.”)

Important: If the file server supports NFS and is configured to restrict access to particular IP addresses, make sure that the IP address specified for **File Director IP Address for Server Access** is an IP address that is allowed to access the server. You must also configure the NFS file server to allow root access from that IP address. For example, on Linux, make sure each export has the following option specified in */etc/exports*:

```
<file_director_ip_address> (no_root_squash,no_subtree_check)
```

Or, for example, make sure each export has the following option specified in */etc/dfs/dfstab* on Solaris or */etc/exports* on Network Appliance:

```
root=<file_director_IPAddress>
```

Important: If you want the File Director to be part of a cluster, you must set the scope to **Cluster** for the IP address you are using for **File Director IP Address for Server Access**. For more information, see “Setting Local or Cluster Scope” on page 7-2.

7. For **Protocol Types**, select the protocols that the file server supports—NFS, CIFS, or both.
8. For a CIFS file server, do the following in the **Step 2** section of the **File Server Configuration** page:

- For **Admin Name** and **Admin Password**, enter the administrator name and password for the CIFS file server.

Important: Be sure to enter the administrator name and password correctly. The File Director uses the name and password to log into the CIFS file server whenever it needs to:

- Access information about migrated files on the CIFS server. The File Director stores migration information in a special directory called **.neopath** in the root of each share. If the File Director can’t log in, it can’t access this information and consequently it won’t allow clients to access *any* files on the server.
- Migrate files to or from the CIFS file server. If the File Director can’t log in, it won’t be able to copy the files and the migration will fail.
- Retrieve and display the list of CIFS shares on the CIFS file server. This occurs when you open the **File Server Configuration** page (see the next section, “Displaying and Matching Shares and Exports for a File Server.”), or when you use the **show share** CLI command (see the *File Director CLI Reference Guide*). If the File Director can’t log in, the list of CIFS shares for the file server won’t be listed.
- For **CIFS Domain Name**, enter the domain name of the CIFS file server.

- For **CIFS Proxy IP Address**, specify an IP address assigned to the File Director.

You don't have to specify a proxy IP address for this server if you already specified a proxy IP address for the File Director. For more information, see "Setting a Proxy IP Address for the File Director" on page 2-9.

Note: Client users will not use the Proxy IP address of the file server. The File Director uses this IP address internally only to meet certain CIFS protocol requirements.

- For **Supported CIFS Protocols**, select the CIFS protocols that the file server supports—**NetBIOS**, **SMB Direct Host**, or both. Most CIFS servers support both.

9. Click **Submit**.

If all of the settings are correct, the file server is shown in the **List of File Servers** section of the **File Server Configuration** page.

Displaying the File Server Access Configuration

To display the file server access configuration:

1. In the Navigation frame, expand **Server Configuration** and then click the **File Servers** link.

The **File Server Configuration** page appears.

2. Click **Display** next to the file server whose access configuration you want to display.

The **File Server Configuration** page appears.

Note: The **File Director IP Address for Server Access** field in the **NFS Configuration** section displays the File Director IP address that the File Director is using to access the selected file server.

Displaying and Matching Shares and Exports for a File Server

You can display the currently configured physical shares and exports for a file server.

Displaying Shares and Exports for a File Server

To display shares and exports for a file server:

1. In the Navigation frame, expand **Server Configuration** and then click the **File Servers** link.

The **File Server Configuration** page appears.

2. Click **Display** next to the file server whose shares or exports you want to display.

The **File Server Configuration** page appears.

3. Scroll down the **File Server Configuration** page to see the currently defined list of shares and exports for the file server you selected.

For example:

PHYSICAL SHARES & EXPORTS		
Page: 1 2 3 4		HELP
CIFS PHYSICAL SHARE	NFS PHYSICAL EXPORT	
common	/export/common	Migration Stats
	/export/src	Migration Stats
homes	/export/home	Migration Stats
	/export/scratch	Migration Stats
software	/export/software	Migration Stats
	/vol/vol1/film6	Migration Stats
SEARCH		
1 - 6 of 20 records		Page: 1 2 3 4
MAX ROWS 6		
Match Share to Export		

Tip: You can also display migration statistics for shares and exports. For more information, see “Displaying Migration Statistics for a Physical Share or Export” on page 4-12.

Matching or Unmatching Physical Shares and Exports

Important: On file servers that support both CIFS and NFS, you can export the same sub-tree of the file-system using both CIFS and NFS. You should match the CIFS share name with the NFS export name on the File Director to create an association between the two names. This is important if you want to migrate files in this sub-tree. If you don't match this migration source share name with the migration source export name, the File Director will migrate using either CIFS or NFS (depending on the migration specified), and clients won't be able to access the migrated files via the other protocol.

To match or unmatch physical shares and exports for a file server:

1. In the **File Server Configuration** page, click **Match Share to Export**.

The **Physical Share & Export Configuration** page appears.

<div>File Director Configuration</div> <div>Server Configuration</div> <div>File Servers</div> <div>Virtual Servers</div> <div>Synthetic Directories</div> <div>File Migration</div>	PHYSICAL SHARE & EXPORT CONFIGURATION User: admin	
	<div>MATCH SHARE TO EXPORT</div> <div>HELP</div>	
	<div>CIFS Share Name</div> <div></div>	<div>NFS Export Path</div> <div></div>
	<div>Match as Same File System</div> <div>Un-Match</div> <div>Cancel</div>	

2. From the **CIFS Share Name** drop-down menu, choose the CIFS physical share you want to match to, or unmatch from, an NFS physical export.

Note: You cannot match a share to more than one export.

3. From the **NFS Export Path** drop-down menu, choose the NFS physical export you want to match to, or unmatch from, the selected CIFS physical share.

Note: You cannot match an export to more than one share.

4. Do one of the following:
 - To match the selected share or export, click **Match as Same File System**.
 - To unmatch the selected share or export, click **Un-Match**.

If you matched the share or export, they appear on the same row in the table on the **File Server Configuration** page. If you unmatched them, they appear on separate rows.

Specifying Virtual Servers and Virtual Shares

After specifying a physical file server, you can specify one or more virtual servers that you want to present to clients.

Important: The File Director must have one IP address for each virtual server you want to use. Virtual servers cannot share IP addresses. Also, if the File Director is in a cluster, you must first set the scope to Cluster for each of the IP addresses you want to use for virtual servers. For more information on adding and configuring IP addresses, see “Configuring the File Director IP Addresses” on page 2-6.

Specifying a Virtual Server

To specify a virtual server:

1. In the Navigation frame, expand **Server Configuration** and then click the **Virtual Servers** link.

The **Virtual Server Configuration** page appears.

2. Click **Add New Virtual Server** under the **List of Virtual Servers** section.

The **Configure a Virtual Server** page appears.

3. For **Virtual Server Name**, enter the name of the virtual server.

For example, enter *Engr*. The virtual server name is the name that clients will use to access the virtual server.

Note: Add this virtual server name to your DNS server.

4. Choose the File Director IP address you want to use for the virtual server from the **Virtual IP Address** drop-down menu.

The IP addresses listed on the menu are configured on the **Network Settings** page. For more information, see “Configuring the File Director IP Addresses” on page 2-6.

5. For **File Services**, select the protocol(s) that you want to allow the virtual server to share or export to clients.

For example, you can only add a CIFS virtual share to a virtual server that supports CIFS.

6. (Optional) For **Comment**, enter a comment for the virtual server.

The comment can be any text string. The comment is displayed for the virtual server when CIFS clients browse the network. NFS clients do not view the comment. For both CIFS and NFS, the comment also appears on the **Virtual Server Configuration** page for the virtual server as information for administrators.

7. For **Default Physical Server Name**, enter the name of the default file server you want to use for the virtual server.

CIFS requests sent to this virtual server that are not directed to any particular virtual share will be directed to this default file server. In particular, requests sent to the special IPC\$ share will be directed to this server.

8. For **Supported CIFS Protocols**, select the CIFS protocols that you want the virtual server to support—**NetBIOS**, **SMB Direct Host**, or both.

Note: The file server you specify for **Default Physical Server Name** must support at least one of the protocols you select for **Supported CIFS Protocols**.

9. Click **Submit**.

If all of the settings are correct, the virtual server is shown in the **List of Virtual Servers** section of the **Virtual Server Configuration** page.

Specifying a Virtual Share or Virtual Export

After specifying a virtual server, you can specify one or more virtual shares or virtual exports that you want to map to the physical shares or exports on file servers.

Note: On each file server, you must configure the appropriate permissions for the physical shares or exports you want the File Director to present.

To specify a virtual share or virtual export:

1. In the Navigation frame, expand **Server Configuration** and then click the **Virtual Servers** link.

The **List of Virtual Servers** appears in the **Virtual Server Configuration** page.

- Click **Display** next to the virtual server whose virtual shares or exports you want to specify.

The currently defined list of virtual shares and exports appears for the virtual server you selected.

VIRTUAL SERVER CONFIGURATION User: admin

VIRTUAL SERVER HELP

Server Configuration:

Virtual Server Name	Engr
Virtual IP Address	172.22.1.244
Comment	
File Services	<input checked="" type="checkbox"/> NFS <input type="checkbox"/> CIFS

CIFS Configuration:

Default Physical Server Name	
Supported CIFS Protocols	<input type="checkbox"/> NetBIOS <input type="checkbox"/> SMB Direct Host

Edit Back

VIRTUAL SHARES & EXPORTS HELP

CIFS VIRTUAL SHARE NAME	NFS VIRTUAL EXPORT NAME
None.	

0 records MAX ROWS 20

Create CIFS Virtual Share Create NFS Virtual Export

- Do either of the following:
 - To create a CIFS virtual share, click **Create CIFS Virtual Share** under the **Virtual Shares & Exports** section. The **CIFS Virtual Share Configuration** page appears.
 - To create an NFS export, click **Create NFS Virtual Export**. The **NFS Virtual Export Configuration** page appears.

- Do one of the following:

- For **CIFS Virtual Share Name**, enter the name of the CIFS virtual share that you want clients to access.

A CIFS virtual share name can be any text string. Do not include a backslash. For example, enter *home*.

- For **NFS Virtual Export Name**, enter the name of the NFS virtual export that you want clients to access.

An NFS export name can be any absolute path name that includes a slash. For example, enter */home*.

Note: The name of the virtual share or export *do not* have to be the same as the physical shares and exports.

5. For **Export Type**, do one of the following:
 - Choose **Physical Share** or **Physical Export** to create a virtual share or export that is mapped to a physical share or export. Proceed to step 6.
 - Choose **Synthetic Directory** to create a virtual share or export to present a synthetic directory to clients. Proceed to step 7. For information on synthetic directories, see Chapter 3, “Working with Synthetic Directories.”
6. If you chose **Physical Share** or **Physical Export** as the export type, do the following steps:
 - From the **CIFS Physical Server** drop-down menu or **NFS Physical Server** drop-down menu, choose the name of the physical file server that contains the physical share or export you want to map to this virtual share or export.
 - For **CIFS Physical Share Name** or **NFS Physical Export Name**, enter the name of the physical share or export that you want to map to this virtual share or export.
7. If you chose **Synthetic Directory** as the export type, choose the name of the synthetic directory from the **Synthetic Directory** drop-down menu.

Note: You must create the synthetic directory before you can present it by using a virtual share or export. For more information, see “Creating a Synthetic Directory” on page 3-8.
8. Click **Submit**.

If all of the settings are correct, the virtual share is shown in the **Virtual Shares & Exports** section of the **Virtual Server Configuration** page for the currently selected virtual server.

You can specify additional virtual shares or exports for the currently selected virtual server by repeating this procedure.

Configuring ACL Entries for an NFS Virtual Export

You can configure optional Access Control List (ACL) entries for an NFS virtual export. An ACL entry allows you to specify which clients can access an NFS virtual export, and what sort of access they are granted.

Each ACL entry consists of a **Clients** setting that specifies which clients are affected by the entry, and one or more **Options** that specify which access options are in effect for these clients.

To configure an NFS virtual export ACL entry:

1. In the Navigation frame, expand **Server Configuration** and then click the **Virtual Servers** link.

The **List of Virtual Servers** appears in the **Virtual Server Configuration** page.
2. Click **Display** next to the virtual server that has the virtual export for which you want to specify an ACL entry.

The currently defined list of virtual shares and exports appears for the virtual server you selected.

- Click **Display** next to the virtual export for which you want to specify an ACL entry.
The **NFS Virtual Export Configuration** page appears.

- Click **Create ACL Entry**.
The **NFS Virtual Export ACL Entry Configuration** page appears.

- For **Clients**, enter information that specifies the clients.

You can enter a single host name, an IP address for a host, a net group, a host name with wildcards, or an IP network.

Here is a summary and examples of the types of client information you can enter:

Type of Client Information	Description	Examples of Client Information Values to Enter
Single host	An abbreviated host name	server123
	A fully qualified domain name	server123.company.com
	An IP address	172.22.1.4
Net group	Net group is of the form: @<netgroupname>	@engineering

Type of Client Information	Description	Examples of Client Information Values to Enter
Wildcard	A host name with the wildcard character * or ?, or both characters	<code>*.company.com</code> <code>192.168.1.*</code> <code>engr?.company.com</code> <code>engr?*.com</code>
IP network	IP network is of the form: <code><ip-address>/<netmask></code> where the netmask can be in a dotted-quad format or as a contiguous mask length	<code>172.22.1.0/24</code> <code>172.22.1.0/255.255.255.0</code>

6. For **Options**, enter a comma-separated list of the following export options. Blank spaces are not allowed in the list.

Export Options	Description
<code>secure</code>	Enter <code>secure</code> to require that requests originate on an internal port less than 1024. To turn off this requirement, enter <code>insecure</code> . The default is <code>secure</code> .
<code>insecure</code>	
<code>rw</code>	Enter <code>rw</code> to allow both read and write access on this NFS virtual export, or enter <code>ro</code> to prevent write access. Normal UNIX file permissions still apply. The default is <code>ro</code> .
<code>ro</code>	
<code>sub</code>	Enter <code>sub</code> to allow clients to mount subdirectories of this NFS virtual export. To turn this off, enter <code>nosub</code> . The default is <code>sub</code> .
<code>nosub</code>	
<code>root_squash</code>	Enter <code>root_squash</code> to map root's UID (user ID) to the anonymous UID. This prevents the root user on the client computer from having root access to files on the file server. To turn it off, enter <code>no_root_squash</code> . This allows the root user on the client computer to access files on the file server with root privileges. The default is <code>root_squash</code> .
<code>no_root_squash</code>	

Export Options	Description
all_squash	Enter all_squash to map all UIDs from the client computer to the anonymous UID. This prevents all users on the client computer from accessing files on the file server with their normal privileges, and instead gives them the privileges of the anonymous user. To turn it off, specify no_all_squash . It specifies that UIDs from the client computer should not be remapped. Note: no_all_squash doesn't undo the effect of root_squash . If both no_all_squash and root_squash are specified, root's UID is still remapped to the anonymous UID. The default is no_all_squash .
no_all_squash	
anon=<integer>	Enter anon to specify both the UID and GID (group ID) to use for the anonymous user. This applies to the root_squash and all_squash options. Enter anonuid to set the anonymous UID for the anonymous user. Enter anongid to set the anonymous GID for the anonymous user. Note: If none of the anon , anonuid , or anongid options are specified, -2 is used for both the anonymous UID and GID.
anonuid=<integer>	
anongid=<integer>	

For example, suppose you specified a **Client** setting of **172.22.1.*** and an **Options** setting of **ro,root_squash,anon=99**.

This specifies that all clients whose IP address matches **172.22.1.*** have read-only access to the NFS export. It also specifies that the root UID and GID should be remapped to the anonymous UID and GID for these clients. Lastly, it specifies that **99** should be used as the anonymous UID and GID.

7. Click **Submit**.

Configuring Clients

If you have deployed the File Director by using the existing names and IP addresses of legacy file servers for virtual servers, and assigned new names and IP addresses to the legacy file servers, then you don't have to reconfigure the clients. (For example, see the deployment scenario in "Scenario 1: Balance Storage Capacity Among Legacy File Servers" on page 1-5.) When the clients use the existing names, they will access the File Director virtual servers and be redirected to the file servers.

But if you have created new names for virtual servers and shares, you must configure clients to access the virtual servers and shares.

CHAPTER 3

Working with Synthetic Directories

This chapter contains the following topics that explain how to set up synthetic directories on the File Director:

- About Synthetic Directories, Synthetic Links, and Union Directories
- Scenarios for Setting Up and Using Synthetic Directories
- Creating and Using Synthetic Directories, Synthetic Links, and Union Directories

Note: The examples in this chapter use NFS terminology and syntax, but the same concepts and File Director functionality apply to CIFS file servers. Only the CIFS configuration details are different.

About Synthetic Directories, Synthetic Links, and Union Directories

A synthetic directory is a hierarchy of mappings to physical shares or exports on one or more file servers.

To specify the location of each mapping of a physical share or export in the synthetic directory hierarchy, you use a *synthetic link*. A synthetic link defines a mapping of a physical share or export into a specific location in the synthetic directory hierarchy. In other words, the synthetic link inserts a physical share or export as a directory in the synthetic directory hierarchy. You can think of a synthetic link as a symbolic link or short cut that allows clients to navigate to various physical shares or exports on various file servers. You can specify one or more synthetic links within a synthetic directory, and the links can map to physical shares or exports on different file servers. By specifying the various synthetic links, you create the hierarchy of the synthetic directory.

A union directory, which is within the hierarchy of a synthetic directory, blends the *contents* of multiple physical shares or exports into one synthetic directory. For example, you can create a union directory that is a blend of multiple home directories on multiple file servers. The difference between a synthetic link and a union directory is that the former does not blend the *contents* of multiple directories into one directory. That is, each synthetic link contains the same items that are in the physical share or export. A union directory is a blend of all items from multiple exports into one union directory in a way that does not exist in the physical share or export.

To present the hierarchy of synthetic links or a union directory to clients, you associate the synthetic directory with a virtual share or export on the File Director. Because a synthetic directory can contain mappings to multiple physical shares or exports on multiple file servers, you can use it to combine multiple physical shares or exports into one virtual share or export.

Clients don't mount or access a synthetic directory directly. Instead, they mount or access the virtual server and virtual share or export that is associated with the synthetic directory.

You can also associate the same synthetic directory with more than one virtual share or export. For example, you might want to use different virtual servers for different subnets. You can associate the same synthetic directory with a virtual share or export on each virtual server.

Scenarios for Setting Up and Using Synthetic Directories

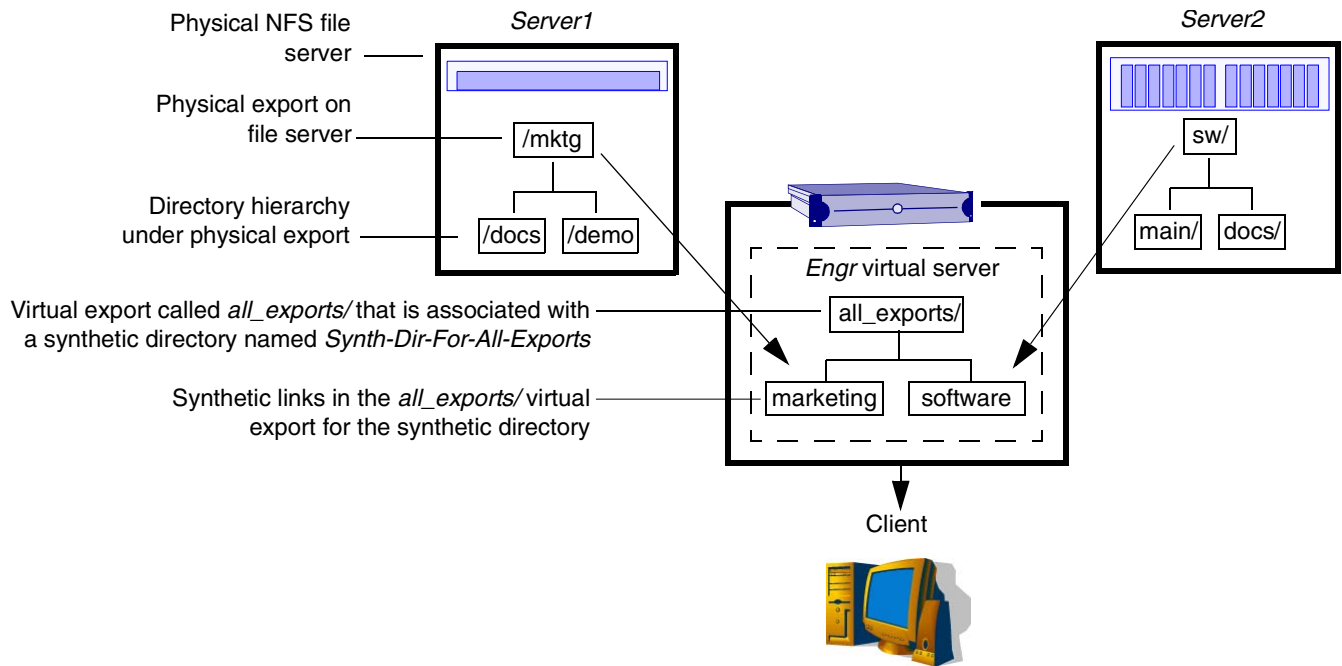
There are three basic scenarios for setting up and using a synthetic directory:

- **Export aggregation**—combines multiple physical shares or exports from multiple file servers into one virtual share or export. See “Synthetic Directory Scenario 1: Export Aggregation,” next.
- **Server aggregation**—through one virtual share or export, clients can navigate through directory hierarchies that are organized based on physical shares or exports on multiple file servers. See “Synthetic Directory Scenario 2: Server Aggregation” on page 3-4.
- **Union directory**—combines the *contents* of multiple physical shares or exports from multiple file servers and presents the combined contents to clients as a single directory in a virtual share or export. See “Synthetic Directory Scenario 3: Union Directory” on page 3-5.

Synthetic Directory Scenario 1: Export Aggregation

You can combine multiple physical shares or exports from multiple file servers into one synthetic directory virtual share or export. In the following example, you can combine the two physical exports called */mktg* and */sw* on *Server1* and *Server2* and present them to clients as one synthetic directory virtual export called */all_exports* on the File Director. The synthetic link *marketing* maps */mktg* into the */all_exports* synthetic directory, and the synthetic link *software* maps */sw*. Note that the names of

the synthetic links (*marketing* and *software*) do not have to match the names of the physical shares or exports (*/mktg* and */sw*).



```
mount -t nfs engr:/all_exports /marketing_and_software
```

A synthetic directory can provide clients with a virtual share or export for accessing data from multiple shares or exports on multiple file servers.

CONFIGURATION OF FILE SERVERS AND THE FILE DIRECTOR FOR THIS SCENARIO

Each of the file servers is configured to provide the exports shown in the following table:

File Server Configuration

File Servers	Physical Exports
Server1	/mktg
Server2	/sw

Here is a summary of the File Director configuration:

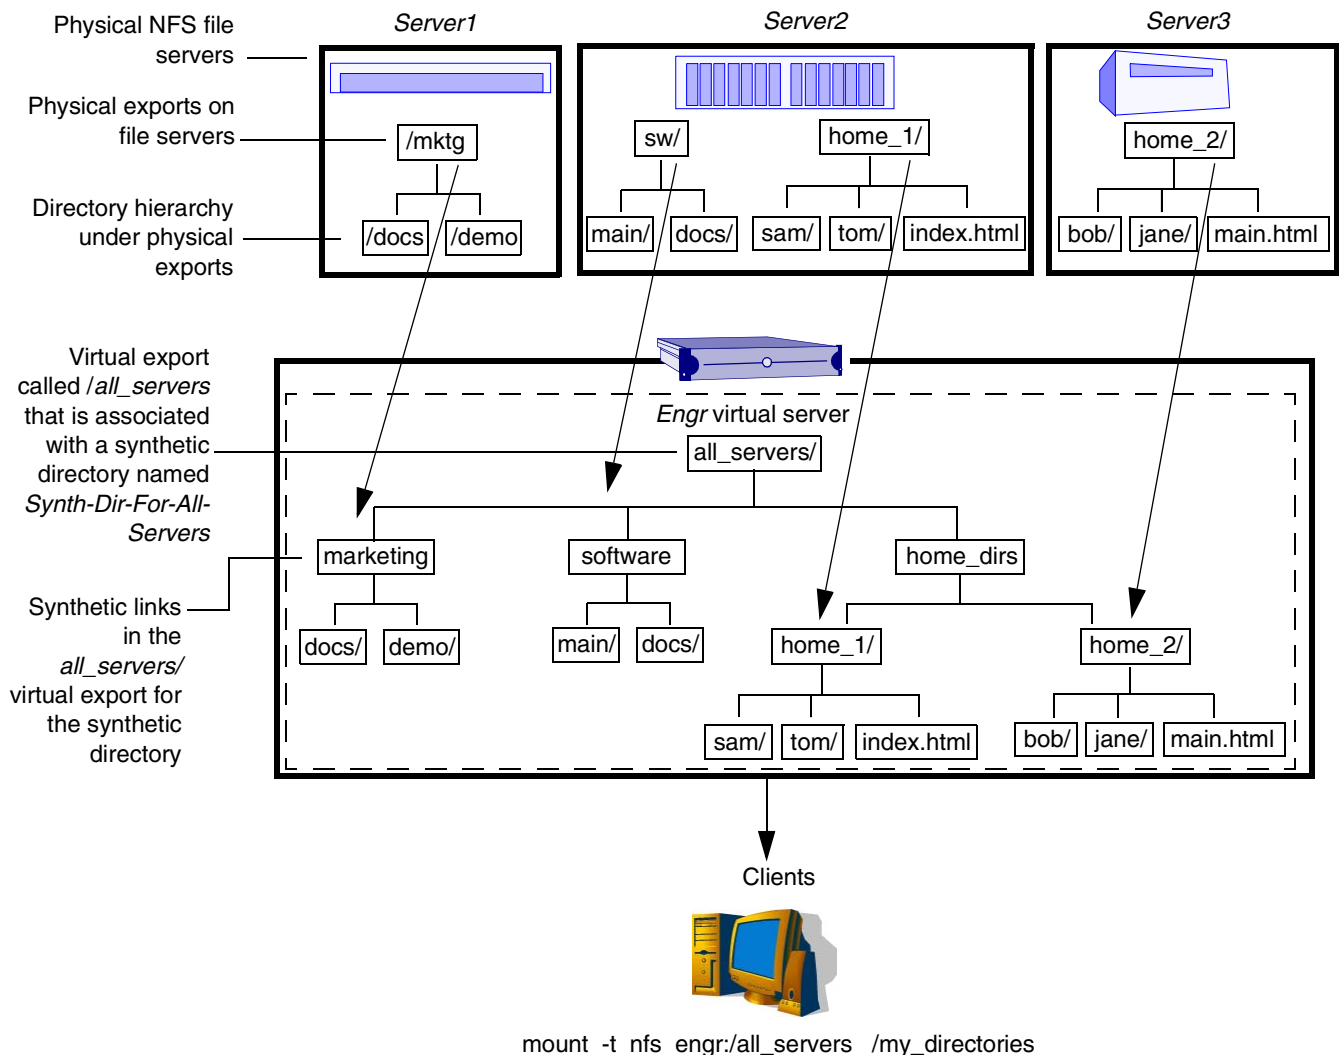
File Director Configuration

This virtual server	And this virtual export	Associated with this synthetic directory	Use these synthetic links on the File Director	To map to these physical servers and exports
Engr	all_exports/	Synth-Dir-For-All-Exports	marketing software	Server1:/mktg/ Server2:/sw/

Synthetic Directory Scenario 2: Server Aggregation

You can create one synthetic directory virtual share or export that allows clients to navigate within the exported directory hierarchies located on different file servers. In the following example, clients can access the */all_servers* virtual export and navigate within a virtual hierarchy that is organized around the physical exports from each file server.

In the following example, *marketing*, *software*, *home_dirs/home_1*, and *home_dirs/home_2* are synthetic links that map to the physical directories on *Server1*, *Server2*, and *Server3*. Note that synthetic links can be at the root level (such as *marketing* and *software*), or they can create a hierarchy in the synthetic directory (such as *home_dirs/home_1* and *home_dirs/home_2*).



A synthetic directory can allow clients to use a virtual share or export to navigate directory hierarchies on multiple file servers.

CONFIGURATION OF FILE SERVERS AND THE FILE DIRECTOR FOR THIS SCENARIO

Each of the file servers is configured to provide the exports shown in the following table:

File Server Configuration

File Servers	Physical Exports
Server1	/mktg
Server2	/sw /home_1
Server3	/home_2

Here is a summary of the File Director configuration:

File Director Configuration

This virtual server	And this virtual export	Associated with this synthetic directory	Use these synthetic links on the File Director	To map to these physical servers and exports
Engr	all_servers/	Synth-Dir-For-All-Servers	marketing software home_dirs/home_1/ home_dirs/home_2/	Server1:/mktg/ Server2:/sw/ Server2:/home_1/ Server3:/home_2/

Synthetic Directory Scenario 3: Union Directory

A union directory on the File Director is a directory under the hierarchy of a synthetic directory. A union directory blends the *contents* of multiple physical shares or exports into one directory. For example, you can create a union directory that is a blend of multiple home directories on multiple file servers.

Similar to a synthetic link, a union directory is contained within a synthetic directory. The difference between a synthetic link and a union directory is that the former does not blend the *contents* of multiple directories into one directory. Each synthetic link contains the same items that are in the physical share or export. A union directory is a blend of all items from multiple exports into one union directory in a way that does not exist in the physical share or export.

As the File Director administrator, keep the following guidelines in mind when you create or change a union directory by using the Management Console or the CLI commands:

- The names of all items in the physical shares or exports for the union directory must be unique before you can create the union directory. For example, the File Director won't let you create a union directory from two physical exports if there are any files or subdirectories in the two physical exports that have the same name.
- The subdirectories of the physical shares or exports for a union directory are not blended. That is, only the items in the top level directory of the physical shares or exports are blended.

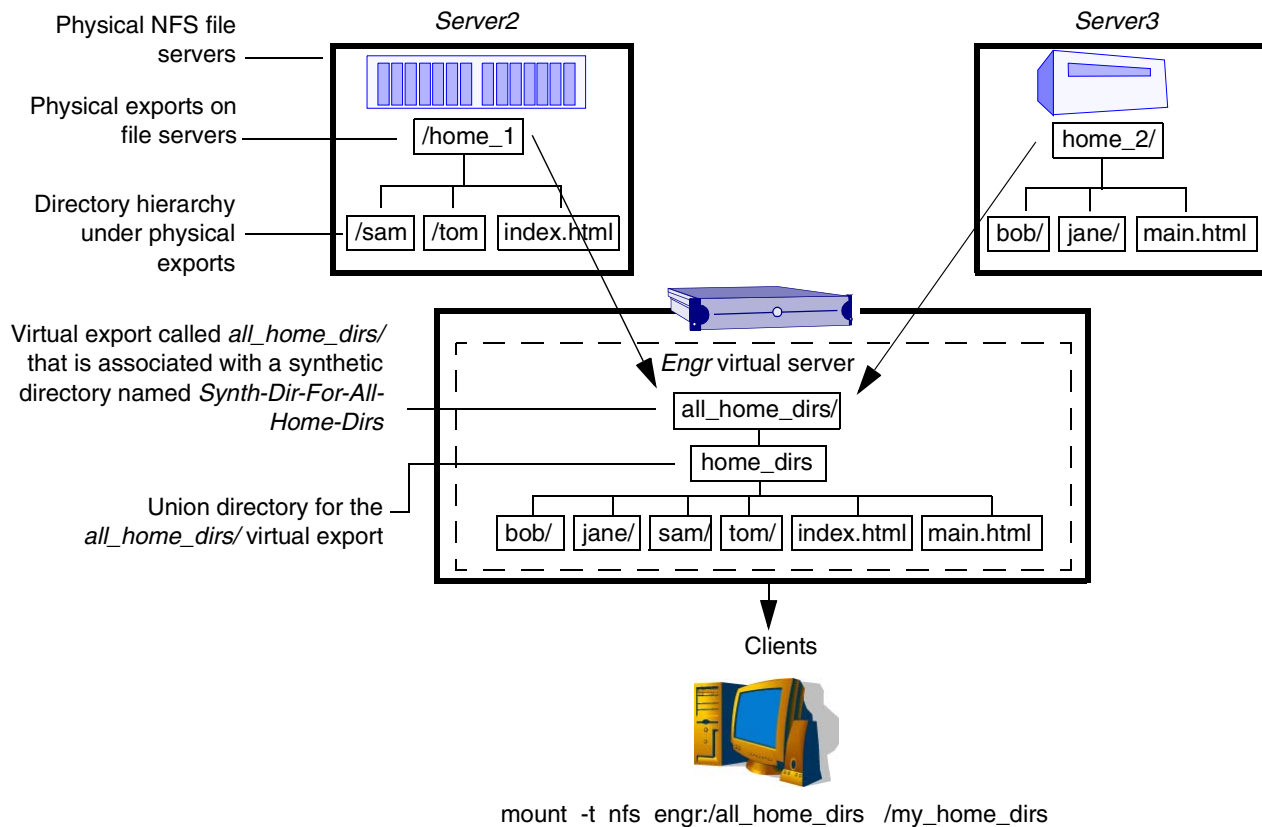
Note: Some special system folders are ignored when creating union directories to prevent file name collisions. The *.snapshot* folder on Network Appliance file servers is ignored, as well as the *lost+found* folder on Linux and Solaris file servers. These folders will not be added to the union directory.

The following restrictions apply to client users who use the File Director to access a union directory:

- Client users cannot use the File Director to create or delete items inside the top level directory of a union directory. This restriction does not apply to items in the subdirectories of the union directory.
- Client users can use the File Director to edit or change existing files inside a union directory.

In the following example, *sam/*, *tom/*, *index.html*, *bob/*, *jane/*, and *main.html* are the contents of the physical exports */home_1* of *Server2* and */home_2* of *Server3*. The

File Director blends these contents into one union directory under the *all_home_dirs/* virtual export.



A union directory presents clients with a single unified directory that contains all of the contents of multiple shares or exports on multiple file servers.

CONFIGURATION OF FILE SERVERS AND THE FILE DIRECTOR FOR THIS SCENARIO

Each of the file servers is configured to provide the exports shown in the following table:

File Server Configuration

File Servers	Physical Exports
Server2	<i>/home_1</i>
Server3	<i>/home_2</i>

Here is a summary of the File Director configuration:

File Director Configuration

This virtual server	And this virtual export	Associated with this synthetic directory	Use this union directory on the File Director	To map to these physical servers and exports
Engr	all_home_dirs/	Synth-Dir-For-All-Home-Dirs	home_dirs home_dirs	Server2:/home_1/ Server3:/home_2/

Creating and Using Synthetic Directories, Synthetic Links, and Union Directories

The general steps for creating and using a synthetic directory, synthetic links, and union directories are:

1. Create the synthetic directory.

For information, see the next section, “Creating a Synthetic Directory.”

2. Do either of the following:

- Create the synthetic links that define the mappings of exports or shares for the synthetic directory.

For information, see “Creating Synthetic Links for a Synthetic Directory” on page 3-9.

- Create one or more union directories for the synthetic directory.

For information, see “Creating Union Directories” on page 3-11.

3. Create a virtual share or export for presenting the synthetic directory.


For information, see “Creating a Virtual Share or Export for a Synthetic Directory” on page 3-13.

Creating a Synthetic Directory

To create a synthetic directory:

1. In the Navigation frame, expand **Synthetic Directories** and then click the **Synthetic Directories** link.

The **Synthetic Directory Configuration** page appears.



2. Click **Create Synthetic Directory** under the **List of Synthetic Directories** section.

The **Synthetic Directory Configuration** page appears.

3. Specify the name of the synthetic directory.

The name of the synthetic directory can be any text string that contains letters, digits, dashes, underscores, percent symbols, or periods. Blank spaces are not allowed.

The File Director uses the name to identify the synthetic directory that contains the synthetic links and union directories that comprise the hierarchy. You use the name when you specify which synthetic directory you want to associate with a virtual share or export that presents the hierarchy to clients. Clients do not use or see the name of the synthetic directory.

For example, the name of the synthetic directory could be *Synth-Dir-For-All-Exports*.

4. Choose an access protocol for the synthetic directory from the **Access Protocol** drop-down menu.

The setting you select from the **Access Protocol** drop-down menu determines the type of virtual shares or exports you can use to present this synthetic directory. For example, if you specify an NFS access protocol, you can only use NFS virtual exports to present the synthetic directory.

5. Click **Submit**.

Creating Synthetic Links for a Synthetic Directory

After you create a synthetic directory, you can create the synthetic links for the directory. Each synthetic link defines a *one-to-one* mapping of a physical share or export into a specific location in the synthetic directory hierarchy.

Important: The data associated with a synthetic link is stored on the physical file server and not on the File Director itself. When you create a synthetic link, make sure that the share or export to be linked is writable by Administrator or root.

To create a synthetic link:


1. If you haven't already done so, specify information on the **File Server Configuration** page about the file server(s) that contain the physical shares or exports you want to present in the synthetic directory.

For more information, see "Specifying a File Server" on page 2-16.

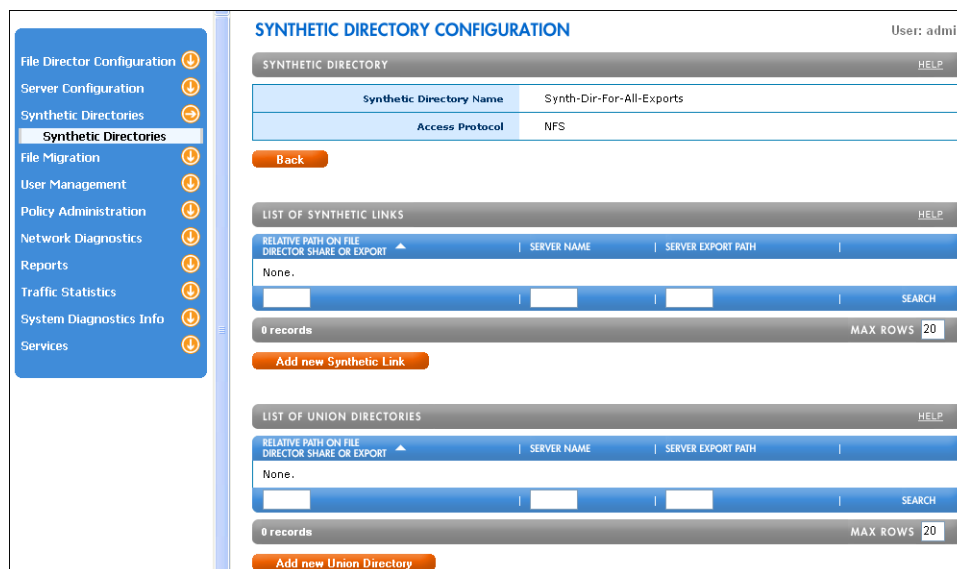
For example, specify information about the *Server1* and *Server2* file servers that are shown in the Export Aggregation example scenario. (See "Synthetic Directory Scenario 1: Export Aggregation" on page 3-2.)

2. In the Navigation frame, expand **Synthetic Directories** and then click the **Synthetic Directories** link.

- On the **Synthetic Directories Configuration** page, click the **Display** link next to the synthetic directory where you want to create the synthetic link.



The **Synthetic Directory Configuration** page appears.



- Click **Add New Synthetic Link** under the **List of Synthetic Links** section.

The **Configure a Synthetic Link** page appears.

- For **Path on File Director Share or Export**, enter the path in the synthetic directory hierarchy you want to present to clients as the symbolic link or shortcut that is mapped to the physical share or export. You can enter either a relative or absolute path.

Clients will use the path you enter to navigate to the data in the physical share or export.

For NFS for example, enter */marketing* for the first synthetic link that is shown in the Export Aggregation example scenario.

- For **File Server Name**, enter the name of the file server that contains the physical share or export you want to map the synthetic link to.

For example, enter *Server1* for the file server that is shown in the Export Aggregation example scenario.

7. For **File Server Share or Export**, enter the name of the physical share or export you want to map the synthetic link to.

For example, enter */mktg* for the physical export on the *Server1* file server that is shown in the Export Aggregation example scenario.

8. Click **Submit**.

If all of the settings are correct, the synthetic link is shown in the **List of Synthetic Links** section.

9. To add another synthetic link to the synthetic directory hierarchy, repeat steps 4 through 8.

For example, to define the second NFS synthetic link that is shown in the Export Aggregation example scenario:

- Enter *software* for the relative path.
- Enter *Server2* for the file server.
- Enter */sw* for the physical export on the *Server2* file server.

Now you're ready to create the virtual share or export to present the synthetic directory. For more information, see "Creating a Virtual Share or Export for a Synthetic Directory" on page 3-13.

Creating Union Directories

A union directory is a special type of synthetic directory that allows you to blend the *contents* of multiple physical shares or exports into one directory. For more information, see "Synthetic Directory Scenario 3: Union Directory" on page 3-5.

Important: The data associated with a union directory is stored on the physical file server and not on the File Director itself. When you create a union directory, make sure that the share or export to be linked is writable by Administrator or root.

To create a union directory:

1. Create a synthetic directory.

For more information, see "Creating a Synthetic Directory" on page 3-8.

For example, specify information about the *Server2* and *Server3* file servers and the *Engr* virtual server that are shown in the Union Directory example. The name of the synthetic directory could be *Synth-Dir-For-All-Home-Dirs*.

- On the **Synthetic Directory Configuration** page, click the **Display** link next to the synthetic directory where you want to create the union directory.

The screenshot shows the 'SYNTHETIC DIRECTORY CONFIGURATION' page. On the left is a navigation menu with 'Synthetic Directories' selected. The main content area has a 'LIST OF SYNTHETIC DIRECTORIES' table with columns for 'DIRECTORY NAME', 'ACCESS PROTOCOL', and actions. One entry is visible: 'Synth-Dir-For-All-Home-Dirs' with protocol 'NFS'. There are 'Delete' and 'Display' buttons for this entry. Below the table, it says '1 - 1 of 1 records' and 'MAX ROWS 20'. A 'Create Synthetic Directory' button is at the bottom.

The **Synthetic Directory Configuration** page appears.

This screenshot shows the 'SYNTHETIC DIRECTORY CONFIGURATION' page after clicking 'Display'. The 'SYNTHETIC DIRECTORY' section shows details for 'Synth-Dir-For-All-Home-Dirs' with 'Access Protocol' set to 'NFS'. A 'Back' button is below. The 'LIST OF SYNTHETIC LINKS' table is empty, showing '0 records'. Below it is an 'Add new Synthetic Link' button. The 'LIST OF UNION DIRECTORIES' table is also empty, showing '0 records', with an 'Add new Union Directory' button at the bottom.

- Click **Add New Union Directory** under the **List of Union Directories** section.

The **Configure a Union Directory** page appears.

- For **Path on Export or Share**, enter the relative or absolute path you want to present as the symbolic link or shortcut for the union directory that contains all of the physical shares or exports.

Clients will use the path you enter to navigate to the data in the set of physical shares or exports.

For NFS, for example, enter *home_dirs* for the union directory that is shown in the Union Directory example scenario.

- For **File Server Name**, enter the name of the file server that contains the physical share or export you want to map the union directory to.

For example, enter *Server2* for the file server that is shown in the Union Directory example scenario.

6. For **File Server Share or Export**, enter the name of the physical share or export you want to map the union directory to.

For example, enter */home_1* for the physical export on the *Server2* file server that is shown in the Union Directory example scenario.

7. Click **Submit**.

If all of the settings are correct, the union directory is shown in the **List of Union Directories** section.

8. To add more physical shares or exports to the union directory, repeat steps 3 through 7 using the same path.

For example, to add the */home_2* physical export to the union directory that is shown in the Union Directory example scenario:

- Enter *home_dirs* for the relative path.
- Enter *Server3* for the file server.
- Enter */home_2* for the physical export on the *Server3* file server.

Now you're ready to create the virtual share or export to present the synthetic directory for the union directory. For example, create the *all_home_dirs/* virtual export that is shown in the Union Directory example scenario. For more information, see "Creating a Virtual Share or Export for a Synthetic Directory," next.

Creating a Virtual Share or Export for a Synthetic Directory

After you create a synthetic directory with at least one synthetic link or union directory, you can create the virtual share or export that the File Director presents to clients for that synthetic directory.

To create a virtual share or export for a synthetic directory:


1. If you haven't already done so, specify information on the **Virtual Server Configuration** page about the virtual server you want to use to present the virtual share or export for the synthetic directory.

For more information, see "Specifying a Virtual Server" on page 2-20.

For example, specify information about the *Engr* virtual server that is shown in the Export Aggregation example scenario.

2. On the **Virtual Server Configuration** page, click **Display** next to the virtual server for the synthetic directory.

For example, click **Display** next to the *Engr* virtual server.



The **Virtual Shares & Exports** section appears at the bottom of the page for the virtual server you selected.

3. Do one of the following:
 - For CIFS, click **Create CIFS Virtual Share** under the **Virtual Shares & Exports** section.

The **CIFS Virtual Share Configuration** page appears.

- For NFS, click **Create NFS Virtual Export**.

The **NFS Virtual Export Configuration** page appears.

Note: The setting you selected from the **Access Protocol** drop-down menu on the **Synthetic Directory Configuration** page determines the type of virtual shares or exports you can use to present the synthetic directory. For example, if you specified an NFS access protocol, you can only use NFS virtual exports to present the synthetic directory. For more information, see “Creating a Synthetic Directory” on page 3-8.

4. Do one of the following:
 - For **CIFS Virtual Share Name**, enter the name of the CIFS virtual share you want clients to access for the synthetic directory.
A CIFS virtual share name can be any text string. Do not include a backslash. For example, enter *all_exports*.
 - For **NFS Virtual Export Name**, enter the name of the NFS virtual export that you want clients to access for the synthetic directory.
An NFS export name can be any absolute path name that includes a slash. For example, enter */all_exports* for the virtual export that is shown in the Export Aggregation example scenario.
5. For **Export Type**, choose **Synthetic Directory**.
6. Choose the name of the synthetic directory from the **Synthetic Directory** drop-down menu.

For example, choose *Synth-Dir-For-All-Exports*.

7. Click **Submit**.

If all of the settings are correct, the virtual share or export is shown in the **Virtual Shares & Exports** section. Clients can now use the virtual server and share or export to access the synthetic directory.

Deleting a Synthetic Directory

Before you can delete a synthetic directory, you must delete all virtual shares or exports that present the synthetic directory.

To delete a synthetic directory:

1. On the **Virtual Server Configuration** page, delete all virtual shares or exports that present the synthetic directory.
2. On the **Synthetic Directory Configuration** page, delete the synthetic directory.

Using the CLI to Work with Synthetic Directories

In some cases, you may wish to use the CLI instead of the Management Console to create synthetic directories. This section provides a task-based overview of the command combinations needed to create synthetic directories. For detailed command descriptions, see the *CLI Command Reference*.

About Directory Storage

When you create a synthetic directory and add synthetic links and union directories, you are essentially specifying a tree structure (or mapping) of paths to the synthetic links and union directories. For example, if you create synthetic links with the following paths:

sales

\homes\home1

\homes\home2

the corresponding artificial tree structure consists of:

**

\sales

\homes

\homes\home1

\homes\home2

When a synthetic directory is exported as a virtual share, this tree structure is presented to users. To users, the end nodes in the structure appear to contain the shares to which the synthetic links point. Union directories work in a similar fashion.

To leverage NeoPath's virtualization technology, File Director actually creates and stores this tree structure within the special **.neopath** directory on a backend share. The share that holds this tree structure provides “directory storage” to the synthetic directory.

In the Management Console, directory storage is created automatically. The share that is used for the first synthetic link (or union directory) added to a synthetic directory provides directory storage for the synthetic directory. If you use CLI commands to create a synthetic directory, you must use the **dstorage** command to manually specify a backend share for directory storage.

For redundancy, NFS allows more than one directory storage to be used for a single synthetic directory. However, no more than one directory storage for the same synthetic directory may come from the same backend server. One directory storage is considered *primary*. All others are labeled *secondary*.

CIFS supports only one directory storage for each synthetic directory. This directory storage is, by default, the primary.

Procedure: Adding a Synthetic Directory

The following steps provide an overview of how to use CLI commands to add a synthetic directory. You can use this as a model for creating your own synthetic directories. For detailed command descriptions, see the *CLI Command Reference*.

1. Use the **sdirectory** command to create a new synthetic directory. For example:

```
add sdirectory Synth-Dir-For-All-Servers NFS
```
2. Use the **dstorage** command to create directory storage. For example:

```
add dstorage Synth-Dir-For-All-Servers server1 /export/dstorage  
primary
```
3. Use the **slink** command to add synthetic links. For example:

```
add slink Synth-Dir-For-All-Servers /marketing server1 /mktg/  
add slink Synth-Dir-For-All-Servers /software server2 /sw/  
add slink Synth-Dir-For-All-Servers /home_dirs/home_1/  
server2 /home_1/  
add slink Synth-Dir-For-All-Servers /home_dirs/home_2/  
server3 /home_2/
```
4. If desired, use the **udirectory** command to add union directories. For example:

```
add udirectory Synth-Dir-For-All-Servers /home_dirs  
server1 /home_1/  
add udirectory Synth-Dir-For-All-Servers /home_dirs  
server2 /home_2/
```
5. Make the directories available to users by exporting them as virtual shares. You can perform this step before or after you add the synthetic links and union directories.

CHAPTER 4

Using the File Director to Migrate Data

This chapter contains the following topics that explain how to use the File Director to migrate data between file servers:

- About Migrating Data
- Configuring File Server Security Before Migrating Data
- Creating a Migration Job
- Checking the Status of a Migration Job that is Running
- Displaying a List of Migration Jobs
- Displaying the File Migration Log
- Retrying a Failed Migration Job
- Displaying Migration Statistics for a Physical Share or Export
- Finding the Current Location of Files or Directories

About Migrating Data

By using the File Director, you can migrate data between file servers at any time with no disruption to clients who may be accessing the data at the same time. (Clients may experience slower access due to migration traffic.)

Before migrating data, you must configure the File Director to recognize the source file server that has the data you want the File Director to migrate, and the destination file server where you want the data to be migrated. For more information, see “Specifying a File Server” on page 2-16.

Note: Migrated files are placed in the special *.neopath* directory on the destination share or export. This is done to prevent NFS or CIFS clients from directly accessing the migrated files, which are now being managed by the File Director. That is, clients cannot use the File Director to access the *.neopath* directory. To clients, the migrated files appear to be in their original location.

Important: A migration job may fail if the source file server has no available disk space for the File Director to store information about the migration task. Additionally, during migration, files are not removed from the source file server until they have all been copied to the destination file server. Consequently, a migration will not free up any disk space on the source file server until the migration job is nearly finished. If the source file server has no available disk space, you can either free up some disk space before running the complete migration job, or else break up the complete migration job into multiple smaller jobs.

Configuring File Server Security Before Migrating Data

Before using the File Director to migrate data, be sure to configure file server security according to the following guidelines.

General CIFS Migration Guideline

The following general guideline applies to all migrations using CIFS file servers.

- Use the same or trusted domains for CIFS migrations.

Each file and directory on a CIFS server has an owner, a primary group, and an access control list (ACL). The ACL contains a list of users and groups, and the access each are permitted or denied. In order to maintain the same file security on migrated files and directories, the destination server must be able to understand the owner, primary group, and ACL of the source file or directory.

A destination server can understand only users and groups that are one of the following:

- Predefined by Microsoft
- Local to the destination server
- Members of the destination server's domain
- Members of a domain that is trusted by the destination server's domain

A destination server cannot understand users or groups that are either:

- Local to a different server
- Members of untrusted domains

Network Appliance CIFS Migration Guidelines

The following guidelines apply to all migrations using Network Appliance CIFS file servers.

- Set the QTree security to NTFS or Mixed.

Network Appliance file servers divide their storage into QTrees. In order to migrate CIFS file-level security settings, the destination QTree must use either the NTFS or Mixed security mode. If a QTree uses the UNIX security mode, then the CIFS migration will fail. For information on how to configure your QTree's security mode, see your Network Appliance documentation.

- Avoid NTFS streams.

Windows NTFS supports a feature called *streams*, which allow an application to store extra information embedded within a file. Network Appliance file servers do not support NTFS streams. To avoid data loss, the File Director will report an error if you try to migrate a file with an NTFS stream to a Network Appliance file server.

- Map migration administrator to root.

Migration requires full access to the source and destination directories. To ensure full access, the Network Appliance file server should be configured to map the CIFS migration administrator account to the NFS root account. You can do this by either:

- Setting the `waf.l.nt_admin_priv_map_to_root` option to On.
- Adding an `/etc/usermap.cfg` entry.

For more information, see your Network Appliance documentation.

Samba CIFS Migration Guidelines

Samba implements CIFS file security by using POSIX file-level security. The following guidelines apply to all migrations to or from a Samba CIFS file server.

- Map domain users/groups to POSIX users/groups.

You must configure Samba to map CIFS domain users and groups to POSIX users and groups. The simplest way to do this is to use the Samba “winbindd” program. For more information, see your Samba documentation.

- Enable support for file-level POSIX access control lists (ACLs).

To support migration to Samba file servers, all of the following must be true:

1. The operating system kernel must support POSIX ACLs.

Check with your operating system vendor or documentation to verify that it supports POSIX ACLs. Note that Linux 2.4 can be compiled with or without support for ACLs. For more information, see: <http://acl.bestbits.at/>

2. The volume must be mounted with ACL support enabled.

See your operating system or file system documentation for instructions on mounting volumes with ACL support. For Linux, here are some ways you can verify whether your volume is mounted with ACL support:

- Run `mount` and see if `acl` is in the options list for the target volume.
- Try using `setfacl` to set permissions for various users. For example:
`setfacl -m user:<user_name>:r foo.txt`

3. The `smbd` executable must be compiled using the `with-acl-support` option. To check this, use the following command:

```
smbd --build-options | grep -i acl
```

The following indicates that `smbd` was compiled with ACL support:

```
HAVE_SYS_ACL_H
```

```
HAVE_POSIX_ACLS
```

The following indicates that `smbd` was compiled without ACL support:

```
HAVE_NO_ACLS
```

■ Configure “winbindd” to use consistent UID/GID.

By default, Samba’s “winbindd” generates a different SID-to-UID/GID mapping on each file server. This can cause local file permissions to change when migrating files.

To prevent these problems, configure “winbindd” to use a consistent mapping from SIDs to UIDs and GIDs across servers. You can do this by setting either of the following parameters in the Samba configuration file:

- “winbind trusted domains only = yes”
- “idmap backend”

For more information on how to configure “winbindd” to use consistent SID-to-UID/GID mappings, see your Samba documentation.

■ Change file owners from groups to users.

CIFS allows a group to be the owner of a file or directory. However, POSIX security only allows users to own a file or directory. Consequently, all files and directories being migrated to a Samba server must be owned by a user, not a group.

■ Map migration administrator to root.

To migrate data, the File Director requires full access to the source and destination directories. To ensure full access, configure Samba to map the CIFS migration administrator account to the NFS root account. You can do this by either:

- Configuring NIS to map the administrator account to UID 0.
- Using the Samba “username map” option.

For more information, see your Samba documentation.

■ Avoid NTFS streams.

Windows NTFS supports a feature called *streams*, which allow an application to store extra information embedded within a file. Samba does not support NTFS streams. To avoid data loss, the File Director will report an error if you try to migrate a file with an NTFS stream to a Samba server.

General Multi-protocol Migration Guidelines

The following general guidelines apply to all migrations using multi-protocol file servers that support both CIFS and NFS.

■ Use the same or trusted CIFS domains.

Since a multi-protocol migration includes CIFS file security information, it has the same domain requirements as a CIFS-only migration. For more information, see “General CIFS Migration Guideline” on page 4-2.

■ Use identical user/group mappings for CIFS and NFS.

CIFS identifies each user and group by a number called a Security Identifier (SID). NFS identifies each user by a user ID number (UID), and each group by a group ID number (GID). A multi-protocol server converts between SIDs and UIDs/GIDs. To maintain the same file security for both CIFS and NFS, the source and destination file servers must use the same SID-to-UID/GID conversion.

- Use the same NIS domain for NFS.

Do the following to have NIS ensure that all NFS servers are using the NFS UIDs and GIDs, and that they use the same CIFS-to-NFS mappings:

- Configure the servers to use the same NIS server.
- Add an NIS entry for each Windows user and group that appears as a CIFS owner, primary group, or ACL entry.

Note: CIFS write access is blocked during a multi-protocol migration and CIFS connections that access migrated files are terminated when a multi-protocol migration is completed.

Windows NFS Multi-protocol Migration Guidelines

The following guidelines apply to all multi-protocol migrations using Windows NFS file servers.

- Add owner and primary group to CIFS ACL.

NFS always reports permissions for the file's owner, primary group, and "others". Windows NFS emulates this by converting the CIFS ACL into NFS permissions. If a file or directory has an ACL that doesn't include the owner or primary group, the Windows NFS server will report permissions that don't reflect the actual user permissions. To avoid problems, add the file owner and primary group to the CIFS ACL.

Windows NFS is one of the few Windows programs that makes use of a file's primary group. Because file primary groups are so little used, most versions of Windows do not include a tool for viewing them. For a Windows server in a domain, a file's primary group is usually "Domain Users". If you want to verify a file's primary group, you can use Microsoft's "subinacl" utility, which is downloadable from the Microsoft web site.

- Map UID 0 to member of Administrators group.

Multi-protocol migration relies on NFS UID 0 to represent a user with administrative or "superuser" privileges. That is not automatically true on a Windows NFS server. To make UID 0 a superuser, you must configure the Windows Services for UNIX "User Name Mapping" service to map UID 0 to a Windows user that is a member of the Administrators. For information on configuring "User Name Mapping," see your Services For Unix documentation.

Network Appliance Multi-protocol Migration Guideline

The following guideline applies to all multi-protocol migrations using Network Appliance file servers.

- Set QTree security to Mixed.

Network Appliance file servers divide their storage into QTrees. In order to support both NFS and CIFS file-level security settings, the destination QTree should use the Mixed security mode. For information on how to configure your QTree's security mode, see your Network Appliance documentation.

Samba and NFS Multi-protocol Migration Guidelines

The following guidelines apply to all multi-protocol migrations using Samba file servers.

- Enable support for file-level POSIX access control lists (ACLs).
- Configure “winbindd” to use consistent UID/GID.

For more information, see the descriptions for these two guidelines in “Samba CIFS Migration Guidelines” on page 4-3.

Creating a Migration Job

This section explains how to manually create and run migration jobs. You can also use policies to automate file migrations. For more information, see Chapter 5, “Using Policies for Automatic File Migration.”

Note: When running a migration job manually, you can only migrate entire directories and not individual files. You can, however migrate an individual file by using a policy.

The File Director runs only one migration job at a time. If you create a series of migration jobs immediately one after another, each migration job will begin only after the previous migration job has finished.

Note: Migrations are always performed from one physical location to another. You cannot specify virtual shares or exports in a migration job.

To create a migration job:

1. In the Navigation frame, expand **File Migration** and then click the **Migrate Files** link.

The **Create New Migration Job** page appears.

The screenshot shows the 'CREATE NEW MIGRATION JOB' page. The left sidebar has a 'File Migration' section with 'Migrate Files' selected. The main area has a title bar 'CREATE NEW MIGRATION JOB' with a 'HELP' link. Below the title bar, there are four steps: Step 1: Enter the export protocols (with checkboxes for NFS and CIFS), Step 2: Enter the source server name, export and path (with fields for Source Server Name, Source Export, and Source Path), Step 3: Enter the destination server name and export (with fields for Destination Server Name and Destination Export), and Step 4: Submit the request or cancel and try again later. At the bottom of Step 4 are 'Submit' and 'Cancel' buttons. The user 'admin' is logged in.

2. For **Protocol Types**, select NFS, CIFS, or both protocols to use for the migration job.

If you select both NFS and CIFS protocols, it is called a *multi-protocol* migration.

Keep the following in mind when selecting the protocol type(s):

- The physical source and destination file servers must support the protocol(s) that you select.
- If you want to create a multi-protocol migration job, be sure you have already done *both* of the following tasks:
 - Selected both NFS and CIFS protocols when you set up the source file server in the File Director configuration. This is called a *multi-protocol* file server. For more information, see “Specifying a File Server” on page 2-16. If you haven’t already done this, you won’t be able to create the migration job.
 - Matched the source CIFS share name in the migration job with the source NFS export name (or vice versa) to create an association between the two names. For more information, see “Matching or Unmatching Physical Shares and Exports” on page 2-19. If you haven’t already done this, you won’t be able to create the migration job.
- Suppose the source file server is set up to be multi-protocol in the File Director configuration. If you *have* matched the migration source share with another CIFS share name or NFS export name, then you *cannot* create a single-protocol migration job. If you *have not* matched the migration source share, then you *can* create a single-protocol migration job by confirming the job after you click **Submit**.
- Be careful when choosing whether or not to use a multi-protocol migration or a single-protocol migration. If a physical share or export is accessible via both CIFS and NFS, then all migrations to and from that share or export should also be multi-protocol migrations. If an NFS-only or CIFS-only migration is performed, the migrated files will no longer be accessible via the other protocol. This applies even if the File Director is not re-exporting the share or export using the other protocol when the migration is performed.

3. Choose the source file server from the **Source Server Name** drop-down menu.

4. For **Source Export**, enter the physical source share or export where you want to migrate files from.

For an NFS export, enter the export location. For example:

`/export/sw`

Or for a CIFS share, enter the share name. For example:

`mktg`

5. For **Source Path**, enter the path to the directory within the export or share that contains the files you want to migrate.

For a NFS export, for example, enter:

/main

Or for a CIFS share, for example, enter:

\docs

Note: You must specify a subdirectory within the CIFS share or NFS export. Currently, File Director does not support migration of the root directory.

6. Choose the destination file server from the **Destination Server Name** drop-down menu.
7. For **Destination Export**, enter the destination share or export where you want to migrate files to.
8. Click **Submit**.

If all of the settings are correct, the File Director creates the migration job and begins migrating the files from the source file server to the destination file server as soon as possible.

Note: If the migration job fails for any reason, such as insufficient disk space on the destination file server, none of the files are migrated from the source file server. That is, a partial migration is *never* performed.

Checking the Status of a Migration Job that is Running

You can check the status of a migration job while it is running. You can also view pie charts that describe conditions before and after the migration job.

To check the status of a migration job that is running:

1. In the Navigation frame, expand **File Migration** and then click the **Migration Status** link.

The **Display Migration Status** page appears.

To show charts of migration data, click this link.

DISPLAY MIGRATION STATUS User: admin

MIGRATION STATUS HELP

Elapsed Time	0H 2M 58S
Creation Date	Wed April 6, 2005
Creation Time	11:30:28
Starting Date	Wed April 6, 2005
Starting Time	11:30:28
Job Status	running
Total Number of Files with Errors	0

Before and After Charts [Show](#)

Export Type	NFS
Source Server Name	server1
Source Export	/export/misc
Destination Server Name	server2
Destination Export	/export/archive
Destination Path	/main

[Refresh](#)

LIST OF MIGRATION SOURCE PATHS HELP

SOURCE PATH ▲

None.

0 records MAX ROWS 20

2. To display the latest status for the job, click **Refresh**.

Displaying a List of Migration Jobs

You can display a list of the currently running migration job, the pending migration jobs, or the completed jobs. Because the File Director runs only one migration job at a time, any additional jobs are listed as a pending job. Migration jobs that result from a policy are also placed in the pending list.

To display a list of migration jobs:

1. In the Navigation frame, expand **File Migration** and then click the **All Migration Jobs** link.

The **Migration Jobs Table** page appears.

JOB CREATION DATE	JOB CREATION TIME	JOB STATUS	SOURCE HOST	SOURCE EXPORT		
Wed April 6, 2005	16:01:00	running	server1	/export/misc	Delete	Display
Wed April 6, 2005	14:06:54	pending	server1	/export/misc	Delete	Display

2. To display the latest status for the job, click **Refresh**.
3. To display details about a job or retry a failed job, click the **Display** link next to the job.
4. To delete a job from the **Migration Jobs Table**, click the **Delete** link next to the job.

Note: Deleting a job from the **Migration Jobs Table** does not abort a running job. It simply deletes it from the list.

5. To display a list of the currently running migration jobs, do either of the following:
 - Click **Show Running Jobs**.
 - In the Navigation frame, expand **File Migration**, expand **Migration Jobs**, and then click the **Running** link.
6. To display a list of pending migration jobs, do either of the following:
 - Click **Show Pending Jobs**.
 - In the Navigation frame, expand **File Migration**, expand **Migration Jobs**, and then click the **Pending** link.
7. To display a list of completed migration jobs, do either of the following:
 - Click **Show Completed Jobs**.
 - In the Navigation frame, expand **File Migration**, expand **Migration Jobs**, and then click the **Completed** link.

Displaying the File Migration Log

The File Director creates a migration log that contains status information about migration jobs that have run. The File Director rotates log files when they reach a certain size or based on a certain schedule. File Director rotates a log file by saving the current log file to *migration.log.1*, and then by saving *migration.log.1* to *migration.log.2*, and so on.

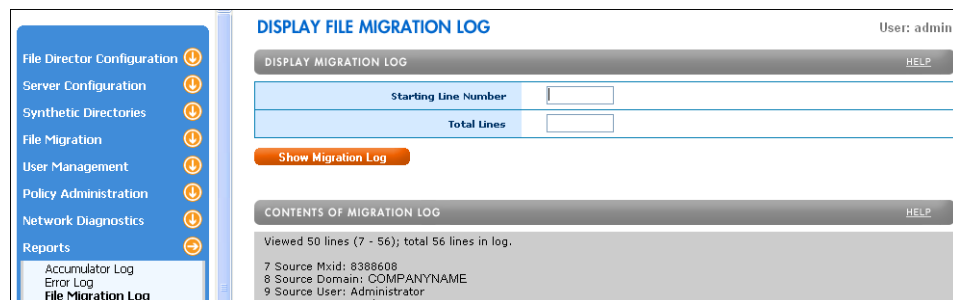
For information on configuring how the File Director rotates log files, see “Configuring Log File Rotation” on page 6-8.

Tip: You can also obtain information about a migration status or failure by using SNMP. For more information, see “Configuring SNMP on the File Director” on page 6-1.

To display the file migration log:

1. In the Navigation frame, expand **Reports** and then click the **File Migration Log** link.

The **Display File Migration Log** page appears showing a log for the migration jobs run since the last time the log was cleared.



2. To display a particular section of the log, enter a starting line number and the total number of lines you want to display, and then click **Show Migration Log**.

Retrying a Failed Migration Job

If a migration job fails, you can retry it after determining why the job failed and taking the appropriate steps to fix the problems. You can only retry jobs that have failed.

To retry a migration job:

1. In the Navigation frame, expand **File Migration** and then click the **All Migration Jobs** link.

The **Migration Jobs Table** page appears.

2. Click **Display** next to the migration job you want to retry.
3. Click **Retry**.

The **Create New Migration Job** page appears with the information from the failed job entered in the fields.

4. Change any information as necessary, and then click **Submit**.

The File Director adds a new entry for the changed migration job to the **Migration Jobs Table**, and the entry for the failed job remains unchanged.

Displaying Migration Statistics for a Physical Share or Export

To display migration statistics for a physical share or export:

1. In the Navigation frame, expand **Server Configuration** and then click the **File Servers** link.

The **File Server Configuration** page appears.

2. Click **Display** next to the file server that has the share or export you want to display statistics for.

The **File Server Configuration** page appears.

3. Scroll down the **File Server Configuration** page to see the currently defined list of shares and exports.

For example:

PHYSICAL SHARES & EXPORTS		Page: 1 2 3 4	HELP
CIFS PHYSICAL SHARE ▲	NFS PHYSICAL EXPORT		
common	/export/common		Migration Stats
	/export/src		Migration Stats
homes	/export/home		Migration Stats
	/export/scratch		Migration Stats
software	/export/software		Migration Stats
	/vol/vol1/film6		Migration Stats
		SEARCH	
1 - 6 of 20 records		Page: 1 2 3 4	MAX ROWS 6
Match Share to Export			

4. Click **Migrations Stats** next to the share or export you want to display statistics for.
The **Share and Export Migration Stats** page appears.
5. To reset the statistics for the share, click **Reset Stats**.

Finding the Current Location of Files or Directories

You can find the current location of a files or directories by specifying the original path. For example, suppose you migrate a directory */export/sw/apps/mkt* from *server1* to *server2*. To find the current location of that directory, you specify the original share name */export/sw* and path of the directory */apps/mkt* on *server1*.

There are two scenarios for finding files or directories:

- If you know the original physical location of the files or directories, or if you're browsing the file server directly and you see a migration point, perform a search for a physical location. For more information, see the next section, "Finding the Location of a Physical File or Directory."
- If you're accessing the files or directories via a virtual share and you want to know the current location, perform a search for a virtual location. For more information, see "Finding the Location of a Virtual File or Directory" on page 4-14.

Finding the Location of a Physical File or Directory

To find the location of a physical file or directory:

1. In the Navigation frame, expand **File Migration**, expand **Find File Location**, and then click the **Physical** link.

The **Physical File Location** page appears.

The screenshot shows the 'Physical File Location' page. On the left is a navigation menu with 'File Migration' expanded, showing 'Physical' as the selected option. The main panel has a title 'PHYSICAL FILE LOCATION' and a user indicator 'User: admin'. Below the title is a 'FIND FILE LOCATION' section with a 'Physical Server' dropdown menu currently showing 'server1'. There are input fields for 'Share or Export Name' and '*Path'. A 'Find' button is located below these fields. Underneath is an '*optional field' section. It contains a 'FILE SOURCE PATH' section with a 'SOURCE PATH' dropdown menu set to 'None'. Below that is a 'LOCATION LIST' section with a 'LOCATION' dropdown menu also set to 'None'. Each of these optional sections has a 'HELP' link.

2. Choose the name of the original file server from the **Physical Server** drop-down menu.

For example, choose *server1*.

3. For **Share or Export Name**, enter the original physical share or export.

For example, for an NFS export, enter:

`/export/sw`

Or for a CIFS share, enter the share name:

`mktg`

4. For **Path**, enter the original path of the physical file or directory as follows:

- For NFS, enter the path in the following format: `/<path>`
- For CIFS, enter the path in the following format: `\<path>`

For an NFS directory, for example, enter:

`/main`

Or for a CIFS directory, for example, enter:

`\docs`

5. Click **Find**.

The File Director displays the current location of the file or directory. If the file or directory has not been migrated, the existing location is returned.

Finding the Location of a Virtual File or Directory

To find the location of a virtual file or directory:

1. In the Navigation frame, expand **File Migration**, expand **Find File Location**, and then click the **Virtual** link.

The **Virtual File Location** page appears.

2. Choose the name of the original virtual server from the **Virtual Server** drop-down menu.

For example, choose *virtualserver1*.

3. For **Share or Export Name**, enter the original virtual share or export.

For example, for a virtual NFS export, enter:

`/virtualexport/sw`

Or for a virtual CIFS share, enter the share name:

`virtualmktg`

4. For **Path**, enter the original path of the virtual file or directory as follows:

- For NFS, enter the path in the following format: `/<path>`
- For CIFS, enter the path in the following format: `\<path>`

5. Click **Find**.

The File Director displays the current location of the file or directory. If the file or directory has not been migrated, the existing location is returned.

CHAPTER 5

Using Policies for Automatic File Migration

This chapter contains the following topics that explain how to create and use a policy to automate file migration:

- About Policies
- Using and Creating Schedules
- Monitoring Traffic for Source Shares or Exports
- Working with Policy Rules
- Working with Policies

About Policies

By using a policy, you can automate the process of file migration from one file server to another based on the file attributes, file access information, and schedule that you specify. For example, you can create a policy that instructs File Director to execute every Saturday night at 2:00 AM to migrate MPEG files greater than 100 MB and have been accessed in the last week.

You can also preview the effects of a policy by creating a report of files that meet the policy specifications, without actually migrating any files. This policy preview process is called a *rehearsal*. By running a policy rehearsal, you can adjust policy conditions if necessary before running the actual policy.

Summary of Steps for Setting Up and Using a Policy

Here are the general steps for setting up and using a policy:

1. (Optional) The File Director includes four system-defined schedules (hourly, daily, weekly, and monthly) used for scheduling re-synchronization of monitored traffic with backend servers and for scheduling policy execution. If necessary, you can also create custom schedules. See “Using and Creating Schedules” on page 5-4.
2. Set up monitoring of client traffic that the File Director processes for the source shares or exports that have been migrated. See “Monitoring Traffic for Source Shares or Exports” on page 5-6.
3. Create a policy rule that contains one or more conditions for selecting the files you want to migrate. See “Working with Policy Rules” on page 5-7.

You can select the files based on various file attributes, such as file size or type, and on dynamic file access characteristics, such as the 50 most frequently accessed files within the last 6 months.

4. Create, rehearse, and run a policy that uses the policy rule. You can schedule a policy by using a system-defined or custom schedule. The policy is run automatically when the system time corresponds to the schedule you specified. For example, a policy scheduled as *hourly* runs at the first minute of every hour.

The policy rule specifies what to migrate and the schedule specifies when to migrate. See “Working with Policies” on page 5-18.

Each policy specifies the following information:

- Source file server and share or export that contains the files you want to migrate
- Name of policy rule you want to use to select the files
- Destination file server and share or export where you want to migrate the files
- Schedule when you want to run the policy and perform the migration

Immediately after you set up monitoring of a source share or export, the File Director begins storing file attribute and access information as it processes client traffic to that share or export. The File Director uses the stored file attribute and access information to determine the files that are qualified by the policy for migration.

You can specify the following attributes when creating conditions in a policy rule for selecting files:

Attribute Category	File Attribute	Valid File Attribute Values	Example of Using Attributes in a Policy Condition
Time	Files last accessed time	hours, days, months	Files last accessed time was more than 9 days ago
	Files last modified time	hours, days, months	Files last modified time was more than 24 hours ago
Size	File size	Bytes, Kilobytes, Megabytes, Gigabytes	<ul style="list-style-type: none">• File size is at least 10 Megabytes• File size is more than 200 Megabytes

Attribute Category	File Attribute	Valid File Attribute Values	Example of Using Attributes in a Policy Condition
Type	File type	AI, AIF, AIFC, AIFF, ASC, AU, AVI, BCPIO, BIN, BMP, CDF, CLASS, CPIO, CPT, CSH, CSS, DCR, DIR, DLL, DMS, DOC, DVI, DXR, EPS, ETX, EXE, EZ, GIF, GTAR, HDF, HQX, HTM, HTML, ICE, IEF, IGES, IGS, JPE, JPEG, JPG, JS, KAR, LATEX, LHA, LZH, M3U, MAN, ME, MESH, MID, MIDI, MIF, MOV, MOVIE, MP2, MP3, MPE, MPEG, MPG, MPGA, MS, MSH, MXU, NC, ODA, PBM, PDB, PDF, PGM, PGN, PNG, PNM, PPM, PPT, PS, QT, RA, RAM, RAS, RGB, RM, ROFF, RPM, RTF, RTX, SGM, SGML, SH, SHAR, SILO, SIT, SKD, SKM, SKP, SKT, SMI, SMIL, SND, SO, SPL, SRC, SV4CPIO, SV4CRC, SWF, T, TAR, TCL, TEX, TEXI, TEXINFO, TIF, TIFF, TR, TSV, TXT, USTAR, VCD, VRML, WAV, WBMP, WBXML, WML, WMLC, WMLS, WMLSC, WRL, XBM, XLS, XML, XPM, XSL, XWD, XYZ, ZIP	File type is one of MPEG, MP3
Access frequency	Files are accessed	Any valid aggregate value by count or by percent. Valid values are expressed as numbers or percentages.	Files are among the 90 percent most frequently accessed within the time period beginning 10 days ago

Policy Execution

Once a policy has executed and submitted a job to the migration server, the job is added to the pending queue. If the migration server is not processing any other migrations at the time and there are no other jobs on the pending queue, the newly-submitted job executes right away. Otherwise, the job remains on the pending queue and executes in turn. The schedule specified in a policy determines when the policy runs and submits the job. The schedule does not determine when the migration executes. For example, if the migration server is currently processing a large export when the new job is submitted, it may take some time before any new jobs execute.

Using and Creating Schedules

The File Director includes several system-defined schedules you can use for scheduling re-synchronization of monitored traffic with backend servers and for scheduling policy execution.

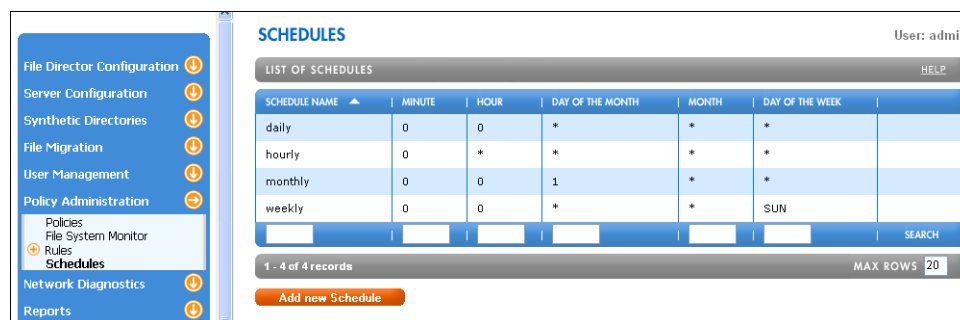
The following are the names of the system-defined schedules which you cannot change or delete:

This system-defined schedule	Executes on
hourly	the first minute of every hour
daily	every day at midnight
weekly	every Sunday at midnight
monthly	every first day of the month at midnight

To create a custom schedule:

1. In the Navigation frame, expand **Policy Administration** and then click the **Schedules** link.

The **Schedules** page includes a list of the system-defined schedules.



2. Click the **Add New Schedule**.
3. For **Schedule Name**, enter a name to identify the schedule.

For example, enter *NoonSchedule*.

4. For **Minute**, **Hour**, **Day of the Month**, **Month**, and **Day of the Week**, enter any of the following values:

- A single number (such as 5)
- A range (such as 0-30)
- A range with intervals, such as every second day or every third day. Use the format: **range/interval**.

For example, enter 1-9/2 to specify the following values: 1, 3, 5, 7, 9

- A comma-separated list of ranges and values (such as 0,5,10-20). The list can also include a **range/interval**.

- A * character (asterisk) to specify all possible values. For example, a * character for **Day of the Month** specifies 1-31.

Note: You must enter a value in each of the fields.

Here is a summary of the values allowed for each schedule parameter:

For this parameter	The allowed values are
Minute	0-59, *
Hour	0-23, *
Day of the Month	1-31, *
Month	1-12, Jan, Feb, Mar, Apr May, Jun, Jul, Aug, Sep, Oct, Nov Dec, *
Day of the Week	0-7, Sun, Mon, Tue, Wed, Thu, Fri, Sat, * Note: Both 0 and 7 specify Sunday for the Day of the Week.

Note: The day in a schedule can be specified by either Day of the Month or Day of the Week. If both parameters are specified by a value other than *, the schedule will run when either parameter matches the current time.

For example, to specify a schedule of every Monday at midnight, enter the following values:

- 0 (zero) for **Minute**
- 0 (zero) for **Hour**
- * for **Day of the Month**
- * for **Month**
- **Mon** for **Day of the Week**

Note: If you specify a day of the month that doesn't occur in a particular month, the policy action does not occur. For example, if you specify 31 as the day of the month, the policy action does not occur in months that have less than 31 days.

5. Click **Submit**.

Monitoring Traffic for Source Shares or Exports

Before creating a policy for migration, you must set up monitoring of client traffic that the File Director processes for the source shares or exports that you want to migrate. The File Director constantly monitors traffic through the File Director to that share or export and stores file attribute and access information. The File Director uses the stored file attribute and access information to determine the files that are qualified by the policy for migration.

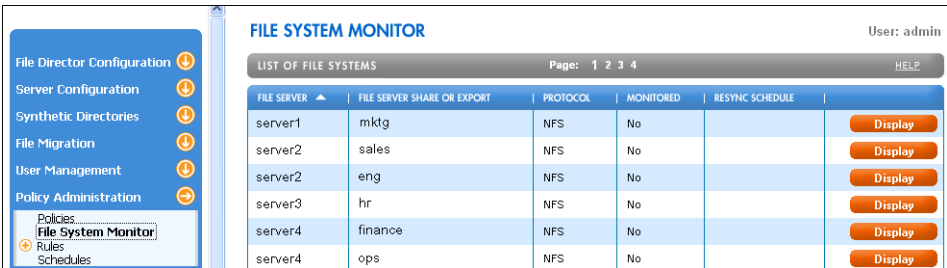
It's possible that during periods of heavy usage, the File Director may be temporarily unable to monitor all traffic. You can have the File Director automatically re-synchronize its stored attribute and access information with the current state of the files in the source shares or exports based on a system-defined or custom schedule.

Note: You must set up monitoring of a source share or export before you can create a policy that migrates that share or export.

To monitor traffic for source shares or exports:

- 1. In the Navigation frame, expand **Policy Administration** and then click the **File System Monitor** link.

The **File System Monitor** page appears with the **List of File Systems**.



FILE SYSTEM MONITOR					
LIST OF FILE SYSTEMS		Page: 1 2 3 4		HELP	
FILE SERVER	FILE SERVER SHARE OR EXPORT	PROTOCOL	MONITORED	RESYNC SCHEDULE	
server1	mktg	NFS	No		Display
server2	sales	NFS	No		Display
server2	eng	NFS	No		Display
server3	hr	NFS	No		Display
server4	finance	NFS	No		Display
server4	ops	NFS	No		Display

- 2. On the **File System Monitor** page, click the **Display** link next to the share or export that you want to monitor for migration with a policy.

The **File System Monitor** page appears showing information about the share or export.

- 3. Click **Edit**.
- 4. For **Monitored**, choose **Yes**.

5. For **Resync Schedule**, choose a schedule that specifies when you want to have the File Director automatically re-synchronize its stored attribute and access information with the current state of the files in the source shares or exports.

By default, the **Resync Schedule** is weekly.

Note: Keep the following guidelines in mind when setting the **Resync Schedule**:

- If the File Director is used heavily, set the **Resync Schedule** to about every two days. Otherwise, a **Resync Schedule** of one week is sufficient. Do not set the **Resync Schedule** to hourly.
- If the source share or export is large, re-synchronization takes more time and you should avoid choosing a frequent schedule.
- Re-synchronization generates extra load on the file server. For file servers that are used heavily, choose a schedule that won't seriously degrade performance.

6. Click **Submit**.

Working with Policy Rules

A policy rule contains the search conditions that select the files you want to migrate based on file attributes and access characteristics. Here is an example of a condition:

File Type is BMP

You can combine one or more conditions in a policy rule by using the Boolean operators AND and OR. For example, the following set of conditions use two AND operators and an OR operator to search for all BMP files larger than 100 MB that either: a.) have been accessed in the last 7 days, or b.) are among the 50 most frequently accessed files in the last 6 months:

File Type is BMP

AND

File Size is more than 100 Megabytes

AND

(Files Last Access Time was less than 7 days ago

OR

Files are among the 50 most frequently accessed within the last 6 months.)

Because a policy can contain only one policy rule, you must combine multiple search conditions in a single rule.

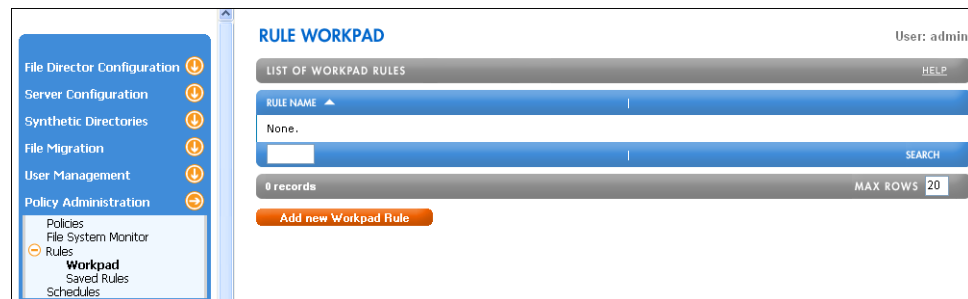
You begin creating a rule on the Workpad, which is a storage place for creating and editing a rule. Rules that you store on the Workpad remain stored after you close your web browser so you can edit them at any time. However, to use a Workpad rule within a policy, you must explicitly save the rule.

Creating a Policy Rule

To create a policy rule:

1. In the Navigation frame, expand **Policy Administration**, expand **Rules**, and then click the **Workpad** link.

The **List of Workpad Rules** appears in the **Rule Workpad** page.



2. Click **Add New Workpad Rule**.

The **Create Workpad Rule** section appears in the **Rule Workpad** page.

3. Enter the name of the rule and then click **Submit**.

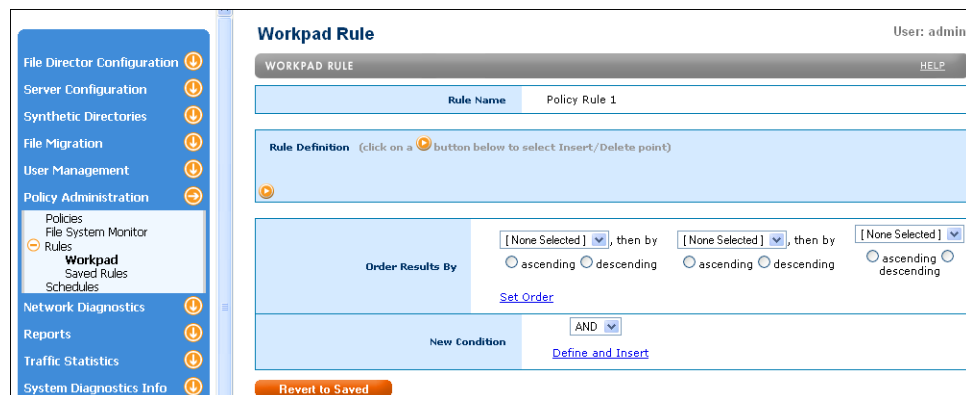
The rule appears in the **List of Workpad Rules** section in the **Rule Workpad** page.

4. Click **Display** next to the rule whose condition(s) you want to specify.

The **Workpad Rule** page appears. This page contains the definition (list of conditions) for the rule.

5. To specify conditions for the rule, click **Edit**.

A new set of options appear for inserting a condition into the rule's definition.



Inserting a Condition into a Rule

After creating a rule on the Workpad, you can insert conditions into the rule.

To insert a condition into a rule:

1. Click the location in the rule definition where you want to insert the condition.

Click the Rule Definition area to select the location where you want to insert a condition.

2. To specify the order you want the files to be selected, choose options in the **Order Results By** section.

For example, choose **File Size** from the left-most drop-down menu, and then select the **Ascending** option below that menu.

3. To define and insert a new condition, click the **Define and Insert** link in the **New Condition** section.

The **Policy Condition** page appears.

- For any *one* of the attribute categories, specify the file attribute values you want to include in the condition by choosing a value from a drop-down menu or list box, or by entering a text value.

For example, to specify a time-based condition, select **Accessed** or **Modified** from the **Last** drop-down menu in the **Time Based** section. Then, choose **Less Than** or **More Than** from the **Time Was** drop-down menu, enter a time value in the text box, and choose a unit of time from the drop-down menu.

Or, to specify a file-type based condition, double-click a file type from the drop-down menus in the **File Type Based** section. You can also enter file types directly in the **Comma-separated List** text box.

An example of a file type condition is:

File Type is one of MPEG, MP3

- Click **Insert** under the section where you specified a condition.

The File Director inserts the condition into the **Rule Definition** section on the **Workpad Rule** page.

New condition —

The screenshot shows the 'Workpad Rule' interface. At the top, it says 'User: admin'. Below that is a 'WORKPAD RULE' header with a 'HELP' link. A 'Rule Name' field contains 'Policy Rule 1'. The 'Rule Definition' section has a sub-header '(click on a [plus] button below to select Insert/Delete point)'. Below this, the condition 'File Type is one of MPEG, MP3' is displayed with a plus icon to its left. Below the condition is the 'Order Results By' section, which includes three columns for 'File Size', '[None Selected]', and '[None Selected]', each with 'ascending' and 'descending' radio buttons and a 'Set Order' link. At the bottom, there is a 'New Condition' section with an 'AND' dropdown and a 'Define and Insert' link. At the very bottom are three buttons: 'Revert to Saved', 'Delete Condition', and 'Commit to Saved Rules'.

Saving a Policy Rule

To use a rule that you have stored on the Workpad in a policy, you must save the rule.

To save a policy rule:

- In the **Workpad Rule** page, click **Commit to Saved Rules**.

Reverting a Workpad Rule to a Saved Rule

After you save a rule and then make changes to the rule on the Workpad, the saved rule and the rule on the Workpad will be different. The differences will remain until you save the rule again. You can also choose to discard the unsaved changes to the rule on the Workpad by reverting back to the saved version. After you revert, the saved rule and the rule on the Workpad will be the same.

To revert to a saved rule:

- In the **Workpad Rule** page, click **Revert to Saved**.

Combining Conditions in a Rule

You can combine multiple conditions in a rule using the Boolean operators AND and OR. For example, you can combine two conditions that select all MPEG files larger than 100 MB.

To combine conditions in a rule:

1. On the **Workpad Rule** page, click the location in the **Rule Definition** section where you want to insert another condition.

A border appears around the selected section.

The screenshot shows the 'Workpad Rule' configuration page. The 'Rule Definition' section is highlighted with a blue border. A condition 'File Type is one of MPEG, MP3' is selected, indicated by a blue border around it. To the left of the condition is a 'Boolean Operator' drop-down menu. The 'New Condition' section at the bottom shows the 'AND' operator selected. The page includes buttons for 'Revert to Saved', 'Delete Condition', and 'Commit to Saved Rules'.

Click the Rule Definition section to select the location where you want to insert another condition.

Boolean Operator drop-down menu

Border around selected section

2. Choose an operator from the Boolean Operator drop-down menu in the **New Condition** section.

For example, choose **AND**.

3. Click the **Define and Insert** link.
4. On the **Policy Condition** page, specify the file attribute values you want to include in the next condition.

An example of a file size condition is:

File Size is more than 100 Megabytes

- Click **Insert** under the section where you specified the condition.

The File Director inserts the Boolean operator and the additional condition into the **Rule Definition** section on the **Workpad Rule** page.

Workpad Rule User: admin

WORKPAD RULE HELP

Rule Name Policy Rule 1

Rule Definition (click on a button below to select Insert/Delete point)

File Type is one of MPEG, MP3

AND

File Size is more than 100 MB

Order Results By

File Size, then by [None Selected], then by [None Selected]

ascending descending ascending descending ascending descending

Set Order

New Condition AND Define and Insert

Revert to Saved Delete Condition Commit to Saved Rules

Nesting Conditions in a Rule

You can nest multiple conditions in a rule to specify the set of files you want to include and exclude from the rule. For example, the following procedure illustrates how to use an AND operator with the nested OR conditions to combine several conditions to select all MPEG files larger than 100 MB that either: a.) have been accessed in the last 7 days, or b.) are among the 50 most frequently accessed files in the last 6 months.

To nest conditions in a rule:

- On the **Workpad Rule** page, click the arrow for the condition where you want to insert nested conditions.

Workpad Rule User: admin

WORKPAD RULE HELP

Rule Name Policy Rule 1

Rule Definition (click on a button below to select Insert/Delete point)

File Type is one of MPEG, MP3

AND

File Size is more than 100 MB

Order Results By

File Size, then by [None Selected], then by [None Selected]

ascending descending ascending descending ascending descending

Set Order

New Condition AND Define and Insert

Revert to Saved Delete Condition Commit to Saved Rules

To nest conditions, click the arrow for the condition where you want to include the nested conditions.

Border around selected section

2. Choose an operator from the Boolean Operator drop-down menu in the **New Condition** section.

For example, choose **AND** to add the nested OR conditions as a set.

3. Click the **Define and Insert** link.

4. On the **Policy Condition** page, specify the file attribute values you want to include in the next condition.

An example of a time-based condition is:

Files Last Access Time was less than 7 days ago

5. Click **Insert** under the section where you specified a condition.

The File Director inserts the Boolean operator and the first condition of the nested set into the **Rule Definition** section on the **Workpad Rule** page.

For example:

The screenshot shows the 'Workpad Rule' configuration page for 'Policy Rule 1'. The 'Rule Definition' section contains three conditions connected by 'AND' operators:

- File Type is one of MPEG, MP3
- AND File Size is more than 100 MB
- AND Last Access Time was less than 7 days ago

Below the conditions, the 'Order Results By' section shows three criteria: 'File Size' (ascending), '[None Selected]' (ascending), and '[None Selected]' (ascending). At the bottom, the 'New Condition' section shows the 'AND' operator selected, with a 'Define and Insert' link. At the very bottom are three buttons: 'Revert to Saved', 'Delete Condition', and 'Commit to Saved Rules'.

6. In the **Workpad Rule** page, click the arrow next to the first nested condition to insert the second nested condition below it.

Click the arrow for the first nested condition where you want to include the next nested condition.

Border around selected section

7. Choose an operator from the Boolean Operator drop-down menu in the **New Condition** section.

For example, choose **OR** to add the second nested OR condition in the set.

8. Click the **Define and Insert** link.

9. On the **Policy Condition** page, specify the values you want to include in the next condition.

For example:

Files are among the 50 most frequently accessed within the last 6 months.

10. Click **Insert** under the section where you specified a condition.

The File Director inserts the Boolean operator and the second condition of the nested set into the **Rule Definition** section on the **Workpad Rule** page.

For example, the following rule uses two AND conditions and a set of nested OR conditions to select all MPEG files larger than 100 MB that either: a.) have been accessed in the last 7 days, or b.) are among the 50 most frequently accessed files in the last 6 months.

Workpad Rule User: admin

WORKPAD RULE HELP

Rule Name Policy Rule 1

Rule Definition (click on a button below to select Insert/Delete point)

File Type is one of MPEG, MP3

AND

File Size is more than 100 MB

AND

Last Access Time was less than 7 days ago

OR

Access Frequency within the last 6 months is among the highest 50 by count

Order Results By

File Size , then by [None Selected] , then by [None Selected]

☒ ascending ☐ descending ☐ ascending ☐ descending ☐ ascending ☐ descending

[Set Order](#)

New Condition AND

[Define and Insert](#)

[Revert to Saved](#) [Delete Condition](#) [Commit to Saved Rules](#)

Deleting a Condition from a Rule

To delete a condition from a rule:

1. On the **Workpad Rule** page, select the individual condition or the set of conditions you want to delete.

Workpad Rule User: admin

WORKPAD RULE HELP

Rule Name Policy Rule 1

Rule Definition (click on a button below to select Insert/Delete point)

File Type can include MPEG

AND

File Size is more than 100 MB

AND

Last Access Time was less than 7 days ago

OR

Access Frequency within the last 6 months is among the highest 50 by count

Order Results By

File Size, then by [None Selected], then by [None Selected]

ascending descending ascending descending ascending descending

Set Order

New Condition AND

Define and Insert

Revert to Saved Delete Condition Commit to Saved Rules

To select a set of nested conditions, click the outer arrow for the set.

Border around selected section

2. Click **Delete Condition**.

Deleting a Saved Rule

To delete a saved rule:

1. In the Navigation frame, expand **Policy Administration**, expand **Rules**, and then click the **Saved Rules** link.

The **Policy Rules** page appears.

POLICY RULES User: admin

LIST OF POLICY RULES HELP

RULE NAME	
Policy Rule 1	Delete Display
Policy Rule 2	Delete Display

SEARCH

1 - 2 of 2 records MAX ROWS 20

Start Rule in Workpad

2. Click **Delete** next to the rule that you want to delete.
3. Click **Confirm Delete**.

Editing a Saved Rule

To edit a saved rule:

1. In the Navigation frame, expand **Policy Administration**, expand **Rules**, and then click the **Saved Rules** link.

The **Policy Rules** page appears.

2. Click **Display** next to the rule that you want to edit.
3. Click **Edit Workpad Version**.

The **Workpad Rule** page appears.

4. Do either of the following to make changes to the rule:
 - Insert a new rule. For more information, see “Inserting a Condition into a Rule” on page 5-9.
 - Delete a rule. For more information, see “Deleting a Condition from a Rule” on page 5-16.
5. To save the changes, click **Commit to Saved Rules**.

Working with Policies

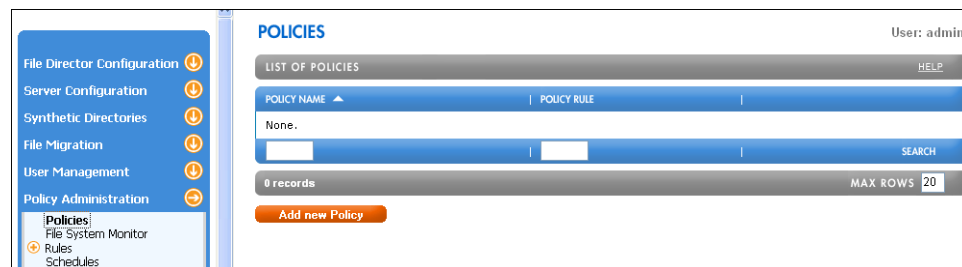
A policy can contain a maximum of one rule for selecting the files for migration. Before you can create the policy, you must create the rule and set up monitoring of traffic to and from the physical shares and exports that you want to migrate. For more information, see “Working with Policy Rules” on page 5-7.

Creating and Running Policy

To create and run policy:

1. In the Navigation frame, expand **Policy Administration** and then click the **Policies** link.

The **Policies** page appears.



2. Click **Add New Policy**.

The **Configure a Policy** page appears.

3. For **Policy Name**, enter a name to identify the policy.

Important: The policy name can contain letters, numbers, dashes (-), or underscores (_). Spaces are not allowed, and the name cannot begin with a dash.

For example, enter *Nightly-Migration*.

4. For **File Server**, enter the name of the source file server that contains the files you want to migrate.
5. For **File Server Share or Export**, enter the share or export where the source files are located.
6. From the **Rule** drop-down menu, choose the name of the saved policy rule you want to use to select the files for the policy.

Note: Only saved policy rules appear on the **Rule** drop-down menu. For more information, see “Saving a Policy Rule” on page 5-10.

7. For **Destination Server Name for Migration**, enter the server name of the destination file server where you want to migrate the files.
8. For **Destination Share or Export for Migration**, enter the name of the share or export where you want to migrate the files.

9. From the **Run Schedule** drop-down menu, choose a schedule that specifies when you want to run the policy and perform the migration.
For more information on schedules, see “Using and Creating Schedules” on page 5-4.
10. (Optional) For **Maximum Bytes for Migration**, enter the maximum number of bytes moved per migration for this policy.
11. (Optional) For **Maximum Files for Migration**, enter the maximum number of files moved per migration for this policy.
12. Click **Submit** to save and run the policy according to the specified schedule.
If all of the settings are correct, the policy is shown on the **Policies** page.

Rehearsing a Policy

You can preview the effects of a policy by creating a report of files that meet the policy specifications, without actually migrating any files. This policy preview process is called a *rehearsal*. By running a policy rehearsal, you can adjust policy conditions if necessary before running the actual policy according to the specified schedule in the policy.

To rehearse a policy:

1. In the Navigation frame, expand **Policy Administration** and then click the **Policies** link.
The **Policies** page appears.
2. Click **Display** next to the policy you want to rehearse.
The **Policies** page appears.
3. Click **Run Rehearsal** under the **Rehearsal Result Files** section.

The files that the policy selects appear in the **Rehearsal Result Files** section.

Note: The results of a rehearsal is dependent on the qualified files that exist when the rehearsal is run. Because the list of qualified files may be different when the policy is actually run, the results of a rehearsal may be different than the results of the policy.

For more accurate policy rehearsals, use the following guidelines:

- Monitor the share with a certain schedule.
- Create the policy with a **Disabled** schedule.
- Run the rehearsal when one of the following conditions is met:
 - There is at least some client traffic to the monitored share as a result of the File Director’s virtual share that exports the monitored share.
 - Share monitoring per schedule has executed at least once.

Editing a Policy

To edit a policy:

1. On the **Policies** page, click **Display** next to the policy you want to edit.
2. Click **Edit**.
3. Make the necessary changes and then click **Submit**.

Deleting a Policy

To delete a policy:

1. On the **Policies** page, click **Delete** next to the policy you want to delete.
2. Click **Confirm Delete**.

Displaying the Policy Execution Log

The policy execution log contains informational, warning, and error messages about the execution of policies.

To display the policy execution log:

1. In the Navigation frame, expand **Reports** and then click the **Policy Execution Log** link.

The **Display Policy Execution Log** page appears showing a log for the policies that have been executed since the last time the log was cleared.

2. To display a particular section of the log, enter a starting line number and the total number of lines you want to display, and then click **Show Policy Execution Log**.

CHAPTER 6

Monitoring and Troubleshooting the File Director

This chapter contains the following topics that explain how to monitor and troubleshoot the File Director:

- Configuring SNMP on the File Director
- Displaying Log Files for the File Director
- Displaying Traffic Statistics for the File Director
- Troubleshooting the File Director
- Restarting and Shutting Down the File Director
- Backing Up and Restoring Configuration Information for the File Director
- Updating and Rolling Back the Software on the File Director

Configuring SNMP on the File Director

You can use a Simple Network Management Protocol (SNMP) management station to monitor the File Director and obtain statistics and configuration information. The File Director supports SNMP version 1 and version 2.

The File Director Management Information Base (MIB) supports standard MIB-II and NeoPath Networks private MIB.

Note: You can download and save the MIB file from the File Director. For more information, see “Downloading and Saving the MIB File from the File Director” on page 6-5. You can also obtain the latest version of the MIB file from the Support area of the NeoPath Networks web site: www.neopathnetworks.com

There are two types of SNMP communities you can configure on the File Director:

- **Trap community**—groups trap destinations (management stations). (See the next section, “Configuring an SNMP Trap Community.”)
- **Access community**—controls access to MIBs and communications with management stations. (See “Configuring an SNMP Access Community” on page 6-4.)

Configuring an SNMP Trap Community

To configure an SNMP trap community:

1. In the Navigation frame, expand **File Director Configuration**, expand **System**, and then click the **SNMP** link.

The **SNMP Configuration** page appears.

The screenshot displays the 'SNMP CONFIGURATION' page. On the left is a navigation pane with 'File Director Configuration' expanded, showing 'System' and 'SNMP' selected. The main panel has a header 'SNMP CONFIGURATION' and 'User: admin'. It contains three main sections: 'MIB FILE' with a table showing 'neopath-v2.mib' and a 'Download' button; 'LIST OF TRAP COMMUNITIES' with a table containing 'TrapCommunity1' and checkboxes for 'v1' and 'v2', along with 'Display' and 'Delete' buttons; and 'LIST OF ACCESS COMMUNITIES' with a table containing 'private' and 'public' entries, each with checkboxes for 'Public' and 'Private' access types, and 'Delete' buttons. At the bottom of each list section is an 'Add new' button. The user is 'admin'.

2. Click **Add New Trap Community**.

The **Trap Community Configuration** page appears.

3. For **Community Name**, enter a trap community string.

This is the community string that management stations need to access File Director traps. Blank spaces are not allowed in the community string.

4. For **Trap Version**, select one or both SNMP versions to send to destinations in the trap community.
5. Click **Submit**.

The trap community is added to the list on the **SNMP Configuration** page.

6. In the **List of Trap Communities**, click the **Display** link next to the new trap community.

The **Trap Community Configuration** page appears.

7. Click **Add New Trap Destination**.
- The **Trap Destinations** page appears.
8. For **IP Address**, enter the IP address of the trap destination, which is typically the IP address of the management station.
 9. Click **Submit**.
 10. To add other trap destinations to the selected trap community, repeat steps 7 through 9.

About the SNMP Traps for the File Director

The File Director supports the standard MIB-II traps. The following SNMP traps are also defined for the File Director:

Type of Trap	Trap Name	Description
Hardware	npTrapHardwareDiskFailure	Disk failure
	npTrapHardwareNetworkCardFailure	Network card failure
	npTrapHardwareTemperatureHigh	Temperature is too high
Back End File Server	npTrapBackendServerAvailable	Back end file server is available
	npTrapBackendServerUnavailable	Back end file server is unavailable
Services	npTrapSvcMigrationSuccessful	Migration is successful
	npTrapSvcMigrationUnsuccessful	Migration failure
Resources	npTrapResourcesLocalFSFull	Local system is almost full (over 95% usage)
	npTrapResourcesDBFull	Database is almost full (over 95% usage)

Deleting a Trap Community

Note: When you delete a trap community, all trap destinations associated with that community are also deleted.

To delete a trap community:

1. In the **List of Trap Communities** on the **SNMP Configuration** page, click the **Delete** link next to the trap community you want to delete.
2. Click **Confirm Delete**.

Configuring an SNMP Access Community

To configure an SNMP access community:

1. In the Navigation frame, expand **File Director Configuration**, expand **System**, and then click the **SNMP** link.

The **SNMP Configuration** page appears.

2. Click **Add New Access Community**.

The **Access Community Configuration** page appears.

3. For **Community Name**, enter an access community string.

This is the access community string that management stations need to access the File Director MIB. Blank spaces are not allowed in the community string.

4. For **Access Types**, select **Public**, **Private**, or both options to specify the type of access to the File Director MIBs you want to allow the community to have.

The **Public** option allows access only to public (or standard) MIBs such as MIB-II. The **Private** option allows access to private MIBs which are under the company tree.

5. Click **Submit**.

The access community is added to the list on the **SNMP Configuration** page.

Deleting an Access Community

To delete an access community:

1. In the **List of Access Communities** on the **SNMP Configuration** page, click the **Delete** link next to the access community you want to delete.
2. Click **Confirm Delete**.

Downloading and Saving the MIB File from the File Director

To download and save the MIB file from the File Director:

1. In the Navigation frame, expand **File Director Configuration**, expand **System**, and then click the **SNMP** link.

The **SNMP Configuration** page appears.

The screenshot shows the 'SNMP CONFIGURATION' page for a user named 'admin'. On the left is a navigation pane with 'File Director Configuration' expanded, showing 'System' > 'SNMP'. The main content area has three sections:

- MIB FILE**: A table with one row:

FILE NAME	
neopath-v2.mib	Download

 A callout line points to the 'Download' link.
- LIST OF TRAP COMMUNITIES**: A table with one row:

COMMUNITY NAME	VERSIONS	
TrapCommunity1	<input checked="" type="checkbox"/> v1 <input checked="" type="checkbox"/> v2	Display Delete
- LIST OF ACCESS COMMUNITIES**: A table with two rows:

COMMUNITY NAME	ACCESS TYPES	
private	<input type="checkbox"/> Public <input checked="" type="checkbox"/> Private	Delete
public	<input checked="" type="checkbox"/> Public <input type="checkbox"/> Private	Delete

2. Click the **Download** link in the MIB file section.

The MIB file appears in a new browser window.

3. To save the MIB file, use the Save As command on the File menu of the browser.

Displaying Log Files for the File Director

You can check system status and troubleshoot the File Director by using the following log files:

- **System log**— contains informational, warning and error messages from standard system services and the operating system. (See the next section, “Displaying the System Log File.”)
- **Error log**—contains error messages from all NeoPath-specific services except for the core switching service. Note that some of the error messages it contains may be duplicates of errors logged in other log files. (See “Displaying the Error Log” on page 6-6.)
- **Accumulator Log**—contains debug information from the accumulator service. (See “Displaying the Accumulator Log” on page 6-7.)
- **Debug Log**—contains system debug information for the File Director Management Console. (See “Displaying the System Debug Log” on page 6-7.)
- **Switching Service log**—contains warning and error messages from the core switching service. (See “Displaying the Switching Service Log” on page 6-7.)
- **Migration Log**—contains informational, warning, and error messages from migration jobs. Use this log to debug any problems with migration. (See “Displaying the File Migration Log” on page 4-11.)
- **Policy Execution Log**—contains informational, warning, and error messages about the execution of policies. (See “Displaying the Policy Execution Log” on page 5-20.)

For information on configuring how the File Director rotates log files, see “Configuring Log File Rotation” on page 6-8.

Displaying the System Log File

You can troubleshoot File Director system problems or connections to file servers by using the system log file.

To display the system log file:

1. In the Navigation frame, expand **Reports** and then click the **System Log** link.
The **Display System Log** page appears.
2. To display a particular section of the log, enter a starting line number and the total number of lines you want to display, and then click **Show System Log**.

Displaying the Error Log

To display the error log file:

1. In the Navigation frame, expand **Reports** and then click the **Error Log** link.
The **Display Error Log** page appears.
2. To display a particular section of the log, enter a starting line number and the total number of lines you want to display, and then click **Show Error Log**.

Displaying the Accumulator Log

The Accumulator log file contains debugging output from the Accumulator service. For more information on the services, see “Restarting or Stopping System Services on the File Director” on page 6-13.

To display the Accumulator log file:

1. In the Navigation frame, expand **Reports** and then click the **Accumulator Log** link.

The **Display Accumulator Log** page appears.

2. To display a particular section of the log, enter a starting line number and the total number of lines you want to display, and then click **Show Accumulator Log**.

Displaying the System Debug Log

To display the system debug log:

1. In the Navigation frame, expand **System Diagnostics Info** and then click the **Debug Log** link.

The **Display Debug Log** page appears showing debug information collected since the last time the log was cleared.

2. To display a particular section of the log, enter a starting line number and the total number of lines you want to display, and then click **Show Debug Log**.
3. To clear the entire log, click **Clear Debug Log**.

Displaying the Switching Service Log

To display the switching service log file:

1. In the Navigation frame, expand **System Diagnostics Info** and then click the **Switching Service Log** link.

The **Display Switching Service Log** page appears.

2. To display a particular section of the log, enter a starting line number and the total number of lines you want to display, and then click **Show Switching Service Log**.

Configuring Log File Rotation

The File Director rotates system log files based on file size or rotation schedule. When rotation occurs, the File Director saves the current log as a rotated log file and then creates a new blank log. You can configure the following options for log rotation:

- Maximum number of rotated log files you want the File Director to keep
- Maximum size of the log file
- Rotation schedule

To configure log file rotation:

1. In the Navigation frame, expand **File Director Configuration** and then click the **Log File Rotation** link.

The **Log File Rotation** page appears.

LOG FILE ROTATION				
LIST OF LOG FILES				
LOG FILE NAME	MAXIMUM ROTATED FILES	MAXIMUM BYTE SIZE	ROTATION SCHEDULE	
accumulator	10	10 MB		Display
error	10	10 MB		Display
migration	10	10 MB		Display
policy_execution	10	10 MB		Display
switching	10	10 MB		Display
system	10	10 MB		Display

2. Click **Display** for the log file you want to configure, and then click **Edit**.
3. For **Maximum Rotated Files**, enter the maximum number of rotated log files you want the File Director to keep. The default is 10.

For example, enter 2 to have the File Director keep a maximum of two rotated log files. You can also enter 0 (zero) to have the File Director keep no rotated log files or **infinite** to keep all rotated logs.

Note: Selecting 0 (zero) or **infinite** may create logs that fill up available disk space.

4. Specify whether you want log rotation to be based on file size or age by doing one of the following, but not both:
 - For **Maximum Byte Size**, enter the maximum size of the log file in one of the following units: B, KB, MB, or GB. The default is 10 MB.
 - Choose a schedule from the **Rotation Schedule** drop-down menu.
5. Click **Submit**.

Displaying Traffic Statistics for the File Director

You can display the following traffic statistics information:

- NFS and NLM traffic statistics—See “Displaying NFS and NLM Traffic Statistics” on page 6-9.
- CIFS traffic statistics—See “Displaying CIFS Traffic Statistics” on page 6-10.
- Network Interface traffic statistics—See “Displaying Traffic Statistics for the File Director Network Interfaces” on page 6-10.

Displaying NFS and NLM Traffic Statistics

To display the NFS and NLM traffic statistics:

1. In the Navigation frame, expand **Traffic Statistics** and then click the **NFS Stats** link.

The **NFS Stats** page appears, which contains NFS server statistics, NFS client statistics, NFS procedure statistics, and NLM procedure statistics.

The screenshot shows the File Director interface with the 'NFS Stats' page selected. The left navigation pane lists various categories, with 'Traffic Statistics' expanded and 'NFS Stats' highlighted. The main content area displays two tables: 'NFS SERVER STATS' and 'NFS CLIENT STATS'.

NFS SERVER STATS

SERVER NAME	TOTAL BYTES RECEIVED	TOTAL BYTES SENT	
server1	0 (0 kB/s)	0 (0 kB/s)	Display

1 - 1 of 1 records MAX ROWS 20

NFS CLIENT STATS

	TOTAL OPS PER SEC	TOTAL BYTES RECEIVED	TOTAL BYTES SENT
TCP	-	0 (0 kB/s)	0 (0 kB/s)
UDP	-	0 (0 kB/s)	0 (0 kB/s)
Total	0	0 (0 kB/s)	0 (0 kB/s)

2. Scroll down to view all of the statistics.
3. For NFS statistics for a particular NFS file server, click the **Display** link next to the server.

Displaying CIFS Traffic Statistics

To display the CIFS traffic statistics:

1. In the Navigation frame, expand **Traffic Statistics** and then click the **CIFS Stats** link.

The **CIFS Stats** page appears, which contains CIFS client and server statistics.

CIFS STATS					User: admin
CIFS CLIENT STATS					HELP
NUMBER OF USERS	NUMBER OF CONNECTIONS	INCOMING	OUTGOING		
0	0	0 bytes/sec, 0 ops/sec	0 bytes/sec, 0 ops/sec		
CIFS SERVER STATS					HELP
SERVER NAME	NUMBER OF USERS	NUMBER OF CONNECTIONS	INCOMING	OUTGOING	
server2	0	0	0 bytes/sec, 0 ops/sec	0 bytes/sec, 0 ops/sec	Display
server3	0	0	0 bytes/sec, 0 ops/sec	0 bytes/sec, 0 ops/sec	Display
					SEARCH
1 - 2 of 2 records					MAX ROWS 20

2. For CIFS server and share statistics for a particular CIFS file server, click the **Display** link next to the server.

Displaying Traffic Statistics for the File Director Network Interfaces

To display traffic statistics for the File Director network interfaces:

- In the Navigation frame, expand **Traffic Statistics** and then click the **Interface Stats** link.

The **Interface Stats** page appears.

INTERFACE STATS		
INTERFACE NAME	BYTES RECEIVED	BYTES SENT
eth0	0 bytes/sec, 0 pkt/sec	0 bytes/sec, 0 pkt/sec
eth1	0 bytes/sec, 0 pkt/sec	0 bytes/sec, 0 pkt/sec
eth2	0 bytes/sec, 0 pkt/sec	0 bytes/sec, 0 pkt/sec
eth3	0 bytes/sec, 0 pkt/sec	0 bytes/sec, 0 pkt/sec

Tip: You can also display the status of the File Director network interfaces. For more information, see “Displaying the Status of the File Director Network Interfaces” on page 6-11.

Troubleshooting the File Director

You can use the following tools for troubleshooting the File Director:

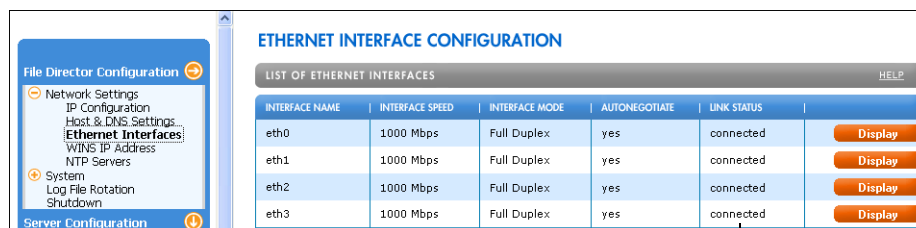
- Check the status of the File Director network interfaces. See “Displaying the Status of the File Director Network Interfaces,” next.
- Ping devices from the File Director to see if they are accessible. See “Performing a Ping Test from the File Director” on page 6-12.
- Check for routing problems. See “Displaying a Routing Table from the File Director” on page 6-12.
- Restart or stop services that are not operating. See “Restarting or Stopping System Services on the File Director” on page 6-13.
- Restart the File Director. See “Restarting and Shutting Down the File Director” on page 6-15.
- Check log files from the File Director. See “Displaying Log Files for the File Director” on page 6-6.

Displaying the Status of the File Director Network Interfaces

To display the status of the File Director network interfaces:

- In the Navigation frame, expand **File Director Configuration**, expand **Network Settings**, and then click the **Ethernet Interfaces** link.

The **Ethernet Interface Configuration** page appears.



ETHERNET INTERFACE CONFIGURATION					
LIST OF ETHERNET INTERFACES					
INTERFACE NAME	INTERFACE SPEED	INTERFACE MODE	AUTONEGOTIATE	LINK STATUS	
eth0	1000 Mbps	Full Duplex	yes	connected	Display
eth1	1000 Mbps	Full Duplex	yes	connected	Display
eth2	1000 Mbps	Full Duplex	yes	connected	Display
eth3	1000 Mbps	Full Duplex	yes	connected	Display

status of network interfaces

Note: To change the default configuration of the File Director network interfaces, click **Display**. For more information see “Configuring the File Director Network Interfaces” on page 2-8.

Tip: You can also display traffic statistics for the File Director network interfaces. For more information, see “Displaying Traffic Statistics for the File Director Network Interfaces” on page 6-10.

Performing a Ping Test from the File Director

You can ping any device (such as a file server) from the File Director to verify that it is available and accessible to the File Director over the network.

To perform a ping test from the File Director:

1. In the Navigation frame, expand **Network Diagnostics** and then click the **Ping Test** link.

The **Ping Test** page appears.

2. Enter the IP address or host name of the device you want to ping.
3. Enter the number of pings to send.
4. Click **Run**.

The File Director displays the results of the ping.

Displaying a Routing Table from the File Director

You can troubleshoot routing problems by using a routing table from the File Director.

To display a routing table from the File Director:

1. In the Navigation frame, expand **Network Diagnostics** and then click the **Routing** link.

The **Routing Table** page appears.

IP ADDRESS	NETMASK	INTERFACE	GATEWAY
0.0.0.0	0.0.0.0	eth0	172.22.1.1
172.22.1.0	255.255.255.0	eth0	0.0.0.0
172.22.83.0	255.255.255.0	eth0	0.0.0.0

2. To display the latest routing information, click **Refresh**.

Restarting or Stopping System Services on the File Director

If the File Director stops operating properly, you can restart, stop, or start individual processes that are running on the File Director. In most cases, you should restart, stop, or start processes only if instructed to do so by a NeoPath Technical Support representative. But, you may want to stop the *admin_ssh* or *admin_www* service for security reasons.

Warning: If you stop both the *admin_www* and *admin_ssh* services, you must use the CLI commands over a console to configure the File Director.

The services on the File Director are:

Service name	Operation	Description
FD_cifs	CIFS switching service	Enables clients to access CIFS virtual shares
FD_nfs	NFS switching service	Enables clients to access NFS virtual shares
accumulator	attribute accumulator	Tracks file attributes and file access patterns for policy-based migrations
admin_snmp	SNMP agent daemon	Allows SNMP clients to query the File Director
admin_ssh	SSH daemon	Allows SSH remote access to use the CLI commands. If you stop this service, future attempts to log in using SSH will be prevented, but existing SSH connections will remain until you log out. After that, you must use a console to use the CLI commands.
admin_www	File Director Management Console	Allows administrative access to the File Director via the File Director Management Console in a web browser. If you stop this service, you will have to use the CLI commands to configure the File Director.

To restart, stop or start a system service on the File Director:

1. In the Navigation frame, expand **Services** and then click the **Services** link.
The **Services** page appears showing the status of the File Director services.

SERVICES			
LIST OF SERVICES			HELP
SERVICE NAME	SERVICE DESCRIPTION	SERVICE STATUS	
FD_cifs	CIFS switching service	Running	Display
FD_nfs	NFS switching service	Running	Display
accumulator	Attribute Accumulator	Running	Display
admin_snmp	SNMP agent daemon	Running	Display
admin_ssh	Administration via SSH	Running	Display
admin_www	Administration via Web	Running	Display

1 - 6 of 6 records MAX ROWS 20

2. Click **Display** next to the service you want to stop or restart.
3. Click **Edit** and then choose **Restart** or **Stop** from the **New Service Operation** drop-down menu. For services that are already stopped, choose **Start** to start the service.
4. Click **Submit**.

Removing a Server

In rare cases, you can use the Management Console or the `no server` CLI command to remove a file server.

Important: Do not remove a server configured for the File Director unless directed by your customer support representative:

- Removing a server will cause the File Director to forget about all of the files that have been migrated away from that server. Even if the server is later reconfigured, files and directories that had previously been migrated away will show up as empty files and directories when accessed through File Director. Make sure you no longer need to use a server before removing it.
- Once a server is removed from the configuration, File Director can no longer access any files on that server, including files that have been migrated to the server from another location. Before removing a server, make sure that no files currently in use have been migrated to that server.


Restarting and Shutting Down the File Director

If the File Director stops operating, you can restart it. Before turning the power off or moving the File Director, be sure to shut it down using the instructions in this section.

To restart or shut down the File Director:

1. In the Navigation frame, expand **File Director Configuration** and then click the **Shutdown** link.

The **File Director Shutdown** page appears.



2. Do one of the following:
 - To restart the File Director, select **Reboot**.
 - To shutdown the File Director and turn the power off, select **Halt**.
3. Click **Shutdown**.

Important: Do not turn the File Director off by using the Power switch. To shut File Director down, always use the **Shutdown** command in the Management Console or the **Halt** command in the CLI.

Backing Up and Restoring Configuration Information for the File Director

You can backup or restore all of the configuration information for the File Director, including IP addresses, file server information, virtual server information, the list of monitored shares and exports, and synthetic and union directory information. Non-configuration items that are not backed up are log files and the file attribute and access information for policies that is accumulated by monitoring shares and exports.

The File Director will create a backup file containing information about the File Director configuration on a file server via NFS, CIFS, SCP or FTP.

Backing Up the File Director Configuration

It is a good practice to back up the File Director configuration in the following situations:

- After any configuration update that changes the set of known file servers or shares and exports. For example, after adding information about a new file server, back up the File Director configuration.
- Before every software upgrade in case there is a problem during the upgrade process.

- After every software upgrade to make sure that you have a recent backup file that is compatible with the new software version.

To backup the File Director configuration:

1. In the Navigation frame, expand **File Director Configuration**, expand **System**, and then click the **Backup** link.

The **System Backup** page appears.

The screenshot shows the 'SYSTEM BACKUP' web interface. On the left is a navigation pane with 'File Director Configuration' expanded, showing 'System' and 'Backup' (highlighted). The main area has a title bar 'SYSTEM BACKUP' with a 'HELP' link. Below are four steps: Step 1: 'Enter the scheme for file transfer' with a 'Scheme' dropdown set to 'FTP'; Step 2: 'Enter the userid and password, if any' with 'User ID' and 'Password' text boxes; Step 3: 'Enter the hostname, path, and encryption password for the backup file' with 'Host Name', 'Path', and 'Encryption Password' text boxes; Step 4: 'Submit the request or cancel and try again later.' with a 'Backup' button.

2. Choose a file transfer scheme from the **Scheme** drop-down menu for transferring the backup file.
3. (Optional) For **User ID** and **Password**, enter a user ID and password for logging in during the FTP or SCP transfer process, or for the file server.
4. For **Host Name**, enter the name of the host where you want to store the backup file.
5. For **Path**, enter the path (including filename) of the backup file.

Tip: Use a filename for the backup file that is meaningful and indicates the software version you are using.

6. For **Encryption Password**, enter a password for encrypting administrator passwords for CIFS file servers in the backup file.

Important: Do not lose the encryption password. To restore the CIFS file server passwords, you must specify this same encryption password when restoring the backup file. If you don't know the encryption password, you can restore the backup file without restoring the CIFS file server passwords. You must then manually specify the **Admin Password** for each CIFS file server. (See "Manually Specifying Server Passwords" on page 6-18.)

7. Click **Backup**.

Restoring the File Director Configuration

You will need to restore configuration information for the File Director in the event of a hardware failure that requires a hardware replacement.

Important: If you have made any changes to the file server access configuration on a File Director and you have not created a new backup file to store those changes, do not restore configuration from an older backup file. Doing so will “roll back” the File Director’s configuration and may lead to data loss. Instead, use the Management Console or CLI commands to manually make any necessary configuration changes. This situation does not apply to a hardware replacement where there is no existing file server access configuration to roll back and no data to lose on the new File Director.

Note: The version of the backup file must be compatible with the software version running on the File Director. If it’s incompatible, the File Director will display an error.

To restore the File Director configuration:

1. In the Navigation frame, expand **File Director Configuration**, expand **System**, and then click the **Restore** link.

The **System Restore** page appears.

The screenshot shows the 'SYSTEM RESTORE' interface. On the left is a navigation pane with 'File Director Configuration' expanded, showing 'System' and 'Restore' (highlighted). The main area has a title bar 'SYSTEM RESTORE' with a 'HELP' button. Below are four steps:

- Step 1:** 'Enter the scheme for file transfer'. A dropdown menu shows 'FTP' selected.
- Step 2:** 'Enter the userid and password, if any'. Fields for 'User ID' and 'Password' are present.
- Step 3:** 'Enter the hostname, path, and decryption password for the backup file'. Fields for 'Host Name', 'Path', and '*Decryption Password' are present. Below these is a radio button group for 'If Password Validation Fails' with options 'Let me try again' (selected) and 'Just clear encrypted data'. A 'Hint' link is also present.
- Step 4:** 'Submit the request or cancel and try again later.' An orange 'Restore' button is at the bottom.

2. Choose a file transfer scheme from the **Scheme** drop-down menu for restoring the backup file.
3. (Optional) For **User ID** and **Password**, enter a user ID and password for logging in during the FTP or SCP transfer process.
4. For **Host Name**, enter the name of the host where the backup file is stored.
5. For **Path**, enter the path (including filename) of the backup file.
6. For **Decryption Password**, enter the password specified during the backup process for encrypting administrator passwords for CIFS file servers in the backup file.

When you created the backup file, you specified an encryption password. (See the previous section, “Backing Up the File Director Configuration.”) To restore the CIFS file servers passwords, you must specify the same password.

Note: If the configuration does not include any CIFS servers, you will be prompted to enter an encryption password, but it will not matter if the password does not match.

7. For **If Password Validation Fails**, select one of the following options:
 - **Let Me Try Again**—Select this option to be allowed to enter the decryption password again in case you typed it incorrectly.
 - **Just Clear Encrypted Data**—Select this option if you don’t know the decryption password. You can restore the backup file, but not the CIFS file server passwords. After completing the restore, you must then manually specify the **Admin Password** for each CIFS file server. For more information, see “Manually Specifying Server Passwords” on page 6-18.
8. Click **Restore**.

The File Director restarts after the configuration information has been restored.

MANUALLY SPECIFYING SERVER PASSWORDS

If you lose the encryption password, you must then manually specify the **Admin Password** for each CIFS file server.

To use the Management Console to reset the password:

1. In the Navigation frame, expand **Server Configuration** then click the **File Servers** link.
2. Select a file server and click **Display**.
3. When the **File Server Configuration Page** appears, click **Edit**.
4. Enter a password and click **Submit**.

For more information, see “Specifying a File Server” on page 2-16.

To use the CLI to reset the password:

1. Run the following command:

```
edit server <name> properties cifs admin <admin-user>
```

where:

<name> is the name of the file server.

<login-name> is the login user name for an administrator account on the file server. A prompt appears asking for a password.

2. Enter a new password at the prompt.

Updating and Rolling Back the Software on the File Director

You can update the software that is installed on the File Director, or roll it back to a previous version.

Updating the Software on the File Director

The File Director can store up to four versions of installed software. One version is the software initially installed at the factory. Note that you will overwrite the factory-version of the software when you use the following procedure to update the software the fourth time. For information on using the `software factory_reset` CLI command to reinstall the factory version of the software, see the *File Director CLI Reference Guide*.


Note: If you are updating the software on the nodes of a cluster, it is recommended (but not required) that you start updating the Primary node first. But, before starting the update process on the other node, be sure to completely finish the update process on one node and wait for the updated node to come up again and join the cluster.

Important: During software installation, the File Director will stop providing service to clients. The installation process will take a few minutes and will be followed by a system restart.

To update the software on the File Director:

1. In the Navigation frame, expand **File Director Configuration**, expand **System**, and then click the **Software** link.

The **Software Information** page appears showing version information about the currently installed File Director software.



2. Click **Install New Version**.

The **Software Information** page appears.

3. Choose a file transfer scheme from the **Scheme** drop-down menu for installing the software.
4. For **User ID** and **Password**, enter a user name and password.

The following notes apply to the interaction between the scheme you chose and the user name and password:

- For the SCP and FTP schemes, the user name defaults to “root”. If you don’t specify a password, a password prompt will be displayed during installation.

- For the NFS scheme, the user name and password are ignored and the access is performed as “root”.
 - For the CIFS scheme, the user name defaults to “Administrator”. If you don’t specify a password, a password prompt will be displayed during installation.
5. For **Host Name**, enter the name of the host where the software you want to install is located.
 6. For **Path Name**, enter the path on the host where the software is located.
 7. Click **Submit**.

Rolling Back the Software on the File Director

You can roll back the software on the File Director to the previously installed software version. The rollback is decremental. That is, if you do two software updates, you can do two software rollbacks. The File Director restarts after the rollback operation is complete.

Important: Rolling back the software also rolls back the configuration information to its previous state that existed when the software was previously updated.

To roll back the software on the File Director:

1. In the Navigation frame, expand **File Director Configuration**, expand **System**, and then click the **Software** link.

The **Software Information** page appears.

2. Click **Rollback to Previous Version**.

The **Software Information** page appears that shows information about the previous version.

3. Click **Confirm Rollback**.

CHAPTER 7

Setting Up a High Availability Cluster of File Directors

This chapter contains the following topics that explain how to configure a cluster of File Directors:

- Overview of High Availability Clusters
- Summary of Steps for Setting Up a Cluster
- Scenario A: Configuring Two New, Unconfigured File Directors as a Cluster
- Scenario B: Configuring an Existing File Director into a Cluster with Another New or Existing File Director
- Checking the Status of Each Cluster Node
- Troubleshooting Problems with Setting Up a Cluster
- Disbanding a Cluster
- Changing the Multicast Base Address

Overview of High Availability Clusters

You can provide clients with uninterrupted access to virtual shares and exports in case a File Director failure occurs by setting up a high availability cluster of File Directors. You can have a maximum of two File Director nodes in a cluster.

Primary and Backup Nodes

When you initially set up a cluster, you configure each File Director node with one of the following roles:

- Primary node—This is the node that you designate to provide services to clients.
- Backup node—This node provides services to clients only if the primary node is not operating.

Each cluster contains a Primary node and a Backup node.

Active and Standby Status

After you have finished configuring the cluster and it is operating normally, the status of the nodes are:

- Active—This node actively provides services to clients.
- Standby—This node *does not* actively provide services to clients, but it is ready to begin providing services if the Active node fails.

When both nodes are operating normally, the Primary node should be Active and the Backup node should be Standby. The Backup node becomes Active only when the Primary node is unavailable.

You configure the Active node with all of the necessary configuration information that specifies access to your file servers, virtual servers, synthetic directories, and so on. The Standby node stores a mirror of the Active node's configuration information. If you make any changes to the Active node's configuration, the Active node automatically copies those changes to the Standby node.

Heartbeat Communications and Transfer of Service

The two nodes use heartbeat messages to constantly monitor each other's availability and operating status. If the Active node fails to respond to the heartbeat messages, the Standby node automatically takes over Cluster IP addresses and continues providing services to clients.

Note: During the transfer of service from the Active node to the Standby node, there may be a brief service interruption. CIFS clients may need to reconnect to the virtual servers on the File Director.

After transfer of service, the roles of the two nodes do not change; that is, once you configure a node with the Primary role and the other node with the Backup role, they retain those roles regardless of which node is actively providing services to clients.

Configuration Changes

When you set up a cluster and the two nodes begin operating, you must make all cluster-wide configuration changes on the Active node. If the Active node fails, then you must make all cluster-wide configuration changes on the other node that is now Active.

When the Primary node recovers from a failure and begins operating, it will automatically become the Active node again when it begins operating. At that time, the Backup node will revert back to the Standby node. Any configuration changes you made on the Backup node while it was Active are automatically copied to the Primary node before it becomes Active.

Note: A CIFS migration job can be run only when both nodes are up, and it will fail if one of the nodes is dead. Policies are run only when the Primary node is the Active node. Do not define migration jobs or policies on the Backup node.

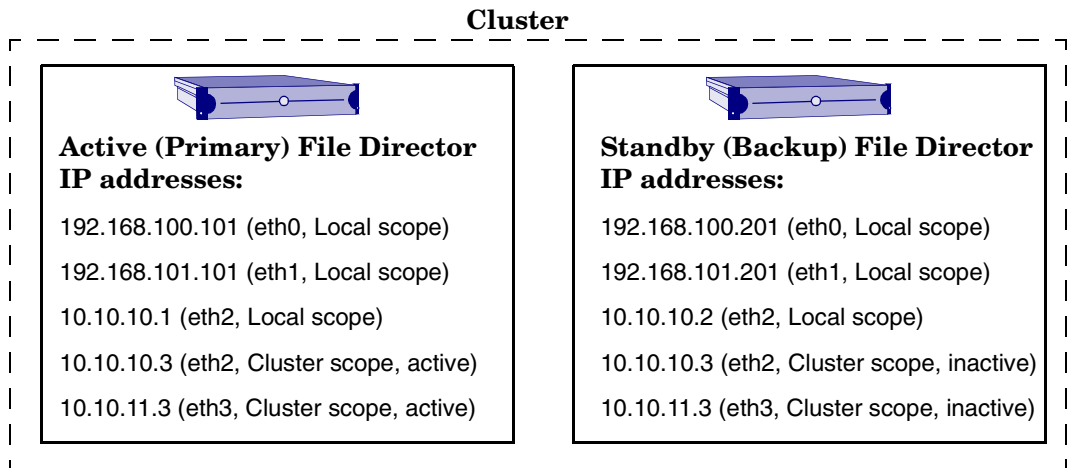
Setting Local or Cluster Scope

One of the steps you must do when configuring a cluster is to set a scope for each IP address on each node. The scope determines whether the Standby node takes over an IP address from a failed Active node. You can set either of the following scopes:

- **Local**—You set Local scope for all IP addresses that you don't want taken over by the Standby node if the Active node fails. If the Active node fails, all Local IP addresses become unavailable to clients.
- **Cluster**—You set Cluster scope for all IP addresses that are used for virtual servers, proxy IP addresses, and for access to file servers. If the Active node fails, the Standby node automatically takes over the Cluster IP addresses to continue providing file access services to clients.

In a typical cluster configuration, some IP addresses on the Active node must have Local scope and some have Cluster scope. Local IP addresses are used for heartbeat communications between the nodes. Cluster IP addresses are used for virtual servers, proxy IP addresses, and for access to file servers.

For example, if the Active node in the following illustration fails, then the Standby node takes over the IP addresses 10.10.10.3 and 10.10.11.3 because they have Cluster scope.



Using Local and Cluster IP address on the Active (Primary) and Standby (Backup) nodes

You assign Cluster IP addresses on the Active node *only*, where they remain in an Active state as long as the Active node is operating. The Active node automatically copies the list of Cluster IP addresses to the Standby node, where they remain in an Inactive state unless the Active node stops operating.

Important: You do not assign any Cluster scope IP addresses on the Standby node.

You can define Local scope IP addresses on the Standby node. Because Local scope IP addresses are not part of the cluster configuration, these changes do not need to be made on the Active node.


Both the Primary and Backup node must have at least one Local scope IP address defined. If one of the nodes does not have any Local scope IP addresses, then it will lose connectivity when it is in Standby mode — making it unable to exchange heartbeat messages with the other node.

Configuration Guidelines for a Cluster


The following are some guidelines for setting up clusters:

- You must set up a cluster as a pair of Primary and Backup nodes. For example, you cannot define one cluster to contain two Primary nodes and another cluster to contain the Backup nodes.

Allowed




Not Allowed




- You cannot configure more than two nodes in a cluster.
- You cannot configure two separate primary nodes to use the same backup node.

Not Allowed



Not Allowed



- You can install two new File Directors to create a cluster. See “Scenario A: Configuring Two New, Unconfigured File Directors as a Cluster” on page 7-9.
- You can combine two existing, stand-alone File Directors to create a cluster. See “Scenario B: Configuring an Existing File Director into a Cluster with Another New or Existing File Director” on page 7-17.
- Both nodes in a cluster should run the exact same version of File Director software.

Summary of Configuration Settings

Important: When you set up a cluster, you must configure all items related to client access (such as file servers, virtual servers, and synthetic and union directories) on the Active node *only*. But you must configure items such as DNS and gateway settings on *both* nodes.

Here is a summary of what you must configure on the Active node *only* or on the Active Primary node *only*, and other items that you must configure on *both* nodes:

On this node	Configure these items	For more information, see
Active Only	• IP addresses that have Cluster scope	“Configuring the File Director IP Addresses” on page 2-6
	• Proxy IP address	“Setting a Proxy IP Address for the File Director” on page 2-9
	• File servers	“Setting Up the File Director to Access File Servers” on page 2-13
	• Virtual servers, shares, and exports	“Specifying Virtual Servers and Virtual Shares” on page 2-20
	• Synthetic directories and union directories	Chapter 3, “Working with Synthetic Directories”
	• SNMP information	“Configuring SNMP on the File Director” on page 6-1
Active Primary Only	• Migration jobs ^a	Chapter 4, “Using the File Director to Migrate Data”
	• Policies ^b	Chapter 5, “Using Policies for Automatic File Migration”
Both Active and Standby	• IP addresses that have Local scope	“Configuring the File Director IP Addresses” on page 2-6
	• User Accounts	“Adding Administrator Accounts” on page 2-4
	• Default gateway	“Setting the Default Gateway for the File Director” on page 2-8
	• DNS settings	“Specifying a DNS Server and DNS Search Domain” on page 2-10
	• WINS server	“Specifying a WINS Server” on page 2-10
	• NTP server	“Specifying an NTP Server” on page 2-12
	• Log Rotation	“Displaying the File Migration Log” on page 4-11

a. A CIFS migration job can be run only if both nodes are up, and it will fail if one of the nodes is dead. Define migration jobs only if both nodes are up. Do not define migration jobs on the Backup node.

b. Policies are run only if the Primary node is the Active node. Define policies only if both nodes are up. Do not define policies on the Backup node.

Note: Although you can view cluster-wide configuration information on the Standby node, you cannot change the cluster-wide configuration settings on the Standby node.

Summary of Steps for Setting Up a Cluster

This section contains a summary of the steps for setting up a cluster. The details for these steps are described in the following sections:

- Scenario A: Configuring Two New, Unconfigured File Directors as a Cluster
- Scenario B: Configuring an Existing File Director into a Cluster with Another New or Existing File Director

Important: If you are combining two stand-alone File Directors into a cluster, make a backup of both configurations first before making any changes.

Scenario A: Combining Two New, Unconfigured File Directors into a Cluster	Scenario B: Combining an Existing, Stand-Alone File Director into a Cluster with Another New or Existing File Director
<p>Task A-1: Run the Initial Configuration Script.</p> <p>Perform this task on both nodes.</p>	<p>Task B-1: Verify Stand-Alone Configuration.</p> <p>Make sure both File Directors are operating as stand-alone nodes.</p> <ul style="list-style-type: none"> ❑ If one of the File Directors is new, run the initial configuration script on that node. ❑ On the node with the configuration you want to keep, verify the following: <ul style="list-style-type: none"> – Proxy IP address – File servers – Virtual servers, shares, and exports – Synthetic directories and union directories – SNMP information ❑ Make a backup of both configurations.
<p>Task A-2: Connect the Heartbeat Interfaces.</p> <p>Perform this task on both nodes.</p> <ul style="list-style-type: none"> ❑ For redundancy connect at least two interfaces between the nodes for heartbeat communications: <ul style="list-style-type: none"> - Create a direct connection between the nodes. - Create a network connection between the nodes. 	<p>Task B-2: Connect the Heartbeat Interfaces.</p> <p>Perform this task on both nodes.</p> <ul style="list-style-type: none"> ❑ For redundancy connect at least two interfaces between the nodes for heartbeat communications: <ul style="list-style-type: none"> - Create a direct connection between the nodes. - Create a network connection between the nodes.
<p>Task A-3: Configure IP Addresses on the Primary.</p> <p>Perform this task on the Primary node only.</p> <ul style="list-style-type: none"> ❑ Configure Local IP addresses for the heartbeat interfaces. ❑ Configure Cluster scope for IP addresses associated with virtual servers and proxy IP addresses. 	<p>Task B-3: Configure IP Addresses.</p> <p>Perform this task on the node with the configuration you want to keep.</p> <ul style="list-style-type: none"> ❑ Configure Local IP addresses for the heartbeat interfaces. ❑ Change the Local scope to Cluster scope for IP addresses associated with virtual servers and proxy IP addresses.

Scenario A: Combining Two New, Unconfigured File Directors into a Cluster	Scenario B: Combining an Existing, Stand-Alone File Director into a Cluster with Another New or Existing File Director
<p>Task A-4: Configure the Multicast Base Address and Set a Backup Node for the Cluster.</p> <p>Perform this task on the Primary node only.</p> <ul style="list-style-type: none"> ❑ (Optional) Set the Multicast Base Address to focus heartbeat communications. ❑ Set a Backup node for the cluster. ❑ Wait until the node Primary node is activated. 	<p>Task B-4: Configure Node Roles and the Multicast Base Address of the Cluster.</p> <p>Perform this task on the node with the configuration you want to keep.</p> <ul style="list-style-type: none"> ❑ (Optional) Set the Multicast Base Address to focus heartbeat communications. ❑ Set a Primary (or Backup) node for the cluster. ❑ Wait until the node you are configuring is activated.
<p>Task A-5: Configure Cluster-Wide Settings.</p> <p>Perform this task on the Primary node only.</p> <ul style="list-style-type: none"> ❑ Configure the following: <ul style="list-style-type: none"> – Proxy IP address – File servers – Virtual servers, shares, and exports – Synthetic directories and union directories – SNMP information 	<p>Task B-5: Configure the Remaining Node.</p> <p>Perform this task on the node with the configuration you <i>do not</i> want to keep.</p> <ul style="list-style-type: none"> ❑ Configure Local IP addresses for heartbeat interfaces ❑ Set the Multicast Base Address to focus heartbeat communications. (If already set on the other node.) ❑ Set the Backup (or Primary) node for the cluster.
<p>Task A-6: Configure the Backup Node.</p> <p>Perform this task on the Backup node only.</p> <ul style="list-style-type: none"> ❑ Configure Local IP addresses for heartbeat interfaces ❑ Set the Multicast Base Address to focus heartbeat communications. (If already set on the Primary.) ❑ Set the Primary node for the cluster. 	<p>Task B-6: Confirm User Account and Network Configuration.</p> <p>Perform this task on both nodes.</p> <ul style="list-style-type: none"> - User Accounts - Default Gateway - DNS Settings - Win Settings - NTP Server - Log Rotation

Scenario A: Combining Two New, Unconfigured File Directors into a Cluster	Scenario B: Combining an Existing, Stand-Alone File Director into a Cluster with Another New or Existing File Director
<p>Task A-7: Complete User Account and Network Configuration.</p> <p>Perform this task on both nodes.</p> <ul style="list-style-type: none"> - User Accounts - Default Gateway - DNS Settings - Win Settings - NTP Server - Log Rotation 	<p>Task B-7: Verify the Cluster Configuration.</p> <ul style="list-style-type: none"> ❑ After a few moments, verify that the Primary node is Active and the Backup node is on Standby.
<p>Task A-8: Verify the Cluster Configuration.</p> <ul style="list-style-type: none"> ❑ After a few moments, verify that the Primary node is Active and the Backup node is on Standby. 	

Scenario A: Configuring Two New, Unconfigured File Directors as a Cluster

This scenario describes the steps you take to create a cluster using two new, unconfigured File Directors.

Task A-1: Run the Initial Configuration Script

- Install and run the initial configuration script on each File Director node.

Run the initial configuration as if each File Director is a stand-alone system. For more information, see “Installing and Configuring the File Director” in the *File Director Quick Start Guide*.

Task A-2: Connect the Heartbeat Interfaces

To allow heartbeat communications to occur, connect at least two interfaces between the File Director nodes in a cluster. The two nodes must be connected to each other either directly via a cable or over a network. To establish a redundant communications path, use both a direct and a network connection between the nodes.

1. Use a network cable to directly connect the two File Directors in the cluster. Connect the cable to an interface port on each File Director.

Important: Although directly connecting the interface ports on the File Directors in a cluster is not required, it is recommended to reduce the possibility of unnecessary failovers in case an external device (such as a switch) fails.

If you use both a direct and a network connection for redundancy instead of two direct connections, it enables you to use the port that is connected to a switch for non-high availability traffic such as management traffic (SSH/HTTPS/SNMP) or NSA traffic.

To avoid errors, connect the same interfaces on each node. For example, make a direct connection between the *eth1* ports on each node, and an indirect connection over a network between the *eth2* ports on each node. You may substitute other ports as long as they are the same set of ports. That is, don't connect the *eth1* port from one node to the *eth4* port on the other node.

Note: A cross-over Ethernet cable is *not* required, but you can use it.




eth1 interface port



Connecting two nodes in a cluster directly via an interface port

2. Connect another set of ports on the two nodes over the network.

For example, make an indirect connection via a switch between the *eth2* interface ports on each node.



Connecting two nodes in a cluster indirectly via a switch

Task A-3: Configure IP Addresses on the Primary

On the node you want to be the Primary node, set the Local and Cluster scope for interfaces.

SETTING LOCAL SCOPE

For heartbeat interfaces:

1. Assign a unique IP address to each interface you connected for heartbeat communications.

- Set the scope to **Local** for each of those IP addresses.

For example, assign a unique **Local** scope IP address to the *eth1* and *eth2* ports on the Primary node.

For information on assigning an IP address to an interface and setting the scope, see “Configuring the File Director IP Addresses” on page 2-6.

The File Director Primary node will use the local IP address to exchange heartbeat messages with the Backup node. For redundancy of communications paths, you can also assign additional local IP addresses to other interfaces, but this is not required.

SETTING CLUSTER SCOPE

On the node you want to be the Primary node, add IP addresses with **Cluster** scope for all of the following:

- IP addresses that are used (or will be used) for virtual servers. (See “Specifying a Virtual Server” on page 2-20.)
- Proxy IP addresses. (See “Setting a Proxy IP Address for the File Director” on page 2-9.)
- File Director IP addresses used for access to file servers. (See step 6 on page 2-17 in the “Specifying a File Server” section.)

This step is necessary to allow the Standby node to automatically take over those Cluster IP addresses if the Active node fails.

To add IP addresses with **Cluster** scope:

- In the Navigation frame on the Primary (Active) node, expand **File Director Configuration**, expand **Network Settings**, and then click the **IP Configuration** link.

The **Network Settings** page appears.

The screenshot displays the 'NETWORK SETTINGS' page in the File Director Administrator interface. On the left, the navigation pane shows 'File Director Configuration' expanded, with 'Network Settings' selected and 'IP Configuration' highlighted. The main panel shows the 'IP ADDRESS LIST' table with two entries:

IP ADDRESS	SUBNET MASK	INTERFACE	SCOPE	ACTIVE	
172.22.1.242	255.255.255.0	eth0	local	Yes	Delete Display
172.22.1.244	255.255.255.0	eth0	local	Yes	Delete Display

Below the table is an 'Add' button. Further down, the 'CLUSTER' section shows 'PRIMARY NODE' and 'BACKUP NODE' both as '(none assigned)' with an 'Assign' button. The 'DEFAULT GATEWAY' section shows 'GATEWAY IP ADDRESS' as 172.22.1.1 and 'INTERFACE' as eth0, with an 'Assign' button.

- To add a new IP address and set the scope to cluster, click **Add** under the **IP Address List** section.

3. From the **Scope** drop-down menu, choose **Cluster**.

For information on setting the other options for an IP address, see “Configuring the File Director IP Addresses” on page 2-6.

4. Click **Submit** to commit the new settings.

Task A-4: Configure the Multicast Base Address and Set a Backup Node for the Cluster

On the node you want to be the Primary node, set a multicast base address you want to use for the cluster (optional) and set a Backup node for the cluster.

Setting a multicast base address, enables you to focus heartbeat communications between File Directors in a cluster. The File Director uses the base address to automatically derive and assign a unique multicast IP address to each interface in the cluster. The two File Director nodes use the multicast IP addresses to exchange messages with each other without communicating with other devices on the network. This prevents unnecessary traffic on the network, and it prevents cross communications between clusters that are on the same network.

Note: Because the procedure for using the Management Console to set the roles and multicast base address is significantly different from the procedure for using the CLI, both procedures are included in this section.

USING THE MANAGEMENT CONSOLE

To use the Management Console to set the Primary node and cluster multicast base address:

1. In the Navigation frame on the node you want to be the Primary node, expand **File Director Configuration**, expand **Network Settings**, and then click the **IP Configuration** link.

The **Network Settings** page appears.

2. Click **Assign** in the **Cluster** section.

The **Cluster Configuration** page appears.

3. For **Backup Node Name**, enter the host name of the File Director you want to configure as the Backup node of the cluster.

Make sure you enter a host name, not an IP addresses.

Note: The **Primary Node Name** (host name of the current File Director you are configuring) is filled-in automatically.

4. (Optional) For **Multicast IP Address**, enter the multicast base IP address.

Enter the address in dotted quad notation. The first three digits *www* in the address *www.xxx.yyy.zzz* must be in the range of 224-239. The range of the digits *xxx* and *yyy* is 0-255. The range of the last three digits *zzz* is 1-240. For example: 239.252.200.100

It is strongly recommended that you use an address from one of the following site-specific ranges:

- 239.252.000.001 to 239.252.255.240
- 239.253.000.001 to 239.253.255.240
- 239.254.000.001 to 239.254.255.240
- 239.255.000.001 to 239.255.255.240

Although setting a multicast base address is optional, it is strongly recommended – especially if you have multiple File Director clusters on the same network.

Important: If there are multiple clusters on the same network, be sure to assign a different multicast base address to each cluster to prevent cross communications between the clusters. Although using the same multicast base address for multiple clusters will not prevent the clusters from operating correctly, it is not efficient and can potentially cause incorrect pairings of nodes in a cluster.

5. Click **Submit** to commit the new settings.

A few moments after you click **Submit**, the Primary node becomes the Active node in the cluster.

Important: Make sure the Primary node status is Active and the Backup node status is Dead before proceeding. It may take a few seconds for the node to activate.

USING THE CLI

To use the CLI to set a cluster multicast base address and a Backup node for the cluster:

1. (Optional) On the Primary node, run the **cluster multicast** CLI command to set the multicast base address for the cluster. For example, **add cluster multicast 239.252.200.100** where **<239.252.200.100>** is a multicast base address.

For more information on using this command, see “cluster multicast” on page 23 in the *File Director CLI Reference Guide*.

Important: Be sure to use this command to set the multicast base address before using the **cluster node** command in the next step.

2. On the Primary node, run the **cluster node** CLI command and set the host name of the File Director you want to configure as the Backup node of the cluster. For example, **add cluster node server-xyz backup**

For more information on using this command, see “cluster node” on page 25 in the *File Director CLI Reference Guide*.

3. On the Primary node, run the `show cluster node` CLI command to verify the status of the nodes.

The output should show the current node as the Primary node. The status of the Backup node will be *dead*; this status will change to *standby* after you set up the Backup node.

Important: Make sure the Primary node status is *active* and the Backup node status is *dead* before proceeding. It may take a few seconds for the node to activate.

Task A-5: Configure Cluster-Wide Settings

As necessary on the Primary node *only* (which is now Active), configure all items related to the following:

- Access to file servers
- Virtual servers
- Virtual shares and exports
- Synthetic directories.

For more information, see the appropriate chapters in this book.

Important: When configuring these items, be sure to use IP addresses that have Cluster scope.

Task A-6: Configure the Backup Node

This section describes how to set up the Backup node:

Warning: After you set up the Backup node, the Primary node automatically copies the appropriate configuration information to the Backup node; therefore, the File Director you want to designate as the Backup node should not store any pre-configured cluster-wide information, such as virtual server configuration. When you set a File Director as having a Backup role, this cluster-wide access information will be overwritten by the cluster-wide information from the Primary node.

On the node you want to be the Backup, do the following:

1. Assign a unique IP address to each interface you connected for heartbeat communications. Set the scope to Local for each of those IP addresses.

For information on assigning an IP address to an interface and setting the scope, see “Configuring the File Director IP Addresses” on page 2-6.

2. Use either the Management Console or CLI commands to set the Primary node and (if applicable) the matching multicast base address.

Note: Because the procedure for using the Management Console to set the roles and multicast base address is significantly different from the procedure for using the CLI, both procedures are included in this section.

USING THE MANAGEMENT CONSOLE

To use the Management Console:

1. In the Navigation frame on the node you want to be the Backup node, expand **File Director Configuration**, expand **Network Settings**, and then click the **IP Configuration** link. The **Network Settings** page appears.
2. Click **Assign** in the **Cluster** section. The **Cluster Configuration** page appears.
3. For **Primary Node Name**, enter the host name of the Primary node.
4. If you specified a multicast base address on the Primary node, set the same multicast base address on the Backup node.
5. Click **Submit** to commit the new settings.

After you click **Submit**, the Primary node automatically copies the appropriate configuration information to the Backup node. After a few seconds, the Backup node becomes the Standby node and is ready to begin providing services if the Primary (Active) node fails. The two nodes are now operating as a cluster.

USING THE CLI

To use the CLI:

1. If you specified a multicast base address on the Primary node, run the **cluster multicast** CLI command on the Backup node and set the same multicast base address (for example, **add cluster multicast 239.252.200.100**).

Important: Be sure to use this command to set the multicast base address before using the **cluster node** command in the next step.

2. On the Backup node, run the **cluster node** CLI command and set the host name of the Primary node. For example, run **add cluster node server-abc primary**

After you run this command, the Primary node automatically copies the appropriate configuration information to the Backup node. After a few seconds, the Backup node becomes the Standby node and is ready to begin providing services if the Primary (Active) node fails. The two nodes are now operating as a cluster.

3. On the Backup node, run the **cluster node** CLI command to verify the status of the nodes.

show cluster node

The output should show the current node as Backup node and the other node as the Primary node. The status of the Primary node will be *active*, and the status of the Backup node will be *standby*.

Task A-7: Complete User Account and Network Configuration

You must configure all items related to client access on the Active node *only*, such as file servers, virtual servers, and synthetic and union directories. But you must configure on *both* nodes network items such as DNS and gateway settings.

- On the Primary node, complete the configuration by setting the following:
 - User Accounts
 - Default Gateway
 - DNS Settings
 - WINS Server
 - NTP Server
 - Log Rotation
- On the Backup node, complete the configuration by setting the following:
 - User Accounts
 - Default Gateway
 - DNS Settings
 - WINS Server
 - NTP Server
 - Log Rotation

For more information, see the appropriate chapters in this book.

Task A-8: Verify the Cluster Configuration

After waiting a few moments to allow the nodes to initialize and synchronize, verify the cluster nodes are operating properly by checking that the status of one node is Active and the other is Standby.

For more information, see “Checking the Status of Each Cluster Node” on page 7-25.

Tip: For help with solving problems related to setting up a cluster, see “Troubleshooting Problems with Setting Up a Cluster” on page 7-26.

Scenario B: Configuring an Existing File Director into a Cluster with Another New or Existing File Director

In this scenario you configure an existing File Director into a cluster with another new (or existing) File Director. For example, you might have a File Director installed and you purchase another new File Director to create a High Availability pair.

In this scenario, it is important to understand how configurations are unified across nodes in a cluster.

When you create a cluster, both nodes in the cluster receive the same configuration. The configuration from one of the File Director nodes is kept, and the configuration from the other node is deleted. The configuration can be kept from either node – it does not always come from the Primary node. Therefore, it is very important to back up any existing node configurations before creating a cluster.

The cluster configuration comes from the node that first becomes active. If you add the Backup node to the cluster first, it will become active. This happens because the Backup node views the Primary node as dead – since you haven't yet configured the Primary node for the cluster. When the Primary is added to the cluster, it replaces its configuration with the configuration of the Backup node before becoming active. Alternatively, if you add the Primary node to the cluster first, the Backup node takes its configuration from the Primary.

Activating a Backup node first is useful if you have an existing File Director running on an older hardware platform and you buy a new File Director running on newer hardware. Typically, you want the node running on newer hardware to become the Primary node, but you want to load it with the configuration that is still on the node with older hardware (which will become the Backup node). In this situation, you would add the node with older hardware (the Backup) to the cluster first before adding the node with newer hardware (the Primary).

In summary, when you create a cluster, make sure you add the node with the desired configuration first, then add the node with the expendable configuration.

Task B-1: Verify Stand-Alone Configuration

Before you create the cluster, make sure both File Director nodes are configured for stand-alone operation:

1. If one of the nodes you are adding to the cluster is new, install and run the initial configuration script on the new File Director node.

Run the initial configuration as if each File Director is a stand-alone system. For more information, see “Installing and Configuring the File Director” in the *File Director Quick Start Guide*.

2. On the node with the configuration you want to keep, verify the following:
 - Proxy IP address
 - File servers
 - Virtual servers, shares, and exports
 - Synthetic directories and union directories
 - SNMP information

3. Make a backup of both configurations.

For more information, see “Backing Up and Restoring Configuration Information for the File Director” on page 6-15.

Task B-2: Connect the Heartbeat Interfaces

To allow heartbeat communications to occur, connect at least two interfaces between the File Director nodes in a cluster. The two nodes must be connected to each other either directly via a cable or over a network. To establish a redundant communications path, use both a direct and a network connection between the nodes.


1. Use a network cable to directly connect the two File Directors in the cluster. Connect the cable to an interface port on each File Director.

Important: Although directly connecting the interface ports on the File Directors in a cluster is not required, it is recommended to reduce the possibility of unnecessary failovers in case an external device (such as a switch) fails.

If you use both a direct and a network connection for redundancy instead of two direct connections, it enables you to use the port that is connected to a switch for non-high availability traffic such as management traffic (SSH/HTTPS/SNMP) or NSA traffic.

To avoid errors, connect the same interfaces on each node. For example, make a direct connection between the *eth1* ports on each node, and an indirect connection over a network between the *eth2* ports on each node. You may substitute other ports as long as they are the same set of ports. That is, don't connect the *eth1* port from one node to the *eth4* port on the other node.


Note: A cross-over Ethernet cable is *not* required, but you can use it.



Connecting two nodes in a cluster directly via an interface port

2. Connect another set of ports on the two nodes over the network.

For example, make an indirect connection via a switch between the *eth2* interface ports on each node.



Connecting two nodes in a cluster indirectly via a switch

Task B-3: Configure IP Addresses

SETTING CLUSTER SCOPE

On the node with the configuration you want to keep, change the **Local** scope setting to **Cluster** scope on all IP addresses that are used for virtual servers, proxy IP addresses, and the File Director IP addresses used for access to file servers. This step is necessary to allow the Standby node to automatically take over those Cluster IP addresses if the Active node fails.

To change the **Local** scope setting to **Cluster**:

1. In the Navigation frame on the Primary (Active) node, expand **File Director Configuration**, expand **Network Settings**, and then click the **IP Configuration** link.

The **Network Settings** page appears.

NETWORK SETTINGS User: admin

IP ADDRESS LIST

IP ADDRESS	SUBNET MASK	INTERFACE	SCOPE	ACTIVE	
172.22.1.242	255.255.255.0	eth0	local	Yes	Delete Display
172.22.1.244	255.255.255.0	eth0	local	Yes	Delete Display

1 - 2 of 2 records MAX ROWS 20

Add

CLUSTER

PRIMARY NODE	BACKUP NODE	MULTICAST IP ADDRESS	
(none assigned)	(none assigned)		Assign

DEFAULT GATEWAY

GATEWAY IP ADDRESS	INTERFACE	
172.22.1.1	eth0	Assign

2. Click **Display** next to the IP Address under the **IP Address List** section. Then click **Edit** to change the scope.
3. From the **Scope** drop-down menu, choose **Cluster**.
4. Click **Submit** to commit the new settings.

Make sure you have changed the scope for all of the following:

- Virtual servers. (See “Specifying a Virtual Server” on page 2-20.)
- Proxy IP addresses. (See “Setting a Proxy IP Address for the File Director” on page 2-9.)
- The File Director IP addresses used for access to file servers. (See step 6 on page 2-17 in the “Specifying a File Server” section.)

Important: Be sure to check the IP address shown in the **File Director IP Address for Server Access** field on the **File Server Configuration** page. (See “Displaying the File Server Access Configuration” on page 2-18.) If you didn’t explicitly set this address when you configured access to the file server, the File Director automatically specifies an IP address and may give it Local scope. Make sure this IP address is set to Cluster scope.

Task B-4: Configure Node Roles and the Multicast Base Address of the Cluster

On the node with the configuration you want to keep, set a multicast base address you want to use for the cluster (optional), then set a Backup or Primary role for the other node.

Setting a multicast base address, enables you to focus heartbeat communications between File Directors in a cluster. The File Director uses the base address to automatically derive and assign a unique multicast IP address to each interface in the cluster. The two File Director nodes use the multicast IP addresses to exchange messages with each other without communicating with other devices on the network. This prevents unnecessary traffic on the network, and it prevents cross communications between clusters that are on the same network.

Note: Because the procedure for using the Management Console to perform this task is significantly different from the procedure for using the CLI, both procedures are included in this section.

USING THE MANAGEMENT CONSOLE

To use the Management Console:

1. In the Navigation frame on the node with the configuration you want to keep, expand **File Director Configuration**, expand **Network Settings**, and then click the **IP Configuration** link.
The **Network Settings** page appears.
2. Click **Assign** in the **Cluster** section.

The **Cluster Configuration** page appears.

3. Set a Primary or Backup role for the other node:
 - If the node you are currently configuring will act as the Primary node, enter the name of the other node in the **Backup Node Name** field.
 - If the node you are currently configuring will act as the Backup node, enter the name of the other node in the **Primary Node Name** field.

Make sure you enter a host name, not an IP address.

4. (Optional) For **Multicast IP Address**, enter the multicast base IP address.

Enter the address in dotted quad notation. The first three digits *www* in the address *www.xxx.yyy.zzz* must be in the range of 224-239. The range of the digits *xxx* and *yyy* is 0-255. The range of the last three digits *zzz* is 1-240. For example: 239.252.200.100

It is strongly recommended that you use an address from one of the following site-specific ranges:

- 239.252.000.001 to 239.252.255.240
- 239.253.000.001 to 239.253.255.240
- 239.254.000.001 to 239.254.255.240
- 239.255.000.001 to 239.255.255.240

Although setting a multicast base address is optional, it is strongly recommended – especially if you have multiple File Director clusters on the same network.

Important: If there are multiple clusters on the same network, be sure to assign a different multicast base address to each cluster to prevent cross communications between the clusters. Although using the same multicast base address for multiple clusters will not prevent the clusters from operating correctly, it is not efficient and can potentially cause incorrect pairings of nodes in a cluster.

5. Click **Submit** to commit the new settings.

After you click **Submit**, the node becomes the Active node in the cluster.

Important: Make sure the node is active before proceeding. It may take a few seconds for the node to activate.

USING THE CLI

To use the CLI:

1. (Optional) On the node with the configuration you want to keep, run the **cluster multicast** CLI command to set the multicast base address for the cluster. For example, add **cluster multicast 239.252.200.100** where **<239.252.200.100>** is a multicast base address.

For more information on using this command, see “cluster multicast” on page 23 in the *File Director CLI Reference Guide*.

Important: Be sure to use this command to set the multicast base address before using the **cluster node** command in the next step.

2. On the node with the configuration you want to keep, run the **cluster node** CLI command and set a Primary or Backup role for the other node:
 - If the node you are currently configuring will act as the Primary node, set the other node as the Backup. For example, add **cluster node server-xyz backup**
 - If the node you are currently configuring will act as the Backup node, set the other node as the Primary. For example, add **cluster node server-xyz primary**Make sure you enter a host name, not an IP address.

For more information on using this command, see “cluster node” on page 25 in the *File Director CLI Reference Guide*.

3. On the node with the configuration you want to keep, run the **show cluster node** CLI command to verify the status of the nodes.

The output should show the current node as active. The status of the other node will be *dead*; this status will change to *standby* after you set up the other node.

Important: Make sure the current node is active before proceeding. It may take a few seconds for the node to activate.

Task B-5: Configure the Remaining Node

Perform this task on the node with the configuration you *do not* want to keep.

1. Assign a unique IP address to each interface you connected for heartbeat communications. Set the scope to Local for each of those IP addresses.

For information on assigning an IP address to an interface and setting the scope, see “Configuring the File Director IP Addresses” on page 2-6.

2. Use either the Management Console or CLI commands to set the role of the other node (either Primary or Backup) and (if applicable) the matching multicast base address.

Note: Because the procedure for using the Management Console to perform this task is significantly different from the procedure for using the CLI, both procedures are included in this section.

USING THE MANAGEMENT CONSOLE

To use the Management Console:

1. In the Navigation frame with the configuration you *do not* want to keep, expand **File Director Configuration**, expand **Network Settings**, and then click the **IP Configuration** link. The **Network Settings** page appears.
2. Click **Assign** in the **Cluster** section. The **Cluster Configuration** page appears.
3. Set a Primary or Backup role for the other node:
 - If the node you are currently configuring will act as the Primary node, enter the name of the other node in the **Backup Node Name** field.
 - If the node you are currently configuring will act as the Backup node, enter the name of the other node in the **Primary Node Name** field.
4. If you specified a multicast base address on the other node, set the same multicast base address on this node.
5. Click **Submit** to commit the new settings.

After you click **Submit**, the Primary node automatically copies the appropriate configuration information to the Backup node. After a few seconds, the Backup node becomes the Standby node and is ready to begin providing services if the Primary (Active) node fails. The two nodes are now operating as a cluster.

USING THE CLI

To use the CLI:

1. If you specified a multicast base address on the other node, run the **cluster multicast** CLI command on the current node and set the same multicast base address (for example, **add cluster multicast 239.252.200.100**).
2. Run the **cluster node** CLI command and set a Primary or Backup role for the other node:

- If the node you are currently configuring will act as the Primary node, set the other node as the Backup. For example, **add cluster node server-xyz backup**
- If the node you are currently configuring will act as the Backup node, set the other node as the Primary. For example, **add cluster node server-xyz primary**

After you run this command, the Primary node automatically copies the appropriate configuration information to the Backup node. After a few seconds, the Backup node becomes the Standby node and is ready to begin providing services if the Primary (Active) node fails. The two nodes are now operating as a cluster.

3. On the node you are currently configuring, run the **cluster node** CLI command to verify the status of the nodes.

show cluster node

The output should show the current node as Backup node and the other node as the Primary node. The status of the Primary node will be *active*, and the status of the Backup node will be *standby*.

Task B-6: Confirm User Account and Network Configuration

You must configure all items related to client access on the Active node *only*, such as file servers, virtual servers, and synthetic and union directories. But you must configure on *both* nodes network items such as DNS and gateway settings.

- On the Primary node, complete the configuration by setting the following:
 - User Accounts
 - Default Gateway
 - DNS Settings
 - Win Settings
 - NTP Server
 - Log Rotation
- On the Backup node, complete the configuration by setting the following:
 - User Accounts
 - Default Gateway
 - DNS Settings
 - Win Settings
 - NTP Server
 - Log Rotation

For more information, see the appropriate chapters in this book.

Task B-7: Verify the Cluster Configuration

After waiting a few moments to allow the nodes to initialize and synchronize, verify the cluster nodes are operating properly by checking that the status of one node is Active and the other is Standby.

For more information, see “Checking the Status of Each Cluster Node” on page 7-25.

Tip: For help with solving problems related to setting up a cluster, see “Troubleshooting Problems with Setting Up a Cluster” on page 7-26.

Checking the Status of Each Cluster Node

After you finish configuring the two nodes in a cluster, the nodes will begin initializing and synchronizing. After a few moments, you can verify the cluster nodes are operating properly by checking that the status of one node is Active and the other is Standby.

To check the status of each cluster node:

1. In the Navigation frame on either node, expand **File Director Configuration**, expand **Network Settings**, and then click the **IP Configuration** link.

The **Network Settings** page appears.

2. Click **Display** in the **Cluster** section.

The current status of each node is displayed in the **Cluster Configuration** page.

Node status	Description
up	This node is initializing the high-availability subsystem.
init	The other node is reported in this state while the high-availability subsystem on this node is initializing.
active	This node is the Active node.
standby	This node is the Standby node.
syncing	This node is synchronizing its configuration and cluster resources from the Active node. After the synchronization is complete, this node is the Standby node.
dead	This node is not operating (missing heartbeats).
incompatible	This node reports the software on the other node is incompatible with the software on this node, which prevents synchronization and fail over in the cluster. The other node can consider itself to have an active, standby, or initialization status, but it will not be able to fail over to the current node.
transition	This node is acquiring or releasing cluster resources and changing its role.

Tip: You can also check the status by using the `show cluster node` CLI command. For more information, see “cluster node” on page 25 in the *File Director CLI Reference Guide*.

Troubleshooting Problems with Setting Up a Cluster

This section contains suggestions for troubleshooting problems you might have while setting up a cluster.

Problem: When you attempt to add a node to a cluster, an error message states that a certain IP address *www.xxx.yyy.zzz* does not have Cluster scope.

Solution: All IP addresses that are used for virtual servers, access to NFS servers, or as a Proxy IP must have Cluster scope. The error message is stating that IP address *www.xxx.yyy.zzz* is being used for such a purpose but it has Local scope. Before you can add the node, you must change the scope of this IP address to Cluster, or else substitute a different IP address that has Cluster scope for the same purpose.

If you didn't explicitly set an IP address when you configured access to a file server, the automatically selected address may have a Local scope. Make sure this IP address is set to Cluster scope. Also, check all backend file servers to make sure the Via IP Address is not set to a local IP address.

Problem: When you attempt to run a CLI command, an error message states that the command is not available and that the node is not Active.

Solution: You can only run CLI commands that modify cluster-wide configuration settings on the Active node. This error message appears if you attempt to run this type of command on the Standby node. Run the command on the Active node that is identified in the error message.

Disbanding a Cluster

If you want to use the File Director nodes as stand-alone systems, you can disband the cluster.

Disbanding a Cluster If Both Nodes are Operating

Important: You must begin the process of disbanding the cluster on the Standby node. If the Standby node is not currently operating but could become operational, see the next section, "Disbanding a Cluster If Only One Node is Operating."

TO USE THE MANAGEMENT CONSOLE TO DISBAND A CLUSTER IF BOTH NODES ARE OPERATING:

1. In the Navigation frame on the Standby node, expand **File Director Configuration**, expand **Network Settings**, and then click the **IP Configuration** link.
The **Network Settings** page appears.
2. Click **Remove** in the **Cluster** section.
3. Click **Confirm Remove Cluster**.
4. Click **Display** in the **Cluster** section.
The **Cluster Configuration** page appears.
5. Click **Remove Multicast IP** under the **Multicast IP Address** section.

6. Click **Confirm Remove Multicast IP**.

The Standby node becomes a stand-alone system and automatically deletes all service-related configuration information that it copied from the Active node, including Cluster IP addresses, file servers, virtual servers, and information about synthetic directories.

7. On the Active node, repeat steps 1 through 6.

The Active node becomes a stand-alone system and retains all service-related configuration information, including Cluster IP addresses, file servers, virtual servers, and information about synthetic directories.

TO USE THE CLI TO DISBAND A CLUSTER IF BOTH NODES ARE OPERATING:

1. On the Standby node, run the following **cluster node** CLI command and set the host name of the Active node:

```
clear cluster node <hostname-of-active-node>
```

For example, on the Standby node called *server-abc*, run this command to remove the Active node called *server-xyz*:

```
clear cluster node server-xyz
```

For more information on using this command, see “cluster node” on page 25 in the *File Director CLI Reference Guide*.

2. (Optional) On the Standby node, run the following **cluster multicast** CLI command to remove the multicast base address:

```
clear cluster multicast
```

For more information on using this command, see “cluster multicast” on page 23 in the *File Director CLI Reference Guide*.

3. On the Active node, run the following **cluster node** CLI command and set the host name of the Standby node:

```
clear cluster node <hostname-of-standby-node>
```

For example, on the Active node called *server-xyz*, run this command to remove the Standby node called *server-abc*:

```
clear cluster node server-abc
```

4. (Optional) On the Active node, run the following **cluster multicast** CLI command to remove the multicast base address:

```
clear cluster multicast
```

The Active node becomes a stand-alone system and retains all service-related configuration information, including Cluster IP addresses, file servers, virtual servers, and information about synthetic directories.

Disbanding a Cluster If Only One Node is Operating

TO USE THE MANAGEMENT CONSOLE TO DISBAND A CLUSTER IF ONLY ONE NODE IS OPERATING:

1. Disconnect the non-operating node from the network or shut it down.
Warning: If one of the nodes is not currently operating (*dead* status) but could become operational, be sure to either disconnect it from the network or shut it down. Otherwise, if the non-operating node starts operating after you have disbanded the cluster using the Active node, the two systems will use the same IP addresses which will cause problems on the network.
2. In the Navigation frame on the Active node, expand **File Director Configuration**, expand **Network Settings**, and then click the **IP Configuration** link.
The **Network Settings** page appears.
3. Click **Remove** in the **Cluster** section.
4. Click **Confirm Remove Cluster**.
5. At the warning message, confirm the removal of the cluster.
6. Click **Display** in the **Cluster** section.
The **Cluster Configuration** page appears.
7. Click **Remove Multicast IP** under the **Multicast IP Address** section.
8. Click **Confirm Remove Multicast IP**.

The Active node becomes a stand-alone system and retains all service-related configuration information, including Cluster IP addresses, file servers, virtual servers, and information about synthetic directories.

Important: Before connecting and using the non-operating node on the network again, disband the cluster on the non-operating node by following the instructions in this section. Also, make sure you delete all service-related configuration information on the non-operating node. You can reconfigure the node off the network by either using the Management Console on a laptop that is connected directly to the node, or by using CLI commands on a serial console that is connected directly to the node.

TO USE THE CLI TO DISBAND A CLUSTER IF ONLY ONE NODE IS OPERATING:

1. Disconnect the non-operating node from the network or shut it down.
Warning: If one of the nodes is not currently operating (*dead* status) but could become operational, be sure to either disconnect it from the network or shut it down. Otherwise, if the non-operating node starts operating after you have disbanded the cluster using the Active node, the two systems will use the same IP addresses which will cause problems on the network.

2. On the Active node, run the following **cluster node** CLI command and set the host name of the non-operating node:

```
clear cluster node <hostname-of-non-operating-node>
```

For example, on the Active node called *server-xyz*, run this command to remove the non-operating node called *server-abc*:

```
clear cluster node server-abc
```

3. (Optional) On the Active node, run the following **cluster multicast** CLI command to remove the multicast base address:

```
clear cluster multicast
```

You only need to run this command if you configured a multicast base address for the cluster.

The Active node becomes a stand-alone system and retains all service-related configuration information, including Cluster IP addresses, file servers, virtual servers, and information about synthetic directories.

Important: Before connecting and using the non-operating node on the network again, disband the cluster on the non-operating node by following the instructions in this section. Also, make sure you delete all service-related configuration information on the non-operating node. You can reconfigure the node off the network by either using the Management Console on a laptop that is connected directly to the node, or by using CLI commands on a serial console that is connected directly to the node.

Changing the Multicast Base Address

Warning: If you have already assigned roles to the nodes and thus created an operating cluster, you must disband a cluster before changing the multicast base address. Otherwise, changing the multicast base address can cause momentary loss of communications and activate the Standby node to take over operations from the Active node.

To change the multicast base address:

1. If you want to change the multicast base address in a cluster that is operating, disband the cluster.

For more information, see “Disbanding a Cluster” on page 7-26.

2. In the Navigation frame on the Active node, expand **File Director Configuration**, expand **Network Settings**, and then click the **IP Configuration** link.

The **Network Settings** page appears.

3. Click **Display** in the **Cluster** section.

The **Cluster Configuration** page appears.

The screenshot shows the 'CLUSTER CONFIGURATION' page. On the left is a navigation menu with categories like File Director Configuration, Network Settings, IP Configuration, Host & DNS Settings, Ethernet Interfaces, WDBS IP Address, NTP Servers, System, Log File Rotation, and Shutdown. The main content area is titled 'CLUSTER CONFIGURATION' and shows the user 'admin'. It contains three sections: 'PRIMARY CLUSTER NODE' with fields for Host Name (server1), Current Status (active), and Software Version (1.0.7-01); 'BACKUP CLUSTER NODE' with fields for Host Name (server2), Current Status (standby), and Software Version (1.0.7-01); and 'MULTICAST IP ADDRESS' with a field for Multicast IP Address (239.252.000.002). At the bottom are buttons for 'Add/Edit Multicast IP', 'Remove Multicast IP', and 'Back'.

PRIMARY CLUSTER NODE	
Host Name	server1
Current Status	active
Software Version	1.0.7-01

BACKUP CLUSTER NODE	
Host Name	server2
Current Status	standby
Software Version	1.0.7-01

MULTICAST IP ADDRESS	
Multicast IP Address	239.252.000.002

Buttons: Add/Edit Multicast IP, Remove Multicast IP, Back

4. Click **Add/Edit Multicast IP** under the **Multicast IP Address** section.
5. For **Multicast IP Address**, enter the multicast base IP address in dotted quad notation.
6. Click **Submit** to commit the new settings.
7. On the Backup node, assign the same multicast base address you assigned to the Primary node by repeating steps 2 through 6 in this procedure.

APPENDIX A

Specifications for the File Director

This appendix contains the following topics about specifications for the File Director:

- File Director Hardware Specifications
- File Director Compatibility
- TCP and UDP Port Specifications

File Director Hardware Specifications

The hardware specifications of the File Director are:

- Includes four 10/100/1000 Ethernet ports and two embedded 10/100/1000 Ethernet ports
- RAID 5 array controller (716 GB)
- Redundant power supply
- 2U rackmount chassis

File Director Compatibility

For NFS, the File Director supports NFS system level authentication (AUTH_SYS/AUTH_UNIX). For CIFS, the File Director supports Plaintext, NTLM and NTLMv2 authentication, and NTLM 0.12 and NT LANMAN 1.0 dialects.

The File Director uses the following standard NAS protocols to communicate with both clients and servers:

- CIFS clients, servers
- NFS (v2, v3) clients, servers

Note: NFS servers must support both v2 and v3.

Compatibility with NFS Clients

- Linux Redhat v.3 WS
- Linux Redhat 7.3, 8, 9
- Solaris 8, 9

Compatibility with NFS Servers

- Network Appliance ONTAP 6.4, 6.5
- EMC Celerra, OS version 5.3.14.2
- Linux Redhat v.3 ES
- Linux Redhat 7.3, 8, 9
- Solaris 8, 9
- Windows 2003 Storage Server (NFS Services)

Compatibility with CIFS Clients

- Microsoft Windows XP Professional (SP1a)
- Microsoft Windows 2000 Professional (SP4)

Compatibility with CIFS Servers

- Network Appliance ONTAP 6.4, 6.5
- EMC Celerra, OS version 5.3.14.2
- Windows Server 2003 Enterprise Edition
- Windows 2000 Advanced Server (SP4)
- Samba 3.0 on Red Hat Enterprise Linux ES v.3

TCP and UDP Port Specifications

If you set up a firewall between the File Director and clients or servers, be aware that the File Director listens on the following ports and needs access through the firewall to operate properly:

Ports	Description
TCP 22	The File Director listens on TCP port 22 on all IP addresses if the <i>admin_ssh</i> service is enabled. (See “Restarting or Stopping System Services on the File Director” on page 6-13.)
TCP 443	The File Director listens on TCP port 443 on all IP addresses if the <i>admin_www</i> service is enabled.
TCP 139/445	The File Director listens on TCP ports 139 and 445 on IP addresses that have a CIFS virtual server configured. The File Director listens on port 139 if the virtual server is configured to support NetBIOS and on port 445 if the virtual server is configured to support SMB Direct Host. If the CIFS proxy IP address is configured, the File Director listens on TCP port 139 on that IP address.
TCP and UDP 111	The File Director listens on TCP and UDP ports 111 for the <i>portmapper</i> service. This is needed to provide NFS services. The File Director listens on all IP addresses, regardless of whether or not there are any NFS virtual servers configured.

Ports	Description
TCP and UDP 1048	The File Director listens on TCP and UDP ports 1048 to provide the <i>mountd</i> service. The File Director listens on only the IP addresses of configured NFS virtual servers.
TCP and UDP 2049	The File Director listens on TCP and UDP ports 2049 to provide NFS service. The File Director listens on all IP addresses, regardless of whether or not there are any NFS virtual servers configured.
Ports for <i>statd</i>	The File Director listens on TCP and UDP ports for the NFS and <i>statd</i> services. To determine the ports of these services, query the portmap service on the File Director by using the standard UNIX <i>rpcinfo</i> command from another host.
UDP ports 137 and 138	The File Director listens on UDP ports 137 and 138 for the NetBIOS name service. This is used to provide the names of CIFS virtual servers to other hosts on the subnet using NetBIOS. To enable or disable the NetBIOS name service, configure the <i>nmbd</i> service by using the service CLI command. (For more information, see the <i>File Director CLI Reference Guide</i> .)
UDP port 161	The File Director listens on UDP port 161 on all IP addresses to provide SNMP service. To enable or disable the SNMP service, configure the <i>admin_snmp</i> service by using the service CLI command.
UDP port 123	If an NTP server has been configured in the File Director, the File Director listens on UDP port 123 on all IP addresses for NTP (Network Time Protocol) messages.

APPENDIX B

Third-Party Software License and Copyright Information

This appendix contains license and copyright information for third-party software included with the File Director.

The following license applies to software known as pysmb:

Copyright (C) 2001-2003 Michael Teo <michaelteo@bigfoot.com>

This software is provided 'as-is', without any express or implied warranty. In no event will the author be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This notice cannot be removed or altered from any source distribution.

The following license applies to software known as SNMP:

Various copyrights apply to this package, listed in 5 separate parts below. Please make sure that you read all the parts. Up until 2001, the project was based at UC Davis, and the first part covers all code written during this time. From 2001 onwards, the project has been based at SourceForge, and Networks Associates Technology, Inc hold the copyright on behalf of the wider Net-SNMP community, covering all derivative work done since then. An additional copyright section has been added as Part 3 below also under a BSD license for the work contributed by Cambridge Broadband Ltd. to the project since 2001. An additional copyright section has been added as Part 4 below also under a BSD license for the work contributed by Sun Microsystems, Inc. to the project since 2003.

Code has been contributed to this project by many people over the years it has been in development, and a full list of contributors can be found in the README file under the THANKS section.

---- Part 1: CMU/UCD copyright notice: (BSD like) ----

Copyright 1989, 1991, 1992 by Carnegie Mellon University

Derivative Work - 1996, 1998-2000

Copyright 1996, 1998-2000 The Regents of the University of California

All Rights Reserved

Permission to use, copy, modify and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appears in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of CMU and The Regents of the University of California not be used in advertising or publicity pertaining to distribution of the software without specific written permission.

CMU AND THE REGENTS OF THE UNIVERSITY OF CALIFORNIA DISCLAIM ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL CMU OR THE REGENTS OF THE UNIVERSITY OF CALIFORNIA BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM THE LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

---- Part 2: Networks Associates Technology, Inc copyright notice (BSD) ----

Copyright (c) 2001-2003, Networks Associates Technology, Inc All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

Neither the name of the Networks Associates Technology, Inc nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 3: Cambridge Broadband Ltd. copyright notice (BSD) ----

Portions of this code are copyright (c) 2001-2003, Cambridge Broadband Ltd. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

The name of Cambridge Broadband Ltd. may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDER ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 4: Sun Microsystems, Inc. copyright notice (BSD) ----

Copyright © 2003 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Use is subject to license terms below.

This distribution may include materials developed by third parties. Sun, Sun Microsystems, the Sun logo and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

Neither the name of the Sun Microsystems, Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 5: Sparta, Inc copyright notice (BSD) ----

Copyright (c) 2003-2004, Sparta, Inc All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

Neither the name of the Networks Associates Technology, Inc nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The following license applies to software known as Click:

Portions of this software have one or more of the following copyrights, and are subject to the license below. The relevant source files are clearly marked; they refer to this file using the phrase ``the Click LICENSE file". The `AUTHORS' file lists the people who have contributed to this software.

(c) 1999-2002 Massachusetts Institute of Technology

(c) 2000-2002 Mazu Networks, Inc.

(c) 2001-2002 International Computer Science Institute

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software. The names and trademarks of copyright holders may not be used in advertising or publicity pertaining to the software without specific prior permission. Title to copyright in this software and any associated documentation will at all times remain with the copyright holders.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

The following license applies to software known as Standard Template Library:

Copyright © 1996-1999

Silicon Graphics Computer Systems, Inc.

Permission to use, copy, modify, distribute and sell this software and its documentation for any purpose is hereby granted without fee, provided that the above copyright notice appears in all copies and that both that copyright notice and this permission notice appear in supporting documentation. Silicon Graphics makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

Copyright © 1994

Hewlett-Packard Company

Permission to use, copy, modify, distribute and sell this software and its documentation for any purpose is hereby granted without fee, provided that the above copyright notice appears in all copies and that both that copyright notice and this permission notice appear in supporting documentation. Hewlett-Packard Company makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

The following license applies to Linux utilities included in this product:

Copyright (c) 1989 The Regents of the University of California.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement:
This product includes software developed by the University of California, Berkeley and its contributors.
4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Index

A

- access community, SNMP 6-4
- accumulator log file, displaying 6-7
- ACL entries, configuring for NFS virtual exports 2-23
- Active node, clusters 7-1
- administrator
 - adding accounts for 2-4
 - name, entering 2-2
- attributes for policies 5-2
- audience ix
- auto-negotiation, configuring 2-8

B

- backing up the File Director configuration 6-15

C

- CIFS
 - file server access configuration, displaying 2-18
 - file server access, configuring 2-16
 - finding files or directories 4-13
 - matching shares to exports 2-18
 - overview of configuration 2-13
 - statistics, displaying 6-10
 - virtual server configuration 2-20
 - virtual share configuration 2-21
- clock, setting on File Director 2-11
- clusters
 - about 7-1
 - Active node 7-1
 - checking status of nodes 7-25
 - configuration on each node 7-5
 - disbanding 7-26
 - multicast base address, setting 7-12, 7-29
 - overview of setup 7-6
 - Primary node, specifying 7-12
 - Standby node 7-1
 - troubleshooting 7-26
- compatibility specifications of the File Director A-1
- conditions in a policy rule
 - combining 5-11
 - creating 5-9
 - deleting 5-16
 - inserting 5-9
 - nesting 5-12

- configuration of File Director
 - backing up 6-15
 - restoring 6-17
- configuring File Director for network 2-5
- conventions x
- copyright and license information, third-party software B-1

D

- date, changing 2-11
- debug file, displaying 6-7
- decryption password, entering 6-17
- default file server, specifying for virtual server 2-21
- default gateway, configuring 2-8
- deployment scenarios 1-5
- DNS server address, specifying 2-10
- domain, specifying for searching 2-10

E

- encryption password, entering 6-16
- error log file, displaying 6-6
- Ethernet interfaces, configuring on the File Director 2-8
- execution log for policies, displaying 5-20
- exports. *See* NFS

F

- features of the File Director 1-4

File Director

- accumulator log file 6-7
- backing up the configuration of 6-15
- CIFS statistics, displaying 6-10
- clusters, configuring 7-1
- compatibility specifications A-1
- configuring for network 2-5
- date and time, changing 2-11
- error log file 6-6
- Ethernet interfaces, configuring 2-8
- features of 1-4
- file servers, setting up access to 2-13
- hardware specifications of A-1
- host name, changing 2-11
- interface statistics, displaying 6-10
- introduction 1-2
- IP addresses, configuring 2-6
- monitoring 6-1
- network interfaces, configuring 2-8
- NFS and NLM statistics, displaying 6-9
- NTP server, specifying 2-12
- restoring configuration 6-17
- rolling back software on 6-20
- scenarios of deployment 1-5
- shutting down and restarting 6-15
- software, updating 6-19
- switching service log file 6-7
- system debug file 6-7
- system log file 6-6
- system services, restarting or stopping 6-13
- tasks for configuring 1-16
- TCP and UDP port specifications of A-2
- third-party license and copyright information B-1
- troubleshooting 6-11
- file migration
 - checking status of migration job 4-9
 - creating migration job 4-6
 - file server security guidelines 4-2
 - list of migration jobs 4-10
 - log of jobs 4-11
 - overview of 4-1
 - retrying failed jobs 4-11
 - statistics 4-12
- file servers
 - removing 6-14
 - setting up access in File Director 2-13
- file system monitor, described 5-6
- filtering information in Management Console 2-4
- finding migrated files or directories 4-13

G

- gateway, configuring 2-8

H

- hardware specifications of the File Director A-1
- Help, displaying and printing xi
- host name, changing 2-11

I

- installing the File Director. See the *File Director Quick Start Guide*
- interfaces
 - displaying status of 6-11
 - statistics, displaying 6-10
- introduction to the File Director 1-2
- IP addresses, configuring on File Director 2-6
- IP route, configuring for default gateway 2-8

L

- license and copyright information, third-party software B-1
- log files
 - configuring rotation 6-8
 - displaying 6-6
- logging into the File Director Management Console 2-1

M

- Management Console, using 2-1
- maximum rows displayed in Management Console, setting 2-3
- MIB file, downloading and saving 6-5
- migration
 - checking status of migration job 4-9
 - creating migration job 4-6
 - file server security guidelines 4-2
 - finding migrated files or directories 4-13
 - list of migration jobs 4-10
 - log of jobs 4-11
 - overview of 4-1
 - policies for 5-1
 - retrying failed jobs 4-11
 - statistics 4-12
- monitoring
 - File Director 6-1
 - shares or exports for policy migrations 5-6
 - multicast base address, setting for cluster 7-12, 7-29

N

- name resolution
 - DNS settings for the File Director 2-10
 - WINS settings for the File Director 2-10
- NAS file servers, problems described 1-1
- NeoPath Networks Technical Support, contacting ix
- network interfaces
 - configuring on the File Director 2-8
 - displaying status of 6-11
- NFS
 - ACL entries, configuring 2-23
 - file server access configuration, displaying 2-18
 - file server access, configuring 2-16
 - finding files or directories 4-13
 - matching exports to shares 2-18
 - overview of configuration 2-13
 - virtual export configuration 2-21
 - virtual server configuration 2-20
- NFS and NLM statistics, displaying 6-9
- NTP server, specifying 2-12

O

- online help, getting xi
- operational tasks for configuring the File Director,
 - described 1-16
- overview of the File Director 1-2

P

- ping test from the File Director 6-12
- policies for migration
 - about 5-1
 - conditions for rules, creating 5-9
 - creating 5-18
 - deleting 5-20
 - displaying execution log 5-20
 - editing 5-20
 - monitoring shares or exports 5-6
 - overview of steps 5-1
 - previewing effects of 5-19
 - rehearsing 5-19
 - rules, creating 5-8
 - schedules for 5-4
 - using 5-18
 - valid attributes of 5-2
- port specifications of the File Director A-2
- Primary node, cluster 7-12
- problems of NAS file servers, described 1-1
- proxy IP address, configuring 2-9

R

- rehearsing policies 5-19
- required tasks for configuring the File Director,
 - described 1-16
- restarting the File Director 6-15
- restoring File Director configuration 6-17
- rolling back the File Director software 6-20
- rotation of log files, configuring 6-8
- routing table from the File Director, displaying 6-12
- rules for policies
 - about 5-7
 - conditions for 5-9, 5-11, 5-12, 5-16
 - creating 5-8
 - deleting 5-16
 - editing 5-17
 - reverting to saved 5-11
 - saving 5-10

S

- scenarios of deployment 1-5
- schedules, creating and using 5-4
- scope, setting to Cluster or Local 2-7
- search domain, specifying for DNS 2-10
- services, restarting or stopping 6-13
- shares. *See* CIFS
- shutting down the File Director 6-15
- SNMP
 - about 6-1
 - access community, configuring 6-4
 - access community, deleting 6-5
 - MIB file, downloading and saving 6-5
 - trap community, configuring 6-2
 - trap community, deleting 6-4
 - trap destination, configuring 6-3
 - traps, list of 6-3
- software on File Director
 - rolling back 6-20
 - updating 6-19
- sorting information in Management Console 2-3
- specifications of the File Director A-1
- Standby node, clusters 7-1
- statistics
 - migration 4-12
 - traffic 6-9
- switching service log file, displaying 6-7

- synthetic directories
 - associating with virtual share or export 3-13
 - creating 3-8
 - deleting 3-15
 - overview of 3-1
 - scenarios of 3-2
- synthetic links
 - creating 3-9
 - overview of 3-1
 - scenarios of 3-2
- system debug file, displaying 6-7
- system log file, displaying 6-6
- system service on the File Director, restarting or stopping 6-13

T

- tasks for configuring the File Director 1-16
- TCP and UDP port specifications of the File Director A-2
- Technical Support at NeoPath Networks, contacting ix
- third-party license and copyright information B-1
- time zone, changing 2-11
- time, changing 2-11
- traffic statistics
 - CIFS statistics, displaying 6-10
 - displaying 6-9
 - interface statistics, displaying 6-10
 - NFS and NLM statistics, displaying 6-9
- trap community, SNMP 6-2
- trap destination, SNMP 6-3
- traps, list of 6-3
- troubleshooting the File Director 6-11

U

- UDP port specifications A-2
- union directories
 - creating 3-11
 - overview of 3-1
 - scenarios of 3-2
- updating the File Director software 6-19

V

- virtual servers, configuring 2-20
- virtual shares and exports
 - associating with synthetic directory 3-13
 - configuring 2-21

W

- WINS server address, specifying 2-10