

F800 Series Filer Cluster Guide

Network Appliance, Inc.
495 East Java Drive
Sunnyvale, CA 94089 USA
Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 4-NETAPP
Documentation comments: doccomments@netapp.com
Information Web: <http://www.netapp.com>

Part number 210-00341
May 2003

Copyright and trademark information

Copyright information

Copyright © 1994–2003 Network Appliance, Inc. All rights reserved. Printed in the U.S.A.

No part of this book covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Network Appliance reserves the right to change any products described herein at any time, and without notice. Network Appliance assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by Network Appliance. The use and purchase of this product do not convey a license under any patent rights, trademark rights, or any other intellectual property rights of Network Appliance.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp and the Network Appliance design are registered trademarks of Network Appliance, Inc. in the United States, Canada, and the European Union. Network Appliance is a registered trademark of Network Appliance, Inc. in Monaco and a trademark of Network Appliance, Inc. in the United States and Canada. FAServer is a registered trademark of Network Appliance, Inc. in the United States and the European Union. NetCache is a registered trademark of Network Appliance, Inc. in the European Union and Japan, and a trademark of Network Appliance, Inc. in the United States. SnapCopy is a registered trademark of Network Appliance, Inc. in the European Union and a trademark of Network Appliance, Inc. in the United States. WAFL is a registered trademark of Network Appliance, Inc. in the United States, the European Union, and Canada. FilerView, SecureShare, SnapManager, SnapMirror and SnapRestore are registered trademarks of Network Appliance, Inc. in the United States. Data ONTAP is a trademark of Network Appliance, Inc. in the United States and Canada. Snapshot is a trademark of Network Appliance, Inc. in the United States and the European Union. NetApp—the Network Appliance Company is a registered trademark of Network Appliance, Inc. in the United States and other countries. ApplianceWatch, BareMetal, Center-to-Edge, ContentDirector, DataFabric, gFiler, MultiStore, NearStore, SecureAdmin, Serving Data by Design, Smart SAN, SnapCache, SnapDrive, SnapVault, vFiler, and Web Filer are trademarks of Network Appliance, Inc. in the United States.

Network Appliance is a licensee of the CompactFlash and CF Logo trademarks.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

Table of Contents

	Preface	vii
	Safety Information (Sicherheitshinweise)	.xi
Chapter 1	Cabling a New Cluster	1
	Cluster installation types	2
	Clusters in system cabinets.	3
	Clusters in equipment racks	4
	Cabling your cluster	5
	Step 1: Preparing your equipment	7
	Step 2: Cabling the Channel A loops.	8
	Step 3: Cabling the nodes together.	11
	Step 4: Completing the Channel B loop	13
	Step 5: Completing setup and testing your cluster	15
Chapter 2	Converting an Existing Installation into a Cluster	21
	System requirements	22
	Required hardware	23
	Required software and firmware	27
	Checking your current configuration.	28
	Hardware and software updates	30
	Updating your hardware	31
	Cabling your system into a cluster	34
	Updating your software	35
Chapter 3	Configuring a Cluster	39
	Bringing up the cluster	40
	Configuring the cluster	42
	Determining current configuration.	43
	Types of interfaces used by takeover.	46
	Configuring the interface for takeover.	48
Chapter 4	Adding an Adapter or Disk Shelves to a Cluster	53
	Adding or replacing an adapter	54

	Requirements for adding a disk shelf	55
	Adding a disk shelf to a cluster that has DS14 disk shelves	57
	Step 1: Preparing to add the disk shelf	58
	Step 2: Cabling the new disk shelf	60
	Hot upgrading a DS14 disk shelf and hot swapping LRC/ESHs	61
Chapter 5	Removing Clustering Capability	65
	Removing clustering software	66
	Removing clustering hardware	67
Chapter 6	Creating Mirrored Clusters	71
	System requirements for mirrored clusters	72
	Required hardware	74
	Required software and firmware	75
	Planning your mirrored cluster	76
	Step 1: Mapping the cluster	77
	Step 2: Determining your storage requirements	78
	Step 3: Determining your FC-AL requirements	79
	Cabling your system into a mirrored cluster	80
	Step 1: Cabling the Channel A loop	81
	Step 2: Cabling the nodes together	83
	Step 3: Cabling the Channel B loop	84
	Step 4: Completing setup and testing your mirrored cluster	86
	Recovering from a disaster	90
Chapter 7	Cabling a MetroCluster	93
	System requirements	94
	Cabling a new Fabric Attached MetroCluster	97
	Step 1: Planning the Fabric Attached MetroCluster	100
	Step 2: Configuring the local switches	101
	Step 3: Cabling the local node	104
	Step 4: Cabling the remote switches and node	107
	Step 5: Configuring the nodes	111
	Cabling a non-switched MetroCluster	113
	Cabling a MetroClusters with QuickLoop	114
	Step 1: Planning the MetroCluster with QuickLoop	115

	Step 2: Installing and preparing the local node and disk shelves	118
	Step 3: Configuring the switch.	119
	Step 4: Cabling the node, disk shelves, and switch	121
	Step 5: Configuring the local node.	124
	Step 6: Setting up and configuring the remote part of the cluster	128
	Converting an existing cluster to a Fabric Attached MetroCluster.	131
	Replacing MetroCluster components.	133
	Interpreting hot swap disk loop error messages	134
Appendix A	Alternative Configuration and Cabling Examples	135
	FC-AL adapter combinations.	136
	Channel B loop cabling alternatives	138
Appendix B	Communications Regulations.	145
	Regulatory notices	146
	Declaration of Conformity	148
	Index	149

Preface

About this guide

This guide describes how to cable new or convert existing Network Appliance™ F810, F820, F825, F840, and F880 filers and Fibre Channel disk shelves into clusters. This guide does not cover basic system or network administration, which is covered in the *System Administration Guide*.

Audience

This guide is for qualified system administrators and service personnel who are familiar with storage appliances and Fibre Channel disk shelves from NetApp®—the Network Appliance Company™, which features Serving Data by Design™, and who

- ◆ Need to cable new filers and disk shelves into a clustered system
- ◆ Have two of the same filers and want to convert these filers into a cluster

Terminology

This guide uses the following terms:

- ◆ *Appliance* refers to an F810, F820, F840, or F880 filer.
- ◆ *Node* refers to a local or remote filer in an F810, F820, F840, or F880 clustered system.
- ◆ *Disk shelf* refers to a Fibre Channel StorageShelf FC7, Fibre Channel StorageShelf FC8, Fibre Channel StorageShelf FC9, or Fibre Channel DiskShelf14.
- ◆ *Mirrored cluster* refers to a cluster running the SyncMirror™ product, where the nodes are separated by between 30 and 500 meters.
- ◆ *Non-switched MetroCluster* refers to a cluster running syncmirror_local and cluster_remote, where the nodes are separated by more than 30 meters and up to 500 meters.
- ◆ *MetroCluster with QuickLoop* refers to a cluster running syncmirror_local and cluster_remote that is cabled to Fibre Channel switches, where the nodes are separated by more than 500 meters.
- ◆ *Fabric Attached MetroCluster* refers to a cluster running syncmirror_local and cluster_remote, where the nodes and disk shelves are attached to two pairs of Fibre Channel switches, where the nodes and disk shelves are not cabled together in traditional cluster fashion, and where the nodes are separated by more than 500 meters.
- ◆ *FC HBA for Disk* or *FC HBA* refers to the dual port Fibre Channel HBA adapter that connects the node to the switch or to the disks.

Command conventions

You can enter filer commands on the system console or from any client that has access to the filer using Telnet. In examples of commands run on a UNIX® workstation, this guide uses the command syntax and output of SunOS 4.1x. If you use a different version of UNIX, the command syntax and output might be different.

Keyboard conventions

This guide uses capitalization and some abbreviations to refer to the keys on the keyboard. The keys on your keyboard might not be labeled exactly as they are in this guide.

What is in this guide...	What it means...
hyphen (-)	Used to separate individual keys. For example: Ctrl-D means holding down the Ctrl key while pressing the D key.
Enter	Used to refer to the key that generates a carriage return, although the key is named Return on some keyboards.
<i>type</i>	Used to mean pressing one or more keys on the keyboard.
<i>enter</i>	Used to mean pressing one or more keys and then pressing the <i>Enter</i> key.

Formatting conventions

The following table lists the kinds of formatting this guide uses to identify special information.

Formatting convention	Type of information
<i>Italic type</i>	◆ Words or characters that require special attention. ◆ Placeholders for information you must supply. For example, if the guide requires you to enter the arp -d hostname command, you enter the characters “arp -d” followed by the actual name of the host. ◆ Book titles in cross-references.

Formatting convention	Type of information
Monospaced font	◆ Command and daemon names. ◆ Information displayed on the system console or other computer monitors. ◆ The contents of files.
Bold monospaced font	Words or characters you type. What you type is always shown in lowercase letters, unless you must type it in uppercase letters for it to work properly.

Special messages

This guide contains special messages that are described as follows:

Note

A note contains important information that helps you install or operate the system efficiently.

Caution

A caution contains instructions that you must follow to avoid damage to the equipment, a system crash, or loss of data.

WARNING

A warning contains instructions that you must follow to avoid personal injury.

Safety Information (Sicherheitshinweise)

Safety rules


You must follow these safety rules when working with this equipment:

WARNING

Failure to follow these directions could result in bodily harm or loss of life.

- ◆ When installing disk shelves and a storage appliance into a movable cabinet or rack, install from the bottom up for best stability.
- ◆ DC-based systems must be installed in a restricted access location and the two input power terminals for the DC power supply must be connected to separate isolated branch circuits.
- ◆ To reduce the risk of personal injury or equipment damage, allow internal components time to cool before touching them and ensure that the equipment is properly supported or braced when installing options.
- ◆ This equipment is designed for connection to a grounded outlet. The grounding type plug is an important safety feature. To avoid the risk of electrical shock or damage to the equipment, do not disable this feature.
- ◆ This equipment is furnished with one or more replaceable batteries. There is danger of explosion if the battery is incorrectly replaced. Replace the battery only with the same or equivalent type recommended by the manufacturer. Dispose of used batteries according to the manufacturer's instructions.

Warning for units with multiple power cords



If your storage appliance or disk shelf has multiple power cords and you need to turn the unit off, heed the following warning:

WARNING

This unit has more than one power supply cord. To reduce the risk of electrical shock, disconnect all power supply cords before servicing.

Sicherheitsvorgaben


Die folgenden Sicherheitshinweise sind beim Betreiben des Geräts unbedingt zu beachten:

Vorsicht

Nichtbeachtung dieser Anweisungen kann zu schweren Körperschäden führen oder tödlich sein.

- ◆ Bei der Montage der Diskettenregale und Archivierungsgeräte, des NetCache™ -Geräts oder des NearStore™ -Systems in bewegliche Schränke oder Regale sind die Geräte von unten nach oben einzubauen, um optimale Stabilität zu gewährleisten.
- ◆ Gleichstrom-Systeme müssen an Betriebsstätte mit beschränktem Zutritt installiert sein und die beiden Eingangsstromklemmen für das Gleichstrom-Netzteil müssen an separate und isolierte Abzweigleitungen angeschlossen sein.
- ◆ Zum Schutz vor Körperverletzung oder Sachschäden am Gerät lassen Sie die inneren Bauteile stets vor dem Berühren abkühlen. Sorgen Sie dafür, dass das Gerät richtig abgestützt ist oder fest aufrecht steht, bevor Sie neues Zubehör einbauen.
- ◆ Dieses Gerät ist für die Einspeisung aus einer geerdeten Netzverbindung ausgelegt. Der Netzstecker mit Erdungsvorrichtung ist ein wichtiger Sicherheitsschutz. Zum Schutz vor elektrischem Schlag oder Sachschäden am Gerät die Erdung nicht abschalten.
- ◆ Das Gerät ist mit einer oder mehreren auswechselbaren Batterien ausgestattet. Bei unsachgemäßem Auswechseln der Batterie besteht Explosionsgefahr. Batterien nur mit dem vom Hersteller empfohlenen Typ oder entsprechenden Typen ersetzen. Gebrauchte Batterien sind gemäß den Anweisungen des Herstellers zu entsorgen.

Warnhinweis für Geräte mit mehrfachen Netzanschlussleitungen



Sollte Ihr Archiviergerät, NetCache-Gerät, NearStore-System oder Diskettenregal mehrfache Netzanschlussleitungen aufweisen und Sie wollen das Gerät abschalten, bitte folgenden Warnhinweis beachten.

ACHTUNG

Gerät besitzt zwei Netzanschlussleitungen. Vor Wartung alle Anschlüsse vom Netz trennen.

About this chapter

This chapter describes the tasks you need to perform to cable a new standard cluster in either NetApp System Cabinets or in equipment racks.

If your configuration is not a standard cluster, go to one of the following chapters for information about creating, cabling and configuring mirrored clusters or MetroClusters.

- ◆ Chapter 6, “[Creating Mirrored Clusters](#),” on page 71
- ◆ Chapter 7, “[Cabling a MetroCluster](#),” on page 93

Topics in this chapter

The following topics are covered in this chapter:

- ◆ “[Cluster installation types](#)” on page 2
- ◆ “[Cabling your cluster](#)” on page 5

Cluster installation types

About cluster installation types

Depending on the requirements of your data center, you can order a clustered system in one or more NetApp System Cabinets, or as a clustered system that you install in a telco-style equipment rack.

For detailed information

The cluster types are described in the following sections:

- ◆ “[Clusters in system cabinets](#)” on page 3
- ◆ “[Clusters in equipment racks](#)” on page 4

Note

If your configuration is not a standard cluster, go to Chapter 6, “[Creating Mirrored Clusters](#),” on page 71, for information about creating, cabling and configuring mirrored clusters, or to Chapter 7, “[Cabling a MetroCluster](#),” on page 93 for information about creating, cabling and configuring MetroClusters.

Cluster installation types

Clusters in system cabinets

How your cluster arrives

If you ordered a cluster in a system cabinet, it comes to you in one or more NetApp System Cabinets. The number of system cabinets you receive depends on how many disk shelves you ordered. All internal adapters, such as networking adapters, Fibre Channel adapters, and other adapters, arrive preinstalled in your nodes.

If you ordered a cluster with six or fewer disk shelves, you receive it in a single system cabinet. This system cabinet comes to you with the Channel A and Channel B disk shelves cabled, and also has the appropriate cluster interconnect adapters cabled. All you need to do is connect the cluster to the network, apply power to the system, and configure your cluster. See the *F800 Series Hardware Guide* and the *42U System Cabinet Guide* for information about connecting the nodes to the network and applying power to the cluster. See the *System Administrator’s Guide* for cluster software configuration information.

If you ordered a cluster with more than six disk shelves, you receive your cluster in two or more system cabinets. You must cable the local node to the partner node’s disk shelves and the partner node to the local node’s disk shelves. You must also cable your nodes together in a cluster. See the *42U System Cabinet Guide* for information about connecting your system cabinets together for clustering. See the *F800 Series Hardware Guide* for information about connecting the nodes to the network and see the *System Administrator’s Guide* for more configuration information.

Basic system cabinet cluster configurations

The following table lists the different combinations of nodes in a cluster and how many disk shelves you can have in each system cabinet.

Number of nodes in the system cabinet	Disk shelves below the node	Disk shelves above the node	Total number of disk shelves in a cabinet
0	N/A	N/A	10 (nodes in other system cabinet)
1	4	4	8
2	3	3	6

Clusters in equipment racks

How your cluster arrives

Depending on the amount of storage you ordered, you need to install the equipment in one or more telco-style equipment racks. These racks can hold one or two nodes on the bottom and eight or more disk shelves. For information about how to install the disk shelves and nodes into the equipment racks, see the appropriate Fibre Channel disk shelf guide and the *F800 Series Hardware Guide*.

The basic configuration you receive with your cluster is a minimum of two Fibre Channel Arbitrated Loop (FC-AL) adapters or one dual Fibre Channel adapter per node. You can also add additional FC-AL or dual Fibre Channel adapters.

The dual Fibre Channel adapter has two ports: one for Channel A and one for Channel B. This adapter enables you to use one FC-AL adapter per node for both the Channel A and Channel B loops, instead of one adapter per channel. Using the dual Fibre Channel adapter in a cluster can free as many as two PCI slots each in your local and partner nodes.

Your node also arrives with one cluster interconnect adapter or one VI cluster adapter per node. These adapters enable you to cluster your nodes together.

Basic cluster configurations

The following table lists recommended slot locations for F800 series clusters with FC-AL adapters or with dual Fibre Channel adapters, in a nonmirrored configuration. Check the NOW Web page at <http://now.netapp.com> to confirm the slot assignment information.

Cluster type	Channel A	Channel B
Cluster with two FC-AL adapters per node	Slot 8	Slot 7
Cluster with four FC-AL adapters per node	Slot 8 and slot 5	Slot 7 and slot 9
Clusters with single dual Fibre Channel adapter per node	Slot 8 (Channel B is on the same adapter)	N/A
Clusters with multiple dual Fibre Channel adapter per node	Slot 8, slot 7, slot 5, and slot 9	N/A

Cabling your cluster

Required tools and manuals

You need the following tools and manuals to prepare your system cabinets or equipment racks for clustering:

- ◆ The *42U System Cabinet Guide*
- ◆ The *F800 Series Hardware Guide*
- ◆ The appropriate Fibre Channel disk shelf guide
- ◆ #1 and #2 Phillips screwdrivers

Required cables

When you received your system cabinets or nodes and disk shelves for installation in equipment racks, you should have received the disk shelf cables listed in the following table.

Cables	DS14	DS14mk2 FC
Filer-to-shelf	HSSDC to DB-9 SC/LC or LC/LC	SC/LC or LC/LC
Shelf-to-shelf	DB-9 to DB-9 HSSDC2-to-HSSDC2	HSSDC2-to-HSSDC2
Shelf-to-switch	LC/SC	LC/SC

Note

See the *DiskShelf14 and DiskShelf14mk2 FC Hardware Guide* for detailed information about the differences between LRC and ESH modules.

You should also receive cluster interconnect cables, one pair for the cluster interconnect adapter or one optical cable for the VI cluster adapter.

Note

If your cluster will be farther apart than 30 meters or more than 30 meters to a switch, when in a MetroCluster, you must use a SC-SC Duplex 50/125 plenum cable to connect them.

For detailed information

For details about the cabling process, see the following topics:

- ◆ [“Step 1: Preparing your equipment”](#) on page 7
- ◆ [“Step 2: Cabling the Channel A loops”](#) on page 8
- ◆ [“Step 3: Cabling the nodes together”](#) on page 11
- ◆ [“Step 4: Completing the Channel B loop”](#) on page 13
- ◆ [“Step 5: Completing setup and testing your cluster”](#) on page 15

Step 1: Preparing your equipment

Setting up the system cabinets

To prepare your system cabinets for clustering, complete the following steps.

Step	Action
1	Remove the front and rear doors and side panels from the system cabinets, as described in the <i>42U System Cabinet Guide</i> . You should only remove the panels from the sides that will be connected. If you have a single system cabinet, go to Step 3.
2	Connect the system cabinets. See the <i>42U System Cabinet Guide</i> .
3	Connect the nodes to your network. See the <i>F800 Series Hardware Guide</i> .
4	Go to section “ Step 3: Cabling the nodes together ” on page 11.

Preparing your equipment racks

To prepare your new equipment for clustering, complete the following steps.

Step	Action
1	Install both nodes in the equipment rack. See the <i>F800 Series Hardware Guide</i> .
2	Label the FC-AL interfaces, where appropriate, then install the disk shelves in the equipment rack, then set the disk shelf IDs. See the appropriate disk shelf guide for your Fibre Channel disk shelf.
3	Connect the nodes to your network. See the <i>F800 Series Hardware Guide</i> .
4	Go to section “ Step 2: Cabling the Channel A loops ” on page 8.

Step 2: Cabling the Channel A loops

Channel A loop cabling exceptions

Generally, the Channel A loop in a cluster is cabled by connecting the FC-AL adapter in the node to the LRC/ESH Input port of the first disk shelf in the loop, followed by cabling the disk shelves together. This process differs only by the type of Fibre Channel adapter you have in your cluster and the type of disk shelves you have.

Note

If you are cabling a cluster in a system cabinet, skip this Step and go to [“Step 3: Cabling the nodes together”](#) on page 11. The Channel A loop should be cabled.


The following table lists exceptions to the general cabling process for the Channel A loop.

If you have...	Then...
Nodes using dual Fibre Channel adapters with any disk shelf type	◆ You use optical cables to connect the nodes to the first disk shelf in the Channel A loop. ◆ The first port (labeled Port 1 on the adapter) of the dual Fibre Channel adapter is reserved for the Channel A loop.
DS14 or DS14mk2 FC disk shelves with nodes using copper FC-AL adapters	You use the HSSDC to DB-9 cable to connect the nodes to the first disk shelf in the Channel A loop.
DS14 or DS14mk2 FC disk shelves with ESH modules	All disk shelves must have ESH modules installed in them and you must not have mixed shelf-type loops. Use the supplied cables for your FC-AL adapters.
FC9 disk shelves with nodes using dual Fibre Channel adapters	Your first disk shelf LRC requires an MIA (media interface adapter) plugged into the Input port.

Cabling the Channel A loops

To cable the Channel A loops, complete the following steps, using the diagram for reference.

Note All diagrams in this book show DS14 disk shelves. If you have other types of disk shelves, the Input and Output ports might be in different locations.



Step	Action
1	Cable the FC-AL adapter in Slot 8, Port 0, to the disk shelf 1 LRC/ESH Input port.
2	Cable the disk shelf 1 Channel A LRC/ESH Output port to the disk shelf 2 Channel A LRC/ESH Input port. Apply solid-colored labels to the cables.
3	Repeat Step 2, connecting output to input, for the remaining disk shelves in the loop.

Step	Action	
4	Set the terminate switches to Off on all but the last disk shelf in the loop, then set the terminate switch on the last disk shelf to On.	
5	If you have...	Then...
	A mix of DS14 and DS14mk2 FC disk shelves with LRC modules	Set the 1/2Gb switch to 1Gb.
	A mix of DS14 and DS14mk2 FC disk shelves with ESH modules	Set the 1/2Gb switch to 1Gb.
	Only DS14mk2 FC disk shelves with ESH modules	Set the 1/2Gb switch to 2Gb.
6	Repeat Steps 1 through 5 for the partner node.	
7	Go to “ Step 3: Cabling the nodes together ” on page 11.	

Step 3: Cabling the nodes together

Cluster cabling differences


Your cluster arrives with one of three different types of cluster adapters, depending on the configuration you are installing. The following table lists the different supported cluster adapters and the major differences between them.

Cluster adapter type	General characteristics
Cluster interconnect adapter	◆ No longer shipped with new systems. ◆ The adapter does not support mirrored clusters.
IB cluster adapter	◆ Uses an optical cable up to 500 meters long. ◆ The adapter can be used in standard or mirrored clusters. ◆ Requires two SC/SC cables and GBICs for distances over 30 meters. These cables are not included.
VI cluster adapter (Rev B)	◆ Uses an optical cable up to 500 meters long. ◆ Can be used in mirrored clusters and non-switched MetroClusters. ◆ Requires an LC/LC cable for distances over 30 meters. This cable is not included.
VI-MC cluster adapter	◆ Used for long-distance Fabric Attached MetroClusters. ◆ Connects directly to Fibre Channel switches. ◆ Requires two SC/SC cables. These cables are not included. See Chapter 7, “Cabling a MetroCluster,” on page 93 for more information.

See the System Configuration Guide on the NOW Web page at <http://now.netapp.com> for current information about supported cluster adapters for specific filers.

Cabling the VI cluster adapter

To cable the VI cluster adapters, complete the following steps, using the diagram for reference.



Step	Action
1	Plug one end of the optical cable into the GBIC of the local node’s VI cluster adapter, then plug the other end into the partner node’s VI cluster adapter. Note_____ You must not cross-cable the cluster interconnect adapter. Cable the local node X port only to the partner X port and cable only the Y ports together. _____
2	Go to “ Step 4: Completing the Channel B loop ” on page 13.

Cabling the IB cluster adapter

To cable the IB cluster adapters, complete the following steps.

Step	Action
1	Plug one end of the optical cable into one of the local node’s IB cluster adapter ports, then plug the other end into the corresponding partner node’s IB cluster adapter port. Repeat for the second cable.
2	Gently pull on the cable ends to ensure that they are properly seated.
3	Go to “ Step 4: Completing the Channel B loop ” on page 13.

Step 4: Completing the Channel B loop

Channel B loop cabling exceptions


Generally, the Channel B loop in a cluster is cabled by connecting the FC-AL adapter in the local node to the LRC/ESH Input port of the first disk shelf in the partner node's Channel B loop, followed by cabling the disk shelves together. This process differs only by the type of Fibre Channel adapter you have in your cluster and the type of disk shelves you have.

The following table lists exceptions to the general cabling process for the Channel B loop.

If you have...	Then...
Nodes using dual Fibre Channel adapters with any disk shelf type	◆ You use optical cables to connect the nodes to the first disk shelf in the Channel A loop. ◆ The second port (labeled Port 1 on the adapter) of the dual Fibre Channel adapter is reserved for the Channel B loop.
DS14 or DS14mk2 FC disk shelves with nodes using copper FC-AL adapters	You use the HSSDC to DB-9 cable to connect the nodes to the first disk shelf in the Channel A loop.
DS14 or DS14mk2 FC disk shelves with ESH modules	All disk shelves must have ESH modules installed in them and you must not have mixed shelf-type loops. Use the supplied cables for your FC-AL adapters.
FC9 disk shelves with nodes using dual Fibre Channel adapters	Your first disk shelf LRC requires an MIA (media interface adapter) plugged into the Input port.

Cabling the Channel B loops

To cable the Channel B loops, complete the following steps, using the diagram for reference.



Step	Action
1	Cable the FC-AL adapter in Slot 8, Port 1 of the local node, to the partner node disk shelf 1 LRC/ESH Input port.
2	Cable the local node's disk shelf 1 Channel B LRC/ESH Output port to the local node's disk shelf 2 Channel B LRC/ESH Input port. Label this cable with a striped label.
3	Repeat Step 2, connecting output to input, for the remaining disk shelves in the loop.
4	Set the terminate switches to Off on all but the last disk shelf in the loop, then set the terminate switch on the last disk shelf to On.
5	Repeat Steps 1 through 4 for the partner node.
6	Go to “Step 5: Completing setup and testing your cluster” on page 15.

Step 5: Completing setup and testing your cluster

Completing clustered system setup

Your cluster is ready to be cabled to the network and to power sources. See the *F800 Series Hardware Guide* for more information.

If your cluster is in systems cabinets, you must reinstall the front and rear doors of your clustered system cabinets to comply with FCC noise reduction requirements. See the *42U System Cabinet Guide* for more information about door installation.

Testing takeover and giveback on your cluster

To test your configured cluster, complete the following steps.

Step	Action
1	Verify the network addresses by entering the following command from one node: ifconfig -a Result: The local node and partner node IP addresses are listed.
2	Verify that you can create and retrieve files on both nodes for each licensed protocol.
3	Enable the clustering licenses on by entering the following command on the local and partner nodes: license add xxxxx xxxxx is the license code for your cluster. Note _____ If you are using a version of Data ONTAP earlier than 6.4, the command is license option=code . _____ Result: Clustering is enabled on your system.

Step	Action
4	From the same node, enter the following command: cf enable Result: Takeover of the partner node is enabled
5	Run the cluster configuration checker (cf-config-check.cgi) script, as described in “Using the cluster configuration checker” on page 43.

Step	Action
6	Test communication between the local node and partner node. For example, you can use the <code>fcadmin device_map</code> command to ensure that one node can access the other node's disks. Example: <pre>filer-1> fcadmin device_map Loop Map for channel 8: Translated Map: Port Count 32 7 16 17 18 19 32 33 34 35 48 49 50 51 24 26 27 57 61 42 43 44 45 28 29 30 52 53 54 37 38 21 22 Shelf mapping: Shelf 2: 22 21 XXX 19 18 17 16 Shelf 3: 30 29 28 27 26 XXX 24 Shelf 4: 38 37 XXX 35 34 33 32 Shelf 5: XXX 45 44 43 42 XXX XXX Shelf 6: 54 53 52 51 50 49 48 Loop Map for channel 5: Translated Map: Port Count 22 7 38 35 22 20 19 30 29 28 27 14 41 42 43 26 25 24 18 16 34 33 32 Shelf mapping: Shelf 2: 22 XXX 20 19 18 XXX 16 Shelf 3: 30 29 28 27 26 25 24 Shelf 4: 38 XXX XXX 35 34 33 32 Shelf 5: XXX XXX XXX 43 42 41 XXX Loop Map for channel 7a: Translated Map: Port Count 10 7 40 42 43 44 45 46 58 59 60 Shelf mapping: Shelf 5: 46 45 44 43 42 XXX 40 Shelf 7: XXX XXX 60 59 58 XXX XXX Loop Map for channel 7b: Translated Map: Port Count 13 7 16 17 18 19 20 32 33 34 35 36 37 38 Shelf mapping: Shelf 2: XXX XXX 20 19 18 17 16 Shelf 4: 38 37 36 35 34 33 32</pre>

Step	Action
7	From the local node console, enter the following command and verify that the partner node halted: cf takeover Result: The local node takes over the partner node. You see a message similar to the following: takeover completed
8	Make sure that all clients and client applications can access data.
9	From the same node, enter the following command: cf giveback Result: The partner node reboots and functions normally. A successful giveback ends with the following message on the local node: giveback completed
10	Repeat Steps 1 through 8 on the other node.

If takeover or giveback fails

If takeover or giveback fails, complete the following steps.

Step	Action
1	Check the Channel B cabling on your disk shelves, and tighten any loose cables.
2	Check the cabling to your cluster adapters, making sure that the cabling is correct, that the cables are properly seated in the adapters, and that the cables are secured to the adapters.
3	Check the communication between the local node and partner node by entering the following command: cf status Result: If there is a cabling error or the X or Y link is down, the system displays a message that tells you where the problem is.

Step	Action
4	Check the licenses on the local and partner nodes. Both nodes must have the cluster license enabled. Add missing licenses.
5	Check the parameters and options on your cluster with the cluster configuration checker script. The script is available through download from the NOW Web site at http://now.netapp.com/NOW/tools. Look for the cluster configuration checker (cf-config-check.cgi) file. Correct any errors you find.
6	Check for network connectivity. See the <i>System Administrator's Guide</i> for more information.
7	Reboot the cluster and rerun the takeover and giveback tests.
8	If you still do not have clustering enabled, call Customer Support.

Where to go next

After cabling your nodes and disk shelves into a cluster, continue configuring your cluster as described in Chapter 3, “[Configuring a Cluster](#),” on page 39. Also see the *First-Time Configuration Guide* for more detailed information about configuring your new cluster.

About this chapter

This chapter describes the tasks you need to perform to convert existing filers into a clustered system. It contains the following information:

- ◆ What the minimum hardware requirements are for converting your filers
- ◆ What the minimum software and firmware version requirements are for converting your filers

How to determine your current filer configuration and software and firmware versions

- ◆ How to update the firmware and software on your hardware
- ◆ How to cable your filers into a cluster

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“System requirements”](#) on page 22
- ◆ [“Hardware and software updates”](#) on page 30
- ◆ [“Cabling your system into a cluster”](#) on page 34

System requirements

About your system requirements

Your current system configuration determines what you need to modify before cabling your individual filers into a clustered system. You should determine the following information:

- ◆ What type of hardware is currently in your filers and what type of disk shelves you have
- ◆ What version of Data ONTAP™ software your filers are running
- ◆ What versions of filer and disk and shelf firmware are on your filers
- ◆ What new hardware you are adding to your filers
- ◆ What version of Data ONTAP, filer firmware, and disk shelf firmware you need to support the new hardware
- ◆ Whether synchronous mirroring will be used
- ◆ The distance between the local and partner nodes. If they are more than 30 meters apart, you should install the cluster_remote license, as described in Chapter 6, “[Creating Mirrored Clusters](#),” on page 71. If the nodes are more than 500 meters apart, you must configure your cluster as described in Chapter 7, “[Cabling a MetroCluster](#),” on page 93.

For detailed information

For details about how to determine what you currently have in your filers and what you might need to update, see the following sections:

- ◆ “[Required hardware](#)” on page 23
- ◆ “[Required software and firmware](#)” on page 27
- ◆ “[Checking your current configuration](#)” on page 28

Required hardware

Hardware required for clustering

The following tables list all the hardware required to convert your filers into a cluster. What you have in your individual filers determines what you should add for clustering. See the NOW Web site at <http://now.netapp.com> for the most current hardware requirements.

Node hardware requirements

The following table lists the minimum and maximum number of Fibre Channel adapters and cluster interconnect adapters used to cluster your filers.

Required hardware	Minimum configuration per node	Maximum configuration per node
Fibre Channel adapters		
◆ FC-AL adapters	2	4
◆ FC-AL adapters—mirrored cluster	4	4
◆ Dual Fibre Channel adapters (each has a Port 1 for Channel A and a Port 2 for Channel B)	1	2
◆ Dual Fibre Channel adapters—synchronous mirroring for clusters	2	4
Cluster adapters:	1	1
◆ Cluster interconnect adapters (each has two ports: an X port and a Y port)		
◆ IB cluster adapter		
◆ VI cluster adapter (Rev. B)		
◆ VI-MC cluster adapter for MetroClusters		
Note _____ Do not mix cluster adapter types.		

Disk shelf hardware requirements

The following table lists the minimum and maximum hardware configurations for your disk shelves. See the NOW Web page at <http://now.netapp.com> for more information about disk shelf capacity.

Required hardware	Minimum	Maximum
LRCs: FC7, FC8, or FC9 disk shelves DS14 disk shelves ESH: DS14 or DS14mk2 FC	4 2 2	48 32 32 standard and 48 for mirroring
Disk shelves of the same type Mix of FC7, FC8, FC9, and DS14 disk shelves Caution _____ FC7 disk shelves cannot use 18-GB, 36-GB, or 72-GB disks. _____	2 2	The number of disk shelves for both nodes must be no more than the maximum storage capacity for the nodes.

Cable requirements

The following table lists the minimum and maximum number of cables required for the various adapters you might have in your nodes. It also lists disk shelves requiring terminators.

Required hardware	Minimum	Maximum
4m cluster interconnect cables for cluster interconnect adapters	2	2
or		
Optical cables for clusters separated by no more than 30 meters.		
IB cluster adapter (with 2 GBICs)	2 (SC/SC)	2
VI cluster adapter	1 (LC/LC)	1
or		
SC-SC Duplex 50/125 plenum cable for VI cluster adapters separated by more than 30 meters and less than 500 meters.		
Optical cables for dual Fibre Channel adapters	4	8
2m filer-to-shelf (HSSDC to DB-9) cables for FC-AL adapters	4	8
Note _____ If your cluster is separated by more than 30 meters, you must use a SC-SC Duplex 50/125 plenum cable to connect the disk shelves from the local node to partner's disk shelves and vice-versa. _____		
0.5m shelf-to-shelf (DB-9 to DB-9) cables for DS14 disk shelves	2	You need $2(n - l)$ number of cables, where n is the number of disk shelves and l is the number of loops.
or		
LC-SC or LC-LC cables		

Required hardware	Minimum	Maximum
Terminators for FC7 and FC8 disk shelves only	2	4

Required software and firmware

Software requirements

You require the following software to convert your filer to a cluster:

- ◆ Data ONTAP 6.1.2 or later for F810 clusters
- ◆ Data ONTAP 6.0.1 or later for F820 clusters
- ◆ Data ONTAP 6.3 or later for F825 clusters
- ◆ Data ONTAP 5.3.6H1 or later for F840 clusters
- ◆ Data ONTAP 6.1.1 or later for F880 clusters
- ◆ Data ONTAP 6.1 or later for dual Fibre Channel adapters or VI cluster adapters
- ◆ Data ONTAP 6.2 or later for mirrored clusters
- ◆ Data ONTAP 6.4 or later for MetroClusters with QuickLoop
- ◆ Data ONTAP 6.4.1 or later for Fabric Attached MetroClusters

Note

The software is on a CD-ROM for new systems, or you can download it from <http://now.netapp.com>. To download the software, you need network connectivity and access to a UNIX or PC system.

Firmware requirements

Your filers and disks drives require the following firmware levels for cluster conversion. If your disk type is not listed, contact your Customer Support representative or see the NOW Web page at <http://now.netapp.com> for the current firmware version.

Firmware type	Version
Filer	See the NOW site at http://now.netapp.com/NOW/tools/releasetable.shtml for current filer firmware.
Disk	See the NOW site at http://now.netapp.com/NOW/download/tools/diskfw for current disk firmware and for a list of supported disk drives and associated disk firmware.

Checking your current configuration

Reasons for checking your system information

Certain versions of your hardware or software might require updating before you can convert your filers to a clustered system. You must determine what hardware, software, and firmware exists on your filers, then check them against the requirements for converting your filers into a cluster. See the NOW Web page at <http://now.netapp.com> for the most current version requirements for clustering.

Commands for gathering cluster information

There are two types of information you should gather: hardware and operating system information and networking information. You can gather this information by using the `sysconfig` and `ifconfig` commands on both filers you are converting.

sysconfig command: The `sysconfig -v` command displays your system hardware configuration. See the *System Administrator's Guide* for more information about this command. When you run this command, you should note the following information displayed in the output:

- ◆ Data ONTAP version currently installed on your filers
- ◆ Contents of each PCI slot in your filers
- ◆ Firmware versions on your filers and disk shelves

ifconfig command: You should also note your current network interface settings, using the `ifconfig -a` command. See the *System Administrator's Guide* for more information about this command. You should note the following information:

- ◆ Interface IP addresses
- ◆ Installed interface types
- ◆ Network IP addresses that might be needed

Note

If you want to implement trunking on your cluster, see the *System Administrator's Guide* for a discussion and procedures for setting up virtual interfaces (vifs).

Checking your system configuration

To check your system information, complete the following steps.

Step	Action
1	At the console of each filer, enter the following command: <pre>sysconfig -v</pre> Result: The filer configuration is displayed, similar to the following: <pre>ha15> sysconfig -v NetApp Release 6.4: Wed Dec 11 01:33:09 PST 2002 System ID: 0033554550 (metro1); partner ID: 0033554438 (ha16) System Serial Number: 6080 (metro2) System Rev: A0 Backplane Part Number: 110-02405 Backplane Rev: A0 Backplane Serial Number: 6080 slot 0: System Board 866 MHz (NetApp System Board VII A0) Model Name: F880 Part Number: 1102460 Revision: A0 Serial Number: 244233 Firmware release: 4.0_i1 Processors: 2 Memory Size: 3072 MB CIOB Revision ID: 3 slot 0: FastEnet-10/100 Ethernet Controller Device Type: 82559 e0 MAC Address: 00:a0:98:00:a1:36 (auto-100tx- fd-up) . . .</pre>
2	Compare the information gathered from Step 1 with the information listed in “System requirements” on page 22 and against the configuration information listed on the NOW Web site. Determine what you need to update, then go to “Hardware and software updates” on page 30.

Hardware and software updates

About updating your hardware

After you determine your filer configuration and the software and firmware levels of your hardware, you might need to update those items that are out of date.

Additional hardware adapters you might need

Depending on the number of disk shelves you have or might add, you might also need additional FC-AL adapters or dual Fibre Channel adapters. You also need to add one cluster interconnect adapter, VI cluster adapter, or IB cluster adapter to each filer. See the *F800 Series Field Service Guide* for more information about adding adapters to your filers.

If you are upgrading your DS14 disk shelves, see Chapter 4, “[Adding an Adapter or Disk Shelves to a Cluster](#),” on page 53 for information.

Required tools and manuals

You need the following tools and manuals to convert your filers into a cluster:

- ◆ The *F800 Series Hardware Guide*
- ◆ The *F800 Series Field Service Guide*
- ◆ The *System Administrator's Guide*
- ◆ The appropriate disk shelf hardware guide
- ◆ #1 and #2 Phillips screwdrivers

For detailed information

For detailed information about updating your system configuration and current software and firmware versions, see the following sections:

- ◆ “[Updating your hardware](#)” on [page 31](#)
- ◆ “[Cabling your system into a cluster](#)” on [page 34](#)
- ◆ “[Updating your software](#)” on [page 35](#)

Updating your hardware

About updating your hardware

When updating your hardware, you might need to add the following to your nodes and disk shelves:

- ◆ One or more FC-AL or dual Fibre Channel adapters to your nodes
- ◆ One cluster interconnect adapter per node. See “[Cluster cabling differences](#)” on page 11 for a description of the cluster interconnect adapters
- ◆ Additional LRCs for your FC7, FC8, or FC9 disk shelves
- ◆ Additional disk shelves to your system
- ◆ Two ESH modules per disk shelf to replace the LRC modules, if upgrading DS14 or DS14mk2 FC disk shelves with LRC modules

You might have to move or remove other adapters when adding the appropriate FC-AL adapters. See the *F800 Series Field Service Guide* for information about installing additional FC-AL adapters.

Before installing additional hardware into your nodes or disk shelves, check the NOW Web page at <http://now.netapp.com> for current configuration information.

Mixed disk shelf environment requirements

When converting your filers into a cluster, you can add new disk shelves to your existing disk shelf loops. In certain instances, you add newer disk shelf types to older disk shelf types.

To successfully add newer disk shelves to older, existing disk shelves, you must be aware of the following requirements:

- ◆ If you have both FC9 and FC7 disk shelves, you must assign an FC7 disk shelf as disk shelf 0, the first physical disk shelf in the loop.
- ◆ If you are converting filers that have both DS14 and FC9 or FC8 disk shelves, you must assign a DS14 disk shelf as the last physical disk shelf in the Channel A and B loops.
- ◆ If you have both DS14 and FC7 disk shelves, you must assign an FC7 disk shelf as disk shelf 0, the first physical disk shelf in the loop, and a DS14 disk shelf as the last physical disk shelf in the Channel A and B loops.
- ◆ If you are converting filers that have FC9 disk shelves and want to use dual Fibre Channel adapters, you must use an MIA (media interface adapter) in the Channel A and B Input ports of the first disk shelf in each loop.

- ◆ You cannot use dual Fibre Channel adapters with FC7 or FC8 disk shelves.
- ◆ If you have a DS14/DS14mk2 FC with ESH modules, all your disk shelves in that loop must have ESH modules.

Installing additional hardware

To install the hardware you need for clustering, complete the following steps.

Note

You should run diagnostics for any new components you add to your cluster. See the *Diagnostics Guide* for more information.

Step	Action
1	Shut down the filers by entering the following command: halt
2	Ground yourself using the antistatic grounding strap, then turn off the power to the filers and disk shelves and unplug them.
3	Install and label additional FC-AL adapters or dual Fibre Channel adapters, as described in the <i>F800 Series Field Service Guide</i> . Make sure to label the FC-AL adapters appropriately. See the System Configuration Guide for information about the maximum number of drives supported by a loop.
4	Install the appropriate cluster interconnect adapter, as described in the <i>F800 Series Field Service Guide</i> .

Step	Action	
5	If you have...	Then...
	DS14 or DS14mk2 FC disk shelves with LRCs and you are converting to ESH modules	Replace the LRC modules with ESH modules. See the <i>DiskShelf14 and DiskShelf14mk2 FC Hardware Guide</i> for installation instructions.
	DS14 or DS14mk2 FC disk shelves and are not changing LRC for ESH modules	Go to Step 6.
	FC7, FC8, or FC9 disk shelves	Install any LRCs, as needed. If you have an FC7, you might have to label your LRCs. If the LRCs are not labeled, label them as follows: ◆ Label the LRC above the EDM card as Channel A. ◆ Label the LRC below the EDM card as Channel B.
6	Repeat Steps 3 through 5 on the partner node.	
7	Go to “ Updating your software ” on page 35.	

Cabling your system into a cluster

About cabling your system into a cluster

After you install any additional adapters or disk shelves, you must cable your filers into a cluster. If your filers have been in service, your Channel A loops are already cabled. All you need to do is cable your filers together and cable the Channel B loops. After cabling, you should update the software and firmware on your cluster. It should then be configured as described in Chapter 3, “[Configuring a Cluster](#),” on page 39, and as described in the *System Administrator's Guide*.

Cabling your converted filers

To cable your converted filers into a cluster, complete the following steps.

Step	Action
1	Cable your Channel A loops, as needed, as described in “ Step 2: Cabling the Channel A loops ” on page 8.
2	Cable the cluster interconnect adapters to both nodes of the cluster, as described in “ Step 3: Cabling the nodes together ” on page 11.
3	Cable your Channel B loops, as described in “ Step 4: Completing the Channel B loop ” on page 13.
4	Test your cabled cluster, as described in “ Step 5: Completing setup and testing your cluster ” on page 15.
5	Configure your cluster, as described in Chapter 3, “ Configuring a Cluster ,” on page 39.

Updating your software

About updating your software

Updating your software is a two-part process:

- ◆ Updating Data ONTAP and filer firmware
- ◆ Updating the firmware for your disks
 - ❖ For Data ONTAP 6.1 and earlier
 - ❖ For Data ONTAP 6.1.1 and later

Updating Data ONTAP and filer firmware

To update your Data ONTAP and filer firmware, complete the following steps.

Step	Action
1	Back up all your data.
2	Compare the version of Data ONTAP and filer firmware to the most current version of both pieces of software. If they are current, go to “Cabling your system into a cluster” on page 34; otherwise, continue with this procedure.
3	Verify that both nodes have network connections.
4	Download and install any filer firmware updates you need.
5	Update each node to the latest version of Data ONTAP. See the <i>System Administrator’s Guide</i> for more information.
6	Go to “Updating disk firmware, Data ONTAP 6.1 and earlier” on page 36.

Updating disk firmware, Data ONTAP 6.1 and earlier

To update disk firmware for Data ONTAP 6.1 and earlier, complete the following steps.

Note

Only disks that have firmware files present in the `/etc/disk_fw` directory are updated during firmware updates.

Caution

Do not update disk firmware on active nodes. Doing so might cause client drop-off. See the `disk_fw_update` man page for more information.

Step	Action
1	Disable your cluster by entering the following command on the local node: cf disable
2	Load any Data ONTAP updates, as described in “ Updating your software ” on page 35.
3	Shut down the partner node by entering the following command from the partner console: halt
4	Install the new disk firmware by entering the following command: disk_fw_update
5	Halt the local node, then boot the partner node. Repeat Steps 3 and 4 on the partner node.
6	Reboot the local node and reboot the partner node by entering the following command on the console of the local and partner nodes: reboot
7	Reenable clustering by entering the following command: cf enable

Updating disk firmware, Data ONTAP 6.1.1, 6.0.3, and later

To update disk firmware for Data ONTAP 6.1.1, 6.0.3, and later, complete the following steps.

Note

Only disks that have firmware files present in the `/etc/disk_fw` directory are updated during firmware updates.

Caution

Do not update disk firmware on active filers. Doing so might cause client drop-off. See the `disk_fw_update` man page for more information.

Step	Action
1	Load any Data ONTAP updates to the local and partner node, as described in “ Updating your software ” on page 35.
2	Reboot the local node and reboot the partner node by entering the following command on the console of the local and partner nodes: reboot
3	Install the new disk firmware on the local node by entering the following command: disk_fw_update
4	Power cycle the disk shelves.
5	After the local node completes loading disk firmware, install the disk firmware on the partner node.

About this chapter

This chapter describes the tasks you need to perform after cabling your cluster:

- ◆ Booting and licensing your filers so that they can be brought up in a cluster
- ◆ Configuring your cluster by setting the necessary system parameters, options, and networking parameters for your cluster
- ◆ Testing your cluster to make sure that the filer takeover succeeds

Topics in this chapter

This chapter describes the following procedures:

- ◆ [“Bringing up the cluster”](#) on page 40
- ◆ [“Configuring the cluster”](#) on page 42

Bringing up the cluster

Bringing up the new cluster

After cabling your cluster, you need to bring up the cluster so that you can configure it.

To bring up the cluster, complete the following steps.

Step	Action
1	Turn on the power to each disk shelf in the cluster.
2	Turn on the power to both filers and press the Delete key during the memory test.
3	From the prompt, run the diagnostics program by entering the following command: sys-diag Result: The diagnostics program loads. You should run the diagnostics that check your FC-AL adapters and cluster interconnect adapter. See the <i>Diagnostics Guide</i> for more information.
4	At the prompt on each console, enter the following command: boot Result: The filer begins the boot process, booting to the system prompt.
5	Check the status of your cluster by entering the following command on each console: cf status Result: You should see the following output: cluster monitor not initialized.
6	Enable the cluster licenses on your filers, as described in “ Enabling and licensing clustering on your filers ” on page 41.

Enabling and licensing clustering on your filers

When you purchased your clustered system, you also purchased software licenses for the system. For clustering, you must have the cluster license enabled on both filers. You can also enable other purchased licenses, such as `syncmirror_local`, at this time. See your NetApp Sales Representative for a list of additional licenses.

To enable clustering on your cluster, complete the following steps.

Step	Action
1	Enable clustering and other licenses on both filers by using the following command on the local and partner nodes: license add xxxxx xxxxx is the license code you received for the cluster feature. Note _____ If you are using a version of Data ONTAP earlier than 6.4, the command is license option=code. _____
2	Reboot both filers by entering the following console command: reboot
3	On each console, enter the following command: cf enable Result: Clustering is enabled, but not configured.
4	Check the status of your cluster from the console of each node by entering the following command: cf status Result: You should see the following output for the remote node on the local node console: Cluster enabled, filer 2 is up.
5	Configure your cluster as described in “Configuring the cluster” on page 42.

Configuring the cluster

Why you must configure your converted filers

You must configure some parameters, configuration files, and network interfaces, and all licenses in such a way that when one filer takes over the partner node, it appears as if it were the partner node; otherwise, the consequences described in the following table can result.

If...	Then...
Specific options or configuration files have different values	The values used in takeover are unpredictable.
A feature such as CIFS is licensed on one filer but not on the partner node	After a takeover, CIFS clients cannot access their files on the filer that has taken over from the failed filer.
Network interface configurations do not match	Users of the failed filer see broken network connections.

In general, the more similar the configuration of a filer is to that of its partner, the more it behaves like its partner in a takeover. This section describes the tasks you need to perform to ensure a smooth partner takeover.

For detailed information

See the following sections for detailed clustering configuration information:

- ◆ [“Determining current configuration”](#) on page 43
- ◆ [“Types of interfaces used by takeover”](#) on page 46
- ◆ [“Configuring the interface for takeover”](#) on page 48

Determining current configuration

Methods of checking cluster configuration

You have two ways to check your cluster configuration before placing your cluster online:

- ◆ Running the cluster configuration checker script
- ◆ Manually, through the command line

Using the cluster configuration checker

The cluster configuration checker is a script that runs under UNIX or as a CGI script from a UNIX Web server. It detects the following configuration problems:

- ◆ Mismatched licenses
- ◆ Mismatched option settings
- ◆ Incorrectly configured network interfaces
- ◆ Different versions of Data ONTAP on the local and partner nodes

The script is available through download from the NOW Web site at <http://now.netapp.com/NOW/tools>. Look for the cluster configuration checker (cf-config-check.cgi) file.

To run this script from the command line, complete the following steps.

Step	Action
1	Collect the IP addresses and login information for the local and partner nodes.
2	Download the script from the NOW Web site, then unzip the script.
3	Run the script, as described in the online instructions on the script Web page. You can run the script from either the command line or through a Web browser that can access both filers.
4	Examine the script output and correct any errors or differences identified in the output.
5	Go to “ Types of interfaces used by takeover ” on page 46.

Using the command line

To set options on the local and partner nodes from the command line, complete the following steps.

Step	Action
1	View the options set on the local node by entering the following command at the filer's console: options Result: The current option settings on the filer are listed. You should see output similar to the following: <pre> autosupport.doit DONT autosupport.enable off . . dns.domainname filer1 (might be overwritten in takeover) dns.enable off (might be overwritten in takeover) ftpd.enable off . . </pre>
2	Note the values of these options for the local node, and repeat Step 1 on the partner node.
3	Verify that the options that have comments in parentheses are set to the same value for both filers. The comments you see are as follows: ◆ Value might be overwritten in takeover ◆ Same value required in local+partner ◆ Same value in local+partner recommended
4	Correct any mismatched options using the following command: options <i>option_name</i> <i>option_value</i>
5	Check the license settings on both filers using the following command: license

Step	Action
6	Correct mismatched license settings using the following command: license add xxxxx Note_____ If you are using a version of Data ONTAP earlier than 6.4, the command is license <i>option=code</i>. _____
7	Go to “ Types of interfaces used by takeover ” on page 46.

Types of interfaces used by takeover

Interface requirements for takeover

In a cluster, you must configure each filer so that if one filer fails, the partner node can assume its identity. To ensure a successful takeover, your cluster must meet the following requirements:

- ◆ Clustered filers must be attached to the same networks and subnetworks.
- ◆ Each local interface on a filer must have one local node address and one partner address, and for only one type of interface.

Interface types

There are two kinds of interface configurations you can have for your cluster. The *System Administrator's Guide* provides more information about the types of NICs that support dual-MAC addresses. The *System Administrator's Guide* also describes how to configure vif.

Interface type	Description
Single, stand-by NICs	Each filer has a pair of dedicated interfaces with a local IP address. One of the NICs is a standby interface with the partner's IP address. This configuration is used for adapters that support only a single MAC (media access control) address.
Dual MAC capability	Each filer has a single NIC that supports a pair of dual MAC interfaces, which duplicate each other. These NICs can be one of the following: ◆ Single-port 10Base-T/100Base-TX Ethernet ◆ Quad-port 10Base-T/100Base-TX Ethernet ◆ FDDI ◆ GbE

Network interface modes

Each network interface in a cluster can act in one of two modes, as follows:

- ◆ Normal mode, which is the state where each filer manages its own NICs.
- ◆ Takeover mode, which is the state where the partner node takes over control of the failed filer's NICs.

Normal mode and takeover mode are described in the following table.

Interface type	Normal mode	Takeover mode
Dedicated	Supports the identity of the local node	Idle
Standby	Idle	Supports the identity of the partner node by assuming the partner's IP address
Shared	Supports the identity of the local node	Supports the identity of ◆ The local node by retaining the local IP address ◆ The failed filer by assuming the partner's IP address

Configuring the cluster

Configuring the interface for takeover

Interface configuration commands

You can configure interfaces using either the `setup` command or the `ifconfig` command. The `setup` command runs a script that prompts you for configuration information. The `ifconfig` command enables you to manually define the interface parameters for your cluster in the `/etc/rc` file.

Information you need to gather

You must gather the following information before configuring the interfaces:

- ◆ The IP address for both the local node and partner node
- ◆ The netmask for both the local node and partner node
- ◆ The Ethernet addresses for both the local node and partner node

Using the setup command

When you use the `setup` command, a series of questions appears that enables you to configure the filer. As part of this configuration, the following questions appear for each interface. Answer these questions on both filers, where *n.n.n.n* is the local node address and the partner node address for *m.m.m.m*.

Note—
The *System Administrator's Guide* provides more information about using the `setup` command.

Step	Action
1	Run the <code>setup</code> command on the local node by entering the following command: setup

Step	Action
2	Answer the questions displayed on your screen, including the following questions: Please enter IP address for Network Interface ns0 [n.n.n.n]: Please enter netmask [255.255.255.0]: Should interface ns0 take over a partner IP address during failover? [y]: Please enter the IP address to take over: [m.m.m.m]:
3	Reboot your filer by entering the following command: reboot
4	Repeat Steps 1 through 4 for the partner node.

How you use the ifconfig command

Use the `ifconfig` command at the command line to change the network interface and network interface type. If done from the command line, these changes are only effective for the current boot session. To make your changes persistent, you must include the commands in the `/etc/rc` file on your local and partner nodes. See the *System Administrator's Guide* for more information about interfaces.

Interface configuration commands

The following table shows the commands that you enter in the `/etc/rc` file for configuring the interfaces on the converted cluster.

Type of interface configuration to facilitate takeover	Command
Requires a pair of NICs:	<code>ifconfig e0 local_IP_address</code>
◆ Dedicated	
◆ Shared	<code>ifconfig e1 local_IP_address partner_IP_address</code>
◆ Standby	<code>ifconfig e2 partner partner_IP_address</code>

Type of interface configuration to facilitate takeover	Command
Requires a single NIC with dual MAC addresses shared	<code>ifconfig e0 local_IP_address partner \ partner_IP_address</code>
Dedicated interface with a hot-standby	<code>ifconfig e0 local_IP_address</code> <code>ifconfig e1 partner partner_IP_address</code>

Configuring a pair of NICs with dedicated, shared, and standby interfaces

To set the interfaces for a cluster that has dedicated, shared, and standby interfaces on each filer, complete the following steps.

Note

The *System Administrator's Guide* provides more information about using the `ifconfig` command.

Step	Action
1	Check the current settings of the local node by entering the following command on the local node console: <pre>ifconfig -a</pre> Result: You should see output similar to the following: <pre>e0: Flags=40043<UP,BROADCAST,RUNNING,UP_1ARY> mtu 1500 inet 11.22.33.44 netmask ffffffff00 broadcast 11.22.33.45 secondary host inet 11.22.33.70 (not in use) ether 00:01:00:00:00:80 (Data Link Provider Interface)</pre>
2	Set the dedicated network interface IP addresses for the local node in the <code>/etc/rc</code> file. Enter the IP addresses using the <code>ifconfig</code> command, as follows: <pre>ifconfig e0 local_IP_address</pre>

Step	Action
3	Set the shared network interface IP addresses for the partner node in the <i>/etc/rc</i> file. Enter the IP addresses using the <i>ifconfig</i> command, as follows: <i>ifconfig e1 local_IP_address partner partner_IP_address</i>
4	Set the standby network interface IP addresses for the partner node in the <i>/etc/rc</i> file. Enter the IP addresses using the <i>ifconfig</i> command, as follows: <i>ifconfig e2 partner partner_IP_address</i>
5	Save the <i>/etc/rc</i> file.
6	Repeat Steps 1 through 5 for the partner node.

Configuring a single NIC with dual MAC interfaces

To set the interfaces for a cluster that has a single NIC with dual MAC addresses on each filer, complete the following steps.

Step	Action
1	Check the current settings of the local node by entering the following command on the console: <i>ifconfig -a</i> Result: You should see output similar to the following: e0: Flags=40043<UP,BROADCAST,RUNNING,UP_1ARY> mtu 1500 inet 11.22.33.44 netmask ffffffff broadcast 11.22.33.45 secondary host inet 11.22.33.70 (not in use) ether 00:01:00:00:00:00:80 (Data Link Provider Interface)
2	Set MAC network interface IP addresses for the local and partner nodes in the <i>/etc/rc</i> file. Enter the IP addresses using the <i>ifconfig</i> command, as follows: <i>ifconfig e0 local_IP_address partner partner_IP_address</i>
3	Save the <i>/etc/rc</i> file.
4	Repeat Steps 1 through 3 for the partner node.

**Configuring
dedicated and hot-
standby interfaces**

To set the interfaces for a cluster that has dedicated and hot-standby interfaces on each filer, complete the following steps.

Step	Action
1	Check the current settings of the local node by entering the following command on the consoles: ifconfig -a Result: You should see output similar to the following: <pre>e0: Flags=40043<UP,BROADCAST,RUNNING,UP_1ARY> mtu 1500 inet 11.22.33.44 netmask ffffffff broadcast 11.22.33.45 secondary host inet 11.22.33.70 (not in use) ether 00:01:00:00:00:00:80 (Data Link Provider Interface)</pre>
2	Set the hot-standby network interface IP address for the partner node in the <code>/etc/rc</code> file. Enter the IP addresses using the <code>ifconfig</code> command, as follows: ifconfig e1 partner partner_IP_address
3	Save the <code>/etc/rc</code> file.
4	Repeat Steps 1 through 3 for the partner node.

Where to go next

After successfully configuring the network interfaces and testing takeover and giveback on your cluster, see the *System Administrator's Guide* for additional cluster configuration and administration information.

See Appendix A, “[Alternative Configuration and Cabling Examples](#),” on page 135 for alternative configuration diagrams and for additional Channel B cabling diagrams.

About this chapter

This chapter describes how to add an adapter to a node in the cluster and how to add a disk shelf to an existing cluster. It also describes how to add an LRC to FC7, FC8, or FC9 disk shelves. It also describes how to modify your cluster from one containing DS14/DS14mk2 FC disk shelves with LRC modules to DS14/DS14mk2 FC disk shelves with ESH modules.

Topics in this chapter

This chapter describes the requirements and steps you must take to add a disk shelf to your cluster:

- ◆ [“Adding or replacing an adapter”](#) on page 54
- ◆ [“Requirements for adding a disk shelf”](#) on page 55
- ◆ [“Adding a disk shelf to a cluster that has DS14 disk shelves”](#) on page 57
- ◆ [“Hot upgrading a DS14 disk shelf and hot swapping LRC/ESHs”](#) on page 61

Adding or replacing an adapter

About adding or replacing an adapter

You can add or replace an adapter in your cluster without taking your system out of service. Before adding an adapter to your cluster, you should review the *System Configuration Guide* on the NOW Web page at <http://now.netapp.com> for currently supported slot assignments for your new adapter. You also need the *Field Service Guide* for adapter installation instructions.

Adding or replacing an adapter

To add or replace an adapter in your cluster, complete the following steps.

Step	Action
1	Have the partner node take over the node receiving the new adapter by entering the following command and verifying that the target node is halted: cf takeover Result: The filer takes over the partner node. You see a message similar to the following: takeover completed
2	Halt the target node, then turn off the power supplies and remove the power cables from the power source.
3	Add or replace the adapter, as described in the <i>Field Service Guide</i> .
4	Reconnect any cabling to the new adapter, then reconnect the power supplies and turn them on.
5	After the partner boots, enable your cluster again by entering the following command on the local node's console: cf giveback Result: The partner node reboots and functions normally. A successful giveback ends with the following message on the local node: giveback completed

Requirements for adding a disk shelf

What to do if you need more storage

One FC-AL interface can connect up to a maximum of eight daisy-chained disk shelves (four for DS14 disk shelves) per FC-AL adapter. You must add an additional FC-AL adapter to each of your filers for every set of eight disk shelves you add, up to the maximum of FC-AL adapters supported by your filer. See the *F800 Series Field Service Guide* for FC-AL installation information.

See the NOW Web page at <http://now.netapp.com> for the maximum number of FC-AL adapters your filer supports.

- Required tools and manuals**
- You must have the following cables and components for each disk shelf added:
- ◆ Two shelf-to-shelf cables; one each for the Channel A and Channel B loops
 - ◆ Two power cords
 - ◆ Two power outlets, on separate circuits if possible
 - ◆ The *F800 Series Field Service Guide* for FC-AL installation instructions
 - ◆ The appropriate Fibre Channel disk shelf guide

Restrictions for adding new disk shelves

The following table lists some of the restrictions for adding disk shelves to your cluster. See the *DiskShelf14* and *DiskShelf14mk2 FC Hardware Guide* for a more detailed discussion of restrictions.

If your cluster has...	And you are adding...	Then...
FC7 or FC8 disk shelves	FC9 disk shelves in the same loop	Disk shelf 0 must be an FC7 disk shelf.
FC7, FC8, or FC9 disk shelves	DS14/DS14mk2 FC disk shelves with LRCs	The DS14/DS14mk2 FC must be the last disk shelf in the loop.

If your cluster has...	And you are adding...	Then...
DS14/DS14mk2 disk shelves with ESH modules	Hot adding a disk shelf	◆ Disks should be pre-zeroed and used as spares. ◆ You can only hot add a DS14/DS14mk2 FC disk shelf to an existing DS14/DS14mk2 FC disk shelf. ◆ The DS14/DS14mk2 FC must be the last disk shelf. ◆ DS14/DS14mk2 FC disk shelves with ESH modules can only be added to disk shelf loops with ESH modules. ◆ If the new disk shelf is a DS14mk2 FC, the speed switch must be set for the appropriate installation.

Adding a disk shelf to a cluster that has DS14 disk shelves

About adding a disk shelf

You can add a DS14 or DS14mk2 FC disk shelf to a cluster that is made up of DS14 disk shelves without powering down your cluster. This enables you to add disk shelves without a service interruption, up to the total amount of storage your cluster supports. See the NOW Web page at <http://now.netapp.com> for the amount of storage your cluster can support.

Caution

You cannot mix disk shelf types with DS14/DS14mk2 FC disk shelves containing ESH modules. All disk shelves in that loop must be DS14/DS14mk2 FC disk shelves with ESH modules.

For detailed information

For details about adding a disk shelf to a cluster that has DS14 disk shelves, see the following sections:

- ◆ “[Step 1: Preparing to add the disk shelf](#)” on page 58
- ◆ “[Step 2: Cabling the new disk shelf](#)” on page 60

Step 1: Preparing to add the disk shelf

Setting up the disk shelf

To prepare your new disk shelf for addition to your existing cluster, complete the following steps.

Step	Action
1	Disable negotiated failover in your cluster by entering the following command from the node console: cf nfo disable disk_shelf
2	Install the new disk shelf in your equipment rack, as described in the <i>DiskShelf14 and DiskShelf14mk2 FC Hardware Guide</i> .
3	Ground the new disk shelf, then give the disk shelf a unique shelf ID.
4	Verify that the new disk shelf ID is unique by entering the following command: fcstat device_map Result: Shelf ID's for all Fibre Channel devices are listed.

Step	Action	
5	If disk shelves in the loop to which you are adding the new disk shelf have...	Then, for loops with...
	LRCs	◆ DS14 only: Go to Step 6. ◆ A mix of DS14 and DS14mk2 FC: Set the loop speed to 1Gb. Caution_____ An incorrectly set loop speed results in an open loop condition. _____
	ESHs	◆ DS14 only or a mix of DS14 and DS14mk2 FC: Set the loop speed to 1 Gb. ◆ DS14mk2 FC only that are running at 1 Gb: Set the loop speed to 1 Gb. ◆ DS14mk2 FC only that are running at 2 Gb: Set the loop speed to 2 Gb. Caution_____ An incorrectly set loop speed causes the storage appliance to panic. _____
6	Plug in the new disk shelf and turn it on.	
7	Go to “ Step 2: Cabling the new disk shelf ” on page 60.	

Step 2: Cabling the new disk shelf

Cabling the new disk shelf

To cable the new disk shelf to the existing Channel A and B loops, complete the following steps.

Step	Action
1	Cable the last disk shelf Channel A LRC/ESH Output port to the new disk shelf Channel A LRC/ESH input port. Label the cable with a solid-colored label.
2	Cable the last disk shelf Channel B LRC/ESH Output port to the new disk shelf Channel B LRC/ESH Input port. Label this cable with a striped label.
3	Set both terminate switches on the new disk shelf to On.
4	Set both terminate switches on the previous disk shelf to Off.
5	Reenable negotiated failover in your cluster by entering the following command from the node console: cf nfo enable disk_shelf
6	Go to “Testing the LRC/ESH modules” on page 63.

Hot upgrading a DS14 disk shelf and hot swapping LRC/ESHs

Tasks for hot upgrading and hot swapping

This section contains information about hot upgrading a DS14/DS14mk2 FC disk shelf to a cluster and how to hot swap an LRC/ESH module in a cluster. It describes the procedures in the following sections:

- ◆ “Hot upgrading your DS14 disk shelves” on page 61
- ◆ “Hot swapping a single LRC/ESH module” on page 62
- ◆ “Testing the LRC/ESH modules” on page 63
- ◆ “Troubleshooting takeover or giveback failures” on page 64

Requirements for upgrading and hot swapping

Upgrading a cluster that contains DS14/DS14mk2 FC disk shelves with LRC modules to a cluster containing DS14/DS14mk2 FC disk shelves with ESH modules requires that you exchange the LRC modules with ESH modules.

Hot upgrading your DS14 disk shelves

To convert your DS14/DS14mk2 FC disk shelves with LRCs into disk shelves with ESH modules, complete the following steps.

Step	Action	
1	If the cluster is a...	Then...
	Standard cluster, mirrored cluster, non-switched MetroCluster, or switched MetroCluster	1. Take over the partner filer’s disk shelves by entering the following command on the local filer: cf takeover 2. Go to Step 2.
	Fabric Attached MetroCluster	Go to Step 2.
2	Note the terminator switch position of the LRC you are replacing. The new ESH module must be set to the same setting as the old LRC.	
3	While grounded, unplug the cabling from the partner filer’s disk shelves. Note the cable locations, especially if your cluster is a Fiber Attached MetroCluster.	

Step	Action	
4	Remove the Channel A loop LRCs of the partner's disk shelves and replace them with ESH modules.	
5	Recable the disk shelves to their original locations before the ESH installation.	
6	Check the operation of the new ESH modules by entering the following command from the local node console: environ shelf Result: The local filer reports the status of the modified shelves as good.	
7	If clustering was...	Then...
	Disabled	1. Return control of the partner filer's disk shelves by entering the following command: cf giveback 2. Repeat Steps 2 through 6 for the Channel B loop.
	Not disabled	Repeat Steps 2 through 6 for the Channel B loop.
8	Repeat Steps 1 through 7 for the local filer.	
9	Test the configuration as described in Chapter 1, " Step 5: Completing setup and testing your cluster ," on page 15.	

Hot swapping a single LRC/ESH module

To hot swap a single LRC/ESH module, complete the following steps.

Step	Action
1	Take over the partner filer's disk shelves by entering the following command on the local filer: cf takeover

Step	Action
2	Note the terminator switch position of the LRC or terminator and loop speed switches of the ESH you are replacing. The new LRC/ESH module must be set to the same settings as the old LRC/ESH.
3	While grounded, unplug the cabling from the target LRC/ESH module, then remove the module.
4	Slide the replacement LRC/ESH module into the slot at the rear of the disk shelf and push the levers of the CAM mechanism into place.
5	Reconnect the cabling to the replacement LRC/ESH.
6	Return control of the partner filer's disk shelves by entering the following command: cf giveback

Testing the LRC/ESH modules

To test your LRC/ESH modules, complete the following steps on both nodes.

Step	Action
1	Verify that all disk shelves with ESH modules are functioning properly by entering the following command: environ shelf Results: The system reports the modules as good.
2	Verify that you can create and retrieve files on both nodes for each licensed protocol.
3	Verify there are no missing disks by entering the following command: vol status -r Result: Local disks displayed on the local node should be displayed as partner disks on the partner node and vice-versa.

Troubleshooting takeover or giveback failures

If takeover or giveback fails, complete the following steps.

Step	Action
1	Check the Channel B cabling on the disk shelves, and tighten any loose cables.
2	Check the cabling to the cluster interconnect adapter. Make sure that the cables are connected to the correct interconnect adapter ports and that they are secure.
3	Check communication between the nodes by entering the following command: cf status Result: If there is a cabling error or the interconnect link is down, the system displays an error message indicating the error.
4	Check the parameters and option settings on your cluster with the cluster configuration checker script. This script is available for download from the NOW Web site at http://now.netapp.com/NOW/tools . Look for the cluster configuration checker (cf-config-check.cgi) file.
5	Check network connectivity. See the appropriate System Administration guide for more information.
6	Reboot the cluster and rerun the takeover and giveback tests.

About this chapter

This chapter describes the tasks you need to perform to disable a standard cluster and return the filers to independent systems. If you have MetroClusters, contact your Customer Support representative for information about returning a MetroCluster to stand-alone systems.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Removing clustering software”](#) on page 66
- ◆ [“Removing clustering hardware”](#) on page 67

Removing clustering software

Disabling clustering functionality

To disable clustering functionality, complete the following steps.

Step	Action
1	On the console of either filer, enter the following command: cf disable Result: This command disables the clustering functionality on both filers.
2	Disable the clustering license by entering the following command: license delete cluster Note _____ If you are using a version of Data ONTAP earlier than 6.4, the command is license cluster=disable . _____
3	Review the /etc/rc file in the root of the local node. Remove any commands that reference a partner node, then save the file. For example, when viewing the /etc/rc file, you might see the following entry: ifconfig e1 partner 10.10.11.12
4	Repeat Steps 2 and 3 on the partner node.
5	Go to “ Removing clustering hardware ” on page 67.

Removing clustering hardware

About removing clustering hardware

Removing clustering hardware is a three-part process:

- ◆ Preparing to remove the clustering hardware
- ◆ Recabling your disk shelves
- ◆ Removing the clustering hardware

Recabling Channel B for redundant disk loops

To recable your Channel B loops and convert your cluster to stand-alone filers, complete the following steps.

Note_____

This procedure assumes that you have DS14 disk shelves and dual Fibre Channel adapters.

Step	Action
1	Halt both filers by entering the following command at each console: halt
1	Ground yourself, then power off the nodes and disk shelves.
2	Unplug the local node from the partner node’s Channel B disk shelf and do the same for the partner node.
3	Connect the local node to the open Channel B loop in its local disk shelves.
4	Repeat Step 3 for the partner node.
5	Go to “ Removing the clustering hardware ” on page 68.

Removing the clustering hardware

To remove the clustering hardware and extra FC-AL adapters, complete the following steps.

Step	Action	
1	Ground yourself, using the antistatic wrist strap.	
2	Remove the interconnect cables from the cluster adapter.	
3	Open the back of the local node, as described in the <i>F800 Series Hardware Guide</i> .	
4	If you have...	Then...
	Cluster interconnect adapters or IB cluster adapters	1. Unscrew the adapter from the chassis. 2. Remove the adapter. 3. Replace the slot cover and secure it with the screw from the adapter.
	VI cluster adapters	1. Unlock and remove the GBIC from the adapter. 2. Unscrew the adapter from the chassis. 3. Remove the adapter. 4. Replace the slot cover and secure it with the screw from the adapter.
5	Remove any extra FC-AL or dual Fibre Channel adapters.	
6	Close and secure the back of the appliance, as described in the <i>F800 Series Hardware Guide</i> .	
7	Power on the disk shelves, then the filer and press the <i>Del</i> key during boot to get to the OK prompt.	
8	Unset the partner sysid by entering the following command at the OK prompt: unsetenv partner-sysid	

Step	Action
9	Boot the filer by entering the following command at the OK prompt: boot
10	Check the status of clustering on the filer by entering the following command on the console: cf status Result: Cluster monitoring is disabled, as indicated by the following message: Cluster disabled
11	Repeat Steps 1 through 10 for the partner node.

About this chapter

This chapter describes the tasks you need to perform to set up a new mirrored cluster and what you need to do to convert an existing cluster to a mirrored cluster. It contains the following information:

- ◆ What the minimum and maximum hardware requirements are for creating a new mirrored cluster
- ◆ What the minimum and maximum hardware requirements are for converting an existing cluster into a mirrored cluster
- ◆ How to plan your disk pools for mirrored clusters
- ◆ How to cable your filers into a mirrored cluster
- ◆ How to set up your cluster for disaster recovery

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“System requirements for mirrored clusters”](#) on page 72
- ◆ [“Planning your mirrored cluster”](#) on page 76
- ◆ [“Cabling your system into a mirrored cluster”](#) on page 80
- ◆ [“Recovering from a disaster”](#) on page 90

System requirements for mirrored clusters

What mirrored clusters contain

A mirrored cluster contains two complete copies of the data volumes or file systems that you indicated as being mirrored volumes or file systems in your cluster. These copies are called plexes and are continuously and synchronously updated every time Data ONTAP writes data to the disks. Plexes are physically separated from each other across different groupings of disks.

Advantages of mirrored clusters

The advantages of having duplicate copies of your data are as follows:

- ◆ Data on your cluster survives a double disk error in one of your disk shelves
- ◆ Data on your cluster survives if one part of the cluster loses connectivity to your disks
- ◆ Data on your cluster survives if your cluster loses connectivity to remote storage
- ◆ Data on your cluster survives the destruction of a mirrored site or destruction of connectivity to the remote site

In any of these situations, the surviving mirrored plex continues to serve data until the problem is corrected. See the *Data Protection Guide* for detailed information about disaster protection through mirroring.

How Data ONTAP works with mirrored clusters

During configuration of the cluster, Data ONTAP identifies spare disks and divides them into separate groupings called pools. These pools of disks are physically separated from each other, allowing for high availability of mirrored volumes. When you add a mirrored volume or add disks to one side of a mirrored volume, Data ONTAP determines how much storage you need for the second half of the mirror, and dedicates that storage from a separate pool to the mirrored volume.

For more detailed information about mirroring disk volumes, see the *Data Protection Guide* and *System Administrator's Guide*.

Mirrored cluster setup requirements and restrictions

The following list specifies the requirements and restrictions you need to be aware of when setting up a new mirrored cluster or converting an existing cluster to a mirrored cluster:

- ◆ Disks in the same volume must be part of the same plex, with its mirror in a different plex.
- ◆ Disk volumes that you want to mirror must have at least an equivalent number of disks in the mirrored plex of the same or greater disk size.
- ◆ Disk pool 0 volumes must be cabled to FC-AL or dual Fibre Channel adapters in filer slots 3 through 7.
- ◆ Disk pool 1 volumes must be cabled to FC-AL or dual Fibre Channel adapters in filer slots 8 and 9.
- ◆ The two halves of the mirrored disk should have no common components.
- ◆ `syncmirror_local` must be licensed and enabled on the local and partner nodes.
- ◆ Ensure that you have sufficient spares in each pool to account for a disk failure.

For detailed information

See the following sections for detailed system requirements information:

- ◆ [“Required hardware”](#) on page 74
- ◆ [“Required software and firmware”](#) on page 75

Required hardware

Required hardware In addition to the standard clustering hardware listed in “[System requirements](#)” on page 22, you also need to add the following number of Fibre Channel adapters for a mirrored cluster.

Cluster type	Fibre Channel adapter type	Total minimum configuration per filer	Total maximum configuration per filer
F810 cluster	FC-AL	0	0
	Dual Fibre Channel	2	2
F820 cluster	FC-AL	4	4
	Dual Fibre Channel	2	2
F825 cluster	FC-AL	2	2
	Dual Fibre Channel	1	2
F840 cluster	FC-AL	4	4
	Dual Fibre Channel	2	4
F880 cluster	FC-AL	4	4
	Dual Fibre Channel	2	4

Required software and firmware

Software requirements

To convert an existing cluster into a mirrored cluster, you must have Data ONTAP 6.2 or later software installed.

Note

The software is on a CD-ROM for new systems, or you can download it from <http://now.netapp.com>. To download the software, you need network connectivity and access to a UNIX or PC system.

Firmware requirements

Your filers and disk drives require the following firmware levels for cluster conversion. If your disk type is not listed, contact your Customer Support representative or see the NOW Web page at <http://now.netapp.com> for the current firmware version.

Type of firmware	Version
Filer	See the NOW site at http://now.netapp.com/NOW/tools/releasetable.shtml for current filer firmware.
Disk drive	See the NOW site at http://now.netapp.com/NOW/download/tools/diskfw/ for current disk firmware and a list of supported disk drives and associated disk firmware.

Planning your mirrored cluster

Planning for a new mirrored cluster

If you are creating a new mirrored cluster, you should determine the following information:

- ◆ The number of volumes your file systems will have
- ◆ The number of volumes that will be mirrored and what the composite size is of the volumes to be mirrored

Planning for mirroring an existing cluster

Your current system configuration determines what you need to modify before cabling your individual filers into a mirrored clustered system. You should determine the following information:

- ◆ The number of volumes you plan to have or already have in your system
- ◆ The number of existing volumes you plan on mirroring
- ◆ The disk sizes you have in your existing system
- ◆ Whether your total storage can be divided evenly (equal amounts of storage in each half)
- ◆ Where your volumes are in your existing cluster

For detailed information

For details about planning your mirrored cluster, see the following sections:

- ◆ [“Step 1: Mapping the cluster”](#) on page 77
- ◆ [“Step 2: Determining your storage requirements”](#) on page 78
- ◆ [“Step 3: Determining your FC-AL requirements”](#) on page 79

Step 1: Mapping the cluster

Mirror mapping considerations

If you have a clustered system and want to set up a mirrored cluster using synchronous mirroring, you must do the following:

- ◆ Determine what volumes you currently have.
- ◆ Determine where each volume is located.
- ◆ Determine the size of each volume you want to mirror.
- ◆ Determine how you will structure your data on the new system by sketching a structure of how the data will look, after you begin storing information on your new cluster.

Mapping your cluster

To map your existing cluster, complete the following steps.

Step	Action
1	List all volumes you have on the local node by entering the following command: vol status -r Result: All volumes and spare disks are listed.
2	Determine the size of each volume by entering the following command: vol status -b Result: The command displays the block size of the volume (in bytes), the RAID volume size, and the WAFL® file system size.
3	Note the size and name of the volumes you want to mirror.
4	Repeat Steps 1 through 3 for the partner node.
5	Go to “Step 2: Determining your storage requirements” on page 78.

Step 2: Determining your storage requirements

Storage considerations

After you determine which volumes you want to mirror and their sizes, you must determine the following:

- ◆ The total storage requirements for the volumes you want to mirror
- ◆ The types of disks you have in your cluster
- ◆ Whether the plexes are the same size and reside on the same size disks
- ◆ Whether you need to add storage or move spare disks from one pool to another

Determining your storage requirements

To determine your storage requirements, complete the following steps.

Step	Action
1	Ensure that the Pool 0 disks have the same size or greater sized disks in Pool 1.
2	Determine the spare disks in the disk pool on both the local and partner nodes by entering the following command: sysconfig -r
3	Compare the data from Step 1 with that of Step 2 and add disks to the local or partner node disk pools, as required.
4	Go to “ Step 3: Determining your FC-AL requirements ” on page 79.

Step 3: Determining your FC-AL requirements

FC-AL considerations

When you set up your cluster for mirroring, you should consider the following:

- ◆ The type and number of FC-AL adapters you have in your system
- ◆ Where the FC-AL adapters are located in your cluster.

Disk pool 0 volumes are cabled to FC-AL or dual Fibre Channel adapters located in filer slots 3 through 7, if available. Disk Pool 1 volumes are cabled to FC-AL or dual Fibre Channel adapters in filer slots 8 and 9. No single FC-AL should become a single point of failure.
- ◆ If you are converting a cluster into a mirrored cluster, whether you need to add more FC-AL adapters to spread them across disk pools

Determining your FC-AL requirements

To determine your FC-AL requirements, complete the following steps.

Step	Action
1	Determine how many disk pools you have for each filer. You should have no more than one disk pool per filer for the local node. You should also have one disk pool for the partner node on the local node side.
2	Determine the number of FC-AL adapters you have and install any additional adapters you need. See the <i>Field Service Guide</i> for installation instructions.

Cabling your system into a mirrored cluster

Required tools and manuals

You need the following tools and guides to prepare your system cabinets or equipment racks for mirrored clustering:

- ◆ The *42U System Cabinet Guide* or the *F800 Series Hardware Guide*
- ◆ The *Data Protection Guide* and the *System Administrator's Guide*
- ◆ The appropriate Fibre Channel disk shelf guide
- ◆ #1 and #2 Phillips screwdrivers

Required cables

When you received your system cabinets or nodes and disk shelves for installation in equipment racks, you should have received the disk shelf cables listed in the following table:

Cables	DS14	DS14mk2 FC
Filer-to-shelf	HSSDC to DB-9 SC/LC or LC/LC	SC/LC or LC/LC
Shelf-to-shelf	DB-9 to DB-9 HSSDC2-to-HSSDC2	HSSDC2-to-HSSDC2
Shelf-to-switch	LC/SC	LC/SC

You should also receive cluster interconnect cables, one pair for the cluster interconnect adapter or one optical cable for the VI cluster adapter. If your cluster will be farther apart than 30 meters, you must use an SC-SC Duplex 50/125 plenum cable to connect them.

For detailed information

For details about the cabling process, see the following topics:


- ◆ [“Step 1: Cabling the Channel A loop”](#) on page 81
- ◆ [“Step 2: Cabling the nodes together”](#) on page 83
- ◆ [“Step 3: Cabling the Channel B loop”](#) on page 84
- ◆ [“Step 4: Completing setup and testing your mirrored cluster”](#) on page 86

Cabling your system into a mirrored cluster

Step 1: Cabling the Channel A loop

Cabling the Channel A loops

To cable the Channel A loops, complete the following steps, using the diagram for reference.



Step	Action
1	Cable the FC-AL in Slot 8, Port 0 to the disk shelf 1 LRC/ESH Input port of disk pool 1, then cable the FC-AL in Slot 6, Port 0 to the disk shelf 1 LRC/ESH Input port of disk pool 0.

Step	Action	
2	Cable the disk shelf 1 Channel A LRC/ESH Output port to the disk shelf 2 Channel A LRC/ESH Input port for both disk pools. Label these cables with solid-colored labels.	
3	Repeat Step 2 connecting output to input, for the remaining disk shelves in the loop.	
4	Set the terminate switches to Off on all but the last disk shelf in the loop, then set the terminate switch on the last disk shelf to On.	
5	If you have...	Then...
	DS14 disk shelves with ESH modules	Set the 1/2GB switch to 1Gb.
	DS14mk2 FC disk shelves	Set the 1/2GB switch to 2Gb if you are using the high capacity setting for the disk shelves. Otherwise, set the switch to 1Gb.
6	Repeat Steps 1 through 5 for the partner node.	
7	Go to “ Step 2: Cabling the nodes together ” on page 83.	

Cabling your system into a mirrored cluster

Step 2: Cabling the nodes together

Cabling the nodes To cable the nodes together through the VI cluster adapter or IB cluster adapter, complete the following steps.


Caution—————
You cannot use the cluster interconnect adapter in a mirrored cluster. You must use either the VI cluster adapter or the IB cluster adapter.

Step	Action
1	Cable the cluster adapters together, as described in “ Step 3: Cabling the nodes together ” on page 11.
2	Go to “ Step 3: Cabling the Channel B loop ” on page 84.

Step 3: Cabling the Channel B loop

Cabling the Channel B loops

To cable the Channel B loops, complete the following steps, using the diagram for reference.



Step	Action
1	Cable the FC-AL adapter in Slot 8, Port 1 of the local node, to the local node disk shelf 1 LRC/ESH Input port of disk pool 1, then cable the FC-AL adapter in Slot 6, Port 1 of the local node to the partner node disk shelf LRC/ESH Input port of the first disk shelf in disk pool 0.
2	Cable the local node's disk shelf 1 Channel B LRC/ESH Output port to the local node's disk shelf 2 Channel B LRC/ESH Input port and repeat for the remaining disk shelves in the loop in both disk pools. Label these cables with a striped labels.
3	Cable the FC-AL adapter in Slot 6, Port 1 of the partner node to the local node disk shelf LRC/ESH Input port of the first disk shelf of disk pool 0, then cable the FC-AL adapter in Slot 8, port 1 of the partner node to the partner node disk shelf 1 LRC/ESH Import of disk pool 1, then cable the remaining disk shelves for the node.
4	Set the terminate switches to Off on all but the last disk shelf in the loop, then set the terminate switch on the last disk shelf to On.
5	Repeat Steps 1 through 4 for the partner node.
6	Go to “Step 4: Completing setup and testing your mirrored cluster” on page 86.

Step 4: Completing setup and testing your mirrored cluster

Completing cluster setup

Completing the setup of your mirrored cluster consists of the following tasks:

- ◆ Cabling the local and partner nodes to the network and to power sources, as described in the *F800 Series Hardware Guide*.
- ◆ Bringing up your cluster.
- ◆ Testing the cluster.
- ◆ If your cluster is in systems cabinets, you must reinstall the front and rear doors of your clustered system cabinets, to comply with FCC noise reduction requirements. See the *42U System Cabinet Guide* for more information about door installation.

Bringing up your cluster

To bring up your mirrored cluster, complete the following steps.

Step	Action
1	Apply power to your disk shelves, then to your filers.
2	Enable the syncmirror_local and clustering licenses by entering the following command: <pre>license add xxxxx</pre> Note_____ If you are using a version of Data ONTAP earlier than 6.4, the command is license option=code. _____

Step	Action
3	At the prompt, enter the following command: <pre>vol create vol-name -m ndisks[@disk-size]</pre> <i>vol-name</i> is the name for the new volume (without the /vol/ prefix). <i>ndisks</i> is the number of disks to use; ndisks must be at least 4 and in even increments. Even increments are required because half of the disks specified by ndisks go to one plex from one disk pool and half go to the other plex from another disk pool. <i>disk-size</i> is a number that specifies the disk size to use, in gigabytes. You must have half of the ndisks available disks of the size you specify in one disk pool and half of the ndisks available disks of the size you specify in another disk pool if you use this option.
4	Reboot both nodes.
5	Complete the configuration, as described in “Configuring the cluster” on page 42.

Testing your cluster To test your cluster, complete the following steps.

Step	Action
1	Verify the network addresses by entering the following command from one filer: <pre>ifconfig -a</pre> Result: The local node and partner node IP addresses are listed.
2	Verify that you can create and retrieve files on both filers for each licensed protocol.
3	From the same filer, enter the following command: <pre>cf enable</pre> Result: Takeover of the partner node is enabled.

Step	Action
4	Test communication between the local node and partner node. For example, you can use the <code>fcadmin device_map</code> command to ensure that one filer can access the other filer's disks. Example: <pre>filer-1> fcadmin device_map Loop Map for channel 8: Translated Map: Port Count 32 7 16 17 18 19 32 33 34 35 48 49 50 51 24 26 27 57 61 42 43 44 45 28 29 30 52 53 54 37 38 21 22 Shelf mapping: Shelf 2: 22 21 XXX 19 18 17 16 . . .</pre>
5	From the local node console, enter the following command and verify that the partner node halted: cf takeover Result: The local node takes over the partner node. You see a message similar to the following: <pre>takeover completed</pre>
6	Verify that the correct volumes are mirrored by entering the following command: vol status -r
7	Check the parameters and options on your cluster with the Cluster Configuration Checker script. The script is available through download from the NOW Web site at http://now.netapp.com/NOW/tools. Look for the Cluster Configuration Checker (cf-config-check.cgi) file. Correct any errors you find.

Step	Action
8	From the same filer, enter the following command: cf giveback Result: The partner node reboots and functions normally. A successful giveback ends with the following message on the local node: giveback completed
9	Repeat Steps 1 through 8 on the other filer.

Recovering from a disaster

About recovering from a disaster

There are several types of disasters you might encounter:

- ◆ Failure of the cluster interconnect between the two sites
- ◆ Failure of access to the partner’s storage
- ◆ Complete loss of communication between the two sites

In any of these situations, you should immediately restrict access to the disaster site’s node to prevent possible data corruption by the filer at the site recovering automatically by reestablishing contact with its partner. See the *Data Protection Guide* for detailed information about these types of disasters.

Recovering from a disaster

To recover from a disaster, complete the following steps.

Note

These procedures work for mirrored clusters and MetroClusters, as described in Chapter 7, “[Cabling a MetroCluster](#),” on page 93.

Step	Action
1	Isolate the partner node (filer at disaster site) by powering it off. If you cannot physically turn off that filer, isolate it by manually fencing it, as describe in the <i>Data Protection Guide</i> or the <i>Data ONTAP 6.2 Release Notes</i> .
2	Determine accessibility to the partner node and disks, as described in the <i>Data Protection Guide</i> or the <i>Data ONTAP 6.2 Release Notes</i> .

Step	Action	
3	If...	Then...
	The local and partner nodes can access their local data, but cannot access or communicate with each other or their remote storage	1. Restore access by replacing or reconnecting the damaged or missing cables. Note_____ Do not force a takeover of the partner. 2. Confirm access to the remote storage by entering the following command: sysconfig -v 3. Confirm connection between the local and partner node, by entering the following command: cf monitor
	The remote node and storage are physically damaged or failed	1. Force the takeover of the partner node by entering the following command: cf forcetakeover -d 2. Complete the disaster recovery procedures, as described in the <i>Data Protection Guide</i> or the <i>Data ONTAP 6.2 Release Notes</i>.

About this chapter

This chapter contains instructions for cabling a new MetroCluster, both nonswitched and MetroClusters with QuickLoop, instructions for cabling a new Fabric Attached MetroCluster, and for converting an existing cluster into a MetroCluster. This chapter also describes hot swap error messages and gives procedures for replacing components in a MetroCluster.

Types of MetroClusters

The following list describes the MetroClusters types:

- ◆ Fabric Attached MetroCluster—This is a clustered system where the local and partner nodes are separated by up to 30 kilometers. Each node and associated disk shelves are connected directly to a pair of switches and not to each other, like a traditional cluster. Communication with each other is through fiber optic cables and Brocade® Silkworm® 3200 or 3800 switches. This feature is new in Data ONTAP 6.4.1 software and is available only on F880 or F825 clustered systems for all newly installed MetroClusters over 500 meters.
- ◆ Non-switched MetroCluster—This is a clustered system where the local and partner nodes are separated by between 30 and 500 meters, without Fibre Channel switches. This feature is supported on all F800 series clusters.
- ◆ MetroClusters with QuickLoop—This is a clustered system where the local and partner nodes are separated by up to 10 kilometers and communicate with each other through fiber optic cables and Brocade Silkworm 3200 or 3800 switches. This feature is supported in Data ONTAP 6.3.1 software and later, and is available only on F880 or F825 clusters.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“System requirements”](#) on page 94
- ◆ [“Cabling a new Fabric Attached MetroCluster”](#) on page 97
- ◆ [“Cabling a non-switched MetroCluster”](#) on page 113
- ◆ [“Cabling a MetroClusters with QuickLoop”](#) on page 114
- ◆ [“Converting an existing cluster to a Fabric Attached MetroCluster”](#) on page 131
- ◆ [“Replacing MetroCluster components”](#) on page 133
- ◆ [“Interpreting hot swap disk loop error messages”](#) on page 134

System requirements

Required hardware and cables

You must have the following equipment for a MetroCluster. See the System Configuration Guide to confirm that this information is current.

Required equipment	MetroCluster	MetroCluster with QuickLoop	Fabric Attached MetroCluster
Distance between nodes	30 to 500 meters	500 meters to 10 kilometers	500 meters to 30 kilometers
Appliance type for nodes	Two of the same type of any F800 series appliance	Two F880 or two F825 appliances	
Storage	Up to the maximum for mirrored clusters. See the System Configuration Guide for storage maximums.	Four or eight TB of storage of DS14 disk shelves in the cluster. The DS14 disk shelves must have optical LRCs.	Up to 32 TB of storage for DS14 and DS14mk2 FC disk shelves.
Cluster adapter supported	IB cluster adapter VI cluster adapter, Rev. B	VI cluster adapter, Rev. B	VI-MC cluster adapter
FC HBA (FC HBA for Disk) adapters	None	Two per node The ports on the FC HBA adapters are labeled 1 and 2. However, the software refers to them as a and b. You see these labeling conventions in the user interface and system messages displayed on the console.	

Required equipment	MetroCluster	MetroCluster with QuickLoop	Fabric Attached MetroCluster
Switches	None	Two 16-port, Brocade 3800 switches or two 8-port Brocade 3200 switches.	Four 16-port, Brocade 3800 switches, or Four 8-port Brocade 3200 switches, or two Brocade 3800 switches on one side of the cluster and two Brocade 3200 switches on the other side of the cluster.
SFP modules	None	Two per ISL, long distance One per switch port used, short distance	Two per ISL, long distance One per switch port used, short distance
Cables	◆ Four SC/LC FC-AL to disk shelf cables ◆ Two SC/LC VI cluster adapter-to-switch cables ◆ Four SC/LC disk shelf-to-switch cables	◆ Four SC/LC FC-AL to disk shelf cables ◆ Two SC/LC VI cluster adapter-to-switch cables ◆ Four SC/LC disk shelf-to-switch cables ◆ Four LC/LC node-to-switch cables ◆ Two LC/LC ISL (Inter Switch Links) cables, not provided in shipment	◆ LC/LC node-to-switch cables ◆ SC/LC (for DS14) or LC/LC (for DS14mk2 FC) disk shelf-to-switch cables ◆ Two LC/LC ISL cables, not provided in shipment ◆ HSSDC2-to-HSSDC2 for DS14 disk shelves with ESH modules or DB-9 to DB-9 shelf-to-shelf cabled

Required software, firmware, and licenses

You must have the following software, firmware, and licenses for a MetroCluster or Fabric Attached MetroCluster:

- ◆ Fabric Attached MetroClusters—Data ONTAP 6.4.1 or later
- ◆ Non-switched MetroClusters—Data ONTAP 6.2 or later
- ◆ MetroClusters with QuickLoop—Data ONTAP 6.3.1 or later
- ◆ Diagnostics 3.1 or later
- ◆ Firmware:
 - ❖ Filer firmware—See the NOW Web site at <http://now.netapp.com/NOW/tools/releasetable.shtml> for current filer firmware.
 - ❖ Disk firmware—See the NOW Web site at <http://now.netapp.com/NOW/download/tools/diskfw/> for current disk firmware.
 - ❖ Switch firmware—See the NOW Web site at <http://now.netapp.com/NOW> for the current switch firmware.
- ◆ Node licenses that must be installed on any MetroCluster:
 - ❖ syncmirror_local
 - ❖ cluster_remote
 - ❖ cluster
- ◆ Extended Fabric license installed on the switches, for Fabric Attached MetroClusters separated by more than 10 kilometers and up to 30 kilometers.

Required manuals

The following list specifies the manuals you should have available when installing a MetroCluster:

- ◆ *Brocade SilkWorm 3200 Hardware Reference*
- ◆ *Brocade SilkWorm 3800 Hardware Reference*
- ◆ *Brocade SilkWorm 3900 Hardware Reference*
- ◆ *Brocade Fabric OS Reference*
- ◆ *Data ONTAP System Administrator's Guide*
- ◆ *Data Protection Guide*
- ◆ *F800 Series Field Service Guide*

Cabling a new Fabric Attached MetroCluster

About Fabric Attached MetroClusters

A Fabric Attached MetroCluster is a cluster where the disk shelves and nodes are connected through two Fibre Channel switches per node, instead of the disk shelves and nodes connected to each other. The disk shelves and nodes communicate with each other and to the partner nodes and disk shelves through the switches.

A Fabric Attached MetroCluster has the following characteristics:

- ◆ Four Fibre Channel switches, either the Brocade 3200 or 3800, two switches of the same type per node.
- ◆ Cluster interconnect traffic passes through the switches.
- ◆ One VI-MC cluster adapter per node.
- ◆ Disk shelves are divided into pools.
- ◆ Disk shelf ownership by a node is determined by where the disk shelves are plugged into the switch and by where the FC HBA is connected to the switch.
- ◆ Disk shelves are dual pathed to the same ports on different switches on the same side of the cluster, for example, the first disk shelf in Pool 0 is cabled to port 3 on Switch 1 and also to Port 3 on switch 2.


How switches are divided

The Fibre Channel switches are divided into banks and quadrants. Each switch has two banks and four quadrants, dividing the switch into equal parts. For example, a Brocade 3800 switch has 16 ports, divided into two banks of eight ports per bank.

The banks are further subdivided into two quadrants of four ports each (two for the Brocade 3200). These quadrants are also identified as Pools, which function the same in a Fabric Attached MetroCluster as in a mirrored cluster disk pool. Quadrants 1 and 3 belong to Pool 0, while quadrants 2 and 4 belong to Pool 1. Each switch bank contains a Pool 0 and Pool 1.

See the *Data Protection Guide* for a detailed discussion of disk pools.

The following illustration shows a Brocade 3800 switch, divided into banks, quadrants, and pools.



What a switch fabric is

A Fabric Attached MetroCluster contains two fabrics. A switch fabric consists of one switch connected to a switch on the partner half of the cluster. The two switches are connected to each other through the long distance ISL cables.

For example, the first fabric in a MetroCluster begins from Switch 1 and is completed by connecting the ISL to the first switch on the partner side—Switch 3. The second fabric consists of the second switch, connected through a second ISL cable, to the second switch on the partner side—Switch 4.

How disk ownership is determined

Disk ownership by a node is determined by where the FC HBAs connect into the switch and by where the disks connect to the switch banks. FC HBAs connect to one bank, and own the disks on the other switch bank. For example, if the local node has its FC HBA adapters connect to Bank 1, it owns the disks connected to Bank 2.

Before you begin

To cable a Fabric Attached MetroCluster, you must be familiar with clustering concepts, the Brocade command-line interface, and synchronous mirroring. You must also have the following information:

- ◆ Unique domain IDs for the switches
- ◆ Ethernet IP address for both the switches and nodes
- ◆ Ethernet subnetmask
- ◆ Gateway address

For detailed information

For detailed information about setting up and cabling your Fabric Attached MetroCluster, see the following sections:

- ◆ [“Step 1: Planning the Fabric Attached MetroCluster”](#) on page 100
- ◆ [“Step 2: Configuring the local switches”](#) on page 101
- ◆ [“Step 3: Cabling the local node”](#) on page 104
- ◆ [“Step 4: Cabling the remote switches and node”](#) on page 107
- ◆ [“Step 5: Configuring the nodes”](#) on page 111

Step 1: Planning the Fabric Attached MetroCluster

Reason to plan your configuration You must identify several pieces of information that you use during configuration procedures. Recording this information can reduce configuration errors.

FA MetroCluster worksheet The Fabric Attached MetroCluster worksheet is a planning aid to record specific cabling information about your Fabric Attached MetroCluster.

Local node name: _____

Cluster adapter Port 1 cabled to Switch 1 Port__

Switch 1 name: _____

Switch 1 IP address:_____

FC HBA 1 port 1 cabled to Switch 1 Port__

FC HBA 1 port 2 cabled to Switch 2 Port__

Pool 0 disk shelves:

Disk 1 connected to Port ____

Disk 2 connected to Port ____

Remote node name: _____

Cluster adapter Port 1 is cabled to Switch 3 Port __

Switch 3 name: _____

Switch 3 IP address:_____

FC HBA 1 port 1 cabled to Switch 3 Port__

FC HBA 1 port 2 cabled to Switch 4 Port__

Pool 0 disk shelves:

Disk 1 connected to Port ____

Disk 2 connected to Port ____

Cluster adapter Port 2 cabled to Switch 2 Port __

Switch 2 name: _____

Switch 2 IP address:_____

FC HBA 2 port 1 cabled to Switch 1 Port__

FC HBA 2 port 2 cabled to Switch 2 Port__

Pool 1 disk shelves:

Disk 3 connected to Port ____

Disk 4 connected to Port ____

Cluster adapter Port 2 is cabled to Switch 4 Port__

Switch 4 name: _____

Switch 4 IP address:_____

FC HBA 2 port 1 cabled to Switch 3 Port

FC HBA 2 port 2 cabled to Switch 4 Port

Pool 1 disk shelves:

Disk 3 connected to Port ____

Disk 4 connected to Port ____

Cabling a new Fabric Attached MetroCluster

Step 2: Configuring the local switches

Tasks for
configuring the
switch

To configure the Brocade 3800 or 3200 switch for clustering, complete the following steps.

Note_____

Check the NOW site at <http://now.netapp.com/NOW> to confirm the following procedure.

Step	Action	
1	Log in to the switch, using the console port and hyperterm, and set the appropriate variables and IP addresses, as described in the appropriate Brocade switch hardware reference manual.	
	Note _____	
	Use telnet for subsequent switch login sessions, using the switch IP address.	
2	Check the switch firmware by entering the following command: version Result: The currently installed switch firmware is displayed. See the NOW Web site at http://now.netapp.com/NOW for the current version of switch firmware.	
3	If the switch firmware is...	Then...
	Current	Go to Step 4.
	Out of date	1. Download the switch firmware from the NOW Web site at http://now.netapp.com/NOW and install it, as described in the appropriate Brocade switch hardware reference manual. 2. Reboot the switch.

Step	Action
4	Disable the switch by entering the following command: switchdisable
5	Clear any existing configuration on the switch by entering the following command: cfgclear
6	Ensure that the switch is configured with default settings by entering the following command: configdefault
7	Set the switch parameters by entering the following command: configure Result: You should only set the following parameters: ◆ Fabric parameters = y ◆ domain_id = XXX (See your system administrator for this value.) ◆ Disable device probing = 1
8	Exit the configuration utility by entering the following command: ctrl-d
9	Reboot the switch so that the default settings take effect by entering the following command: fastboot
10	Log in to the switch and disable the switch by entering the following command: switchdisable
11	Set the switch port speeds to half duplex by entering the following command for each switch port into which a disk is plugged: portcflport port# 0,0,1 Note _____ You only need to do this command on the ports to which disks are attached. _____

Step	Action	
12	If the distance between switches is...	Then...
	Between 500 meters and 10 kilometers	Set the ISL port distance on each ISL port to short distance by entering the following command: portCfgLongDistance ISL-port#, "LE"
	Between 10 and 30 kilometers	Set the ISL port distance on each ISL port to long distance by entering the following command: portCfgLongDistance ISL-port#, "L1"
13	Save the new switch configuration by entering the following command: cfgsave	
14	Enable the switch by entering the following command: switchenable	
15	Set the switch name by entering the following command: switchname Metro1_switch1 <i>Metro1_switch1</i> is the name you gave to the current switch. Use a name that is meaningful, such as node name and switch, as in the previous example.	
16	Check the switch settings by entering the following command: configshow	
17	Repeat Steps 1 through 16 on the second switch on the local node.	
18	Go to “Step 3: Cabling the local node” on page 104	


Cabling a new Fabric Attached MetroCluster

Step 3: Cabling the local node

Cabling the local node

When cabling the node, you must also cable the disk shelves, and switch. To cable the node, complete the following steps, using the illustration for reference.

Note—
You can only cable a maximum of two disk shelves per loop.



Step	Action
1	Cable the first FC HBA to Bank 1 of Switch 1 and Switch 2. In the example, Port A of the FC HBA is cabled to Port 0 of Switch 1 and Port B is cabled to Port 0 of Switch 2.
2	Cable the second FC HBA to Port 1 on both switches, with Port A on Switch 1 and Port B on Switch 2.
3	Connect the Pool 0 disks to Bank 2, Pool 0 of the switches by completing the following steps: 1. Connect the Input port of LRC/ESH A to port in Pool 0 in Bank 1 on Switch 2. 2. Connect the Input port of LRC/ESH B to the port in Pool 0 Bank 1 on Switch 1. This must be the same port number you used for the LRC/ESH A cabling. 3. Connect the first disk shelf to the next disk shelf in the loop by connecting the Output port of the LRC/ESH of the first shelf to the corresponding Input Port of the next disk shelf's LRC/ESH. In the example, the Pool 0 disk shelves are connected to Switch Port 8 on both Switch 1 and 2.
4	Connect the Pool 1 disks to Bank 1, Pool 1 of the switches by completing the following steps: 1. Connect the Input port of LRC/ESH A to port in Pool 1 in Bank 1 on Switch 2. 2. Connect the Input port of LRC/ESH B to port in Pool 1 in Bank 1 on Switch 1. This must be the same port number you used for the LRC/ESH A cabling. 3. Connect the first disk shelf to the next disk shelf in the loop by connecting the Output port of the LRC/ESH of the first shelf to the corresponding Input Port of the next disk shelf's LRC/ESH. In the example, the disk shelves are connected to Switch Port 4 on both Switch 1 and 2.

Step	Action																																																
5	Confirm that the disks are visible and dual pathed by entering the following command on the console: storage show disk -p Results: The output shows the disks connected to the switches, to what port they are connected, and to what disk and to which LRC/ESH they belong, as shown in the following example: Example: <pre>ha15*> storage show disk -p</pre> <table><thead><tr><th>PRIMARY</th><th>PORT</th><th>SECONDARY</th><th>PORT</th><th>SHELF</th><th>BAY</th></tr><tr><th>-----</th><th>----</th><th>-----</th><th>----</th><th>----</th><th>--</th></tr></thead><tbody><tr><td>ha15sw:2.40</td><td>A</td><td>ha16sw:2.40</td><td>B</td><td>5</td><td>0</td></tr><tr><td>ha16sw:2.41</td><td>B</td><td>ha15sw:2.41</td><td>A</td><td>5</td><td>1</td></tr><tr><td>.</td><td></td><td></td><td></td><td></td><td></td></tr><tr><td>.</td><td></td><td></td><td></td><td></td><td></td></tr><tr><td>.</td><td></td><td></td><td></td><td></td><td></td></tr><tr><td>ha15sw:10.52</td><td>B</td><td>ha16sw:10.52</td><td>A</td><td>6</td><td>4</td></tr></tbody></table>	PRIMARY	PORT	SECONDARY	PORT	SHELF	BAY	-----	----	-----	----	----	--	ha15sw:2.40	A	ha16sw:2.40	B	5	0	ha16sw:2.41	B	ha15sw:2.41	A	5	1	.						.						.						ha15sw:10.52	B	ha16sw:10.52	A	6	4
PRIMARY	PORT	SECONDARY	PORT	SHELF	BAY																																												
-----	----	-----	----	----	--																																												
ha15sw:2.40	A	ha16sw:2.40	B	5	0																																												
ha16sw:2.41	B	ha15sw:2.41	A	5	1																																												
.																																																	
.																																																	
.																																																	
ha15sw:10.52	B	ha16sw:10.52	A	6	4																																												
6	Connect one ISL cable to any port on each switch. In the example, Port 13 is used on both switches.																																																
7	Connect the VI-MC cluster adapter to the switches, with Port 1 connecting to Switch 1 and Port 2 to Switch 2.																																																
8	Check to make sure the node can see all disk shelves by entering the following command on the console: vol status -r Results: Both disk pools are displayed. If some disk shelves are not displayed, check the cabling on the disk shelf and the switch.																																																
9	Go to the remote site and cable and set up the remote node.																																																

Step 4: Cabling the remote switches and node


Configuring the remote switches

To configure the remote switches, complete the following steps.

Step	Action
1	Log into Switch 3 and set it up as described in “ Step 2: Configuring the local switches ” on page 101.
2	Repeat Step 1 for Switch 4.

Cabling the remote node and disks

To cable the node, disk shelves, and switch, complete the following steps, using the illustration for reference.



Step	Action
1	Cable the first FC HBA to Bank 2 of Switch 3 and Switch 4. In the example, Port A of the FC HBA is cabled to Port 8 of Switch 1 and Port B is cabled to Port 8 of Switch 4.
2	Cable the second FC HBA to Port 9 on both switches, with Port A on Switch 4 and Port B on Switch 4.

Step	Action
3	Connect the Pool 0 disks to Bank 1, Pool 0 of the switches by completing the following steps: 1. Connect the Input port of LRC/ESH A to port in Pool 0 in Bank 1 on Switch 3. 2. Connect the Input port of LRC/ESH B to port in Pool 0 in Bank 1 on Switch 4. This must be the same port number you used for the LRC/ESH A cabling. 3. Connect the first disk shelf to the next disk shelf in the loop by connecting the Output port of the LRC/ESH of the first shelf to the corresponding Input Port of the next disk shelf's LRC/ESH. In the example, the Pool 0 disk shelves are connected to Switch Port 0 on both Switch 3 and 4.
4	Connect the Pool 1 disks to Bank 2, Pool 1 of the switches by completing the following steps: 1. Connect the Input port of LRC/ESH A to port in Pool 1 in Bank 2 on Switch 3. 2. Connect the Input port of LRC/ESH B to port in Pool 1 in Bank 2 on Switch 4. This must be the same port number you used for the LRC/ESH A cabling. 3. Connect the first disk shelf to the next disk shelf in the loop by connecting the Output port of the LRC/ESH of the first shelf to the corresponding Input Port of the next disk shelf's LRC/ESH. In the example, the Pool 1 disk shelves are connected to Switch Port 12 on both Switch 3 and 4.

Step	Action																																																
5	Confirm that the disks are visible and dual pathed by entering the following command on the console: storage show disk -p Results: The output shows the disks connected to the switches, to what port they are connected, and to what disk and to which LRC/ESH they belong, as shown in the following example: Example: <pre>ha16*> storage show disk -p</pre> <table><thead><tr><th>PRIMARY</th><th>PORT</th><th>SECONDARY</th><th>PORT</th><th>SHELF</th><th>BAY</th></tr><tr><th>-----</th><th>----</th><th>-----</th><th>----</th><th>----</th><th>--</th></tr></thead><tbody><tr><td>ha15sw:2.40</td><td>A</td><td>ha16sw:2.40</td><td>B</td><td>5</td><td>0</td></tr><tr><td>ha16sw:2.41</td><td>B</td><td>ha15sw:2.41</td><td>A</td><td>5</td><td>1</td></tr><tr><td>.</td><td></td><td></td><td></td><td></td><td></td></tr><tr><td>.</td><td></td><td></td><td></td><td></td><td></td></tr><tr><td>.</td><td></td><td></td><td></td><td></td><td></td></tr><tr><td>ha15sw:10.52</td><td>B</td><td>ha16sw:10.52</td><td>A</td><td>6</td><td>4</td></tr></tbody></table>	PRIMARY	PORT	SECONDARY	PORT	SHELF	BAY	-----	----	-----	----	----	--	ha15sw:2.40	A	ha16sw:2.40	B	5	0	ha16sw:2.41	B	ha15sw:2.41	A	5	1	.						.						.						ha15sw:10.52	B	ha16sw:10.52	A	6	4
PRIMARY	PORT	SECONDARY	PORT	SHELF	BAY																																												
-----	----	-----	----	----	--																																												
ha15sw:2.40	A	ha16sw:2.40	B	5	0																																												
ha16sw:2.41	B	ha15sw:2.41	A	5	1																																												
.																																																	
.																																																	
.																																																	
ha15sw:10.52	B	ha16sw:10.52	A	6	4																																												
6	Connect the ISL from Switch 1 to Switch 3 and the ISL from Switch 2 to Switch 4.																																																
7	Connect the VI-MC cluster adapter to the switches, with Port 1 connecting to Switch 3 and Port 2 to Switch 4.																																																
8	Check to make sure the node can see all disk shelves by entering the following command on the console: vol status -r																																																
9	Connect the VI-MC cluster adapter to the switches, with Port 1 connecting to Switch 3 and Port 2 to Switch 4.																																																
10	Go to “ Step 5: Configuring the nodes ” on page 111.																																																

Step 5: Configuring the nodes

Configuring the nodes

After cabling and booting the node, you must configure it and the disk shelves, and install any software that will be running on the local node.

To configure the node, complete the following steps.

Step	Action
1	Install the new Data ONTAP release on the node and run setup.
2	License cluster and syncmirror_local on the local node, by entering the following command: license add xxxxx xxxxx is the license code you received for the cluster feature.
3	Reboot the local node by entering the following command: reboot
4	Repeat Steps 1 through 3 for the remote node.
5	License cluster_remote on both nodes.
6	Halt both nodes, by entering the following command from each console window: halt
7	Reboot the nodes.

Step	Action
8	Set up your mirrored volumes by entering the following command from each console: <pre>vol create vol-name -m ndisks[@disk-size]</pre> Note_____ The root volumes on both nodes must be mirrored. _____ Note_____ If you have existing volumes, use the vol mirror vol-name command. _____
9	Configure both nodes, as described in “ Configuring the cluster ” on page 42.
10	Run the cluster configuration checker, as described in “ Using the cluster configuration checker ” on page 43. Correct any configuration errors that are displayed.
11	Enable clustering by entering the following command on the console: <pre>cf enable</pre>
12	Check to make sure that the node can see all disk shelves by entering the following command on the console: <pre>vol status -r</pre> Results: Both local and partner disks should be displayed. The local volumes should be mirrored and there should be a listing of both disk pools. It should also list spares disks for both pools.

Cabling a non-switched MetroCluster

Cabling a non-switched MetroCluster

To cable a non-switched MetroCluster, complete the following steps.

Step	Action
1	Prepare your sites, as described in “ Planning your mirrored cluster ” on page 76.
2	Cable the Channel A loops of the MetroCluster, as described in “ Step 1: Cabling the Channel A loop ” on page 81.
3	Cable the nodes together, as described in “ Step 2: Cabling the nodes together ” on page 83.
4	Cable the Channel B loops of the MetroCluster, as described in “ Step 3: Cabling the Channel B loop ” on page 84.
5	Configure the cluster, as described in “ Step 4: Completing setup and testing your mirrored cluster ” on page 86. Make sure that you enable the syncmirror_local and cluster_remote licenses on both nodes.

Cabling a MetroClusters with QuickLoop

About this section

This section describes how to configure and cable a MetroCluster using Quickloop. MetroClusters installed after the release of Data ONTAP 6.4.1 are Fabric Attached MetroClusters. See [“Cabling a new Fabric Attached MetroCluster”](#) on page 97 for instructions about cabling a Fabric Attached MetroCluster.

Before you begin

To cable a MetroClusters with QuickLoop, you must be familiar with clustering concepts, the Brocade command-line interface, and synchronous mirroring. You must also have the following information:

- ◆ Unique domain IDs for the switches
- ◆ Ethernet IP address
- ◆ Ethernet subnetmask
- ◆ Gateway address

Note

If the nodes in your cluster are separated by less than 500 meters, see [“System requirements for mirrored clusters”](#) on page 72 for a list of required hardware, software, settings, and setup information.

For detailed information

For detailed information about setting up and cabling your MetroClusters with QuickLoop, see the following sections:

- ◆ [“Step 1: Planning the MetroCluster with QuickLoop”](#) on page 115
- ◆ [“Step 2: Installing and preparing the local node and disk shelves”](#) on page 118
- ◆ [“Step 3: Configuring the switch”](#) on page 119
- ◆ [“Step 4: Cabling the node, disk shelves, and switch”](#) on page 121
- ◆ [“Step 5: Configuring the local node”](#) on page 124
- ◆ [“Step 6: Setting up and configuring the remote part of the cluster”](#) on page 128

Step 1: Planning the MetroCluster with QuickLoop

Reason to plan your configuration You must identify or define several pieces of key information that you use during configuration procedures. Recording this information in advance reduces the possibility of configuration errors when you configure the remote node.

MetroCluster with QuickLoop worksheet The MetroCluster with QuickLoop worksheet is a planning aid for recording cabling information about your MetroCluster. The following example shows a completed worksheet and an illustration showing a cabled MetroCluster.

Local node name: *Metro1*

VI cluster adapter is cabled to switch port number	5
FC HBA adapters are installed in slot numbers	5 (Adapter 1) and 7 (Adapter 2)
Adapter 1, port 1 is cabled to	<i>Ch A, Shelf 1</i>
Adapter 1, port 2 is cabled to	<i>Swt Prt 2</i>
First switch port, port number 1 is cabled to	Second switch port, port number 2 is cabled to
FC-AL adapter in slot __port number__ or	FC-AL adapter in slot 7 port number 2 or
Disk shelf ID number 1 Channel <i>B</i>	Disk shelf ID number __Channel __
Third switch port, port number 3 is cabled to	Fourth switch port, port number 4 is cabled to
FC-AL adapter in slot 5 port number 2 or	FC-AL adapter in slot __port number__ or
Disk shelf ID number __Channel __	Disk shelf ID number 2 Channel <i>A</i>

Local switch name: <i>Metro1_switch</i>	Remote switch name: <i>Metro2_switch</i>
Local switch WWN <i>10:00:00:60:69:00:60:11</i>	Remote switch WWN <i>10:00:00:60:69:00:60:12</i>
Zone 1: Name <i>Metro1_local</i> _____	Port number 1 Local switch domain ID <i>111</i>
Zone 2: Name <i>Metro1_remote</i> _____	Port number 2
Zone 3: Name <i>Metro2_local</i> _____	Port number 3 Remote switch domain ID <i>222</i>
Zone 4: Name <i>Metro2_remote</i> _____	Port number 4
Zone 5: Name <i>VI_Clus</i> _____	Port number 5

ISL 1 cabled to switch port number	14
ISL 2 cabled to switch port number	15

Remote node name: *Metro2*

VI cluster adapter is cabled to switch port number (must match local node connection) 5

FC HBA adapters are installed in slot numbers 5 (Adapter 1) and 7 (Adapter 2)

Adapter 1, port 1 is cabled to *Swt Prt 1*

Adapter 2, port 1 is cabled to *Swt Prt 4*

Adapter 1, port 2 is cabled to *Ch A, Shelf 4*


Adapter 2, port 2 is cabled to *Ch B, Shelf 3*

First switch port, port number 1 is cabled to
FC-AL adapter in slot 5 port number 1 or
Disk shelf ID number Channel

Second switch port, port number 2 is cabled to
FC-AL adapter in slot port number or
Disk shelf ID number 3 Channel A

Third switch port, port number 3 is cabled to
FC-AL adapter in slot port number or
Disk shelf ID number 3 Channel B

Fourth switch port, port number 4 is cabled to
FC-AL adapter in slot 7 port number 1 or
Disk shelf ID number Channel



Using the worksheet

Complete the following worksheet by entering the information for your MetroCluster with QuickLoop installation.

Local node name: _____

VI cluster adapter is cabled to switch port number _____

FC HBA adapters are installed in slot numbers _____ (Adapter 1) and _____ (Adapter 2)

Adapter 1, port 1 is cabled to _____

Adapter 2, port 1 is cabled to _____

Adapter 1, port 2 is cabled to _____

Adapter 2, port 2 is cabled to _____

First switch port, port number ____ is cabled to

Second switch port, port number ____ is cabled to

FC-AL adapter in slot __port number__ or

FC-AL adapter in slot __port number__ or

Disk shelf ID number __Channel __

Disk shelf ID number __Channel __

Third switch port, port number ____ is cabled to

Fourth switch port, port number ____ is cabled to

FC-AL adapter in slot __port number__ or

FC-AL adapter in slot __port number__ or

Disk shelf ID number __Channel __

Disk shelf ID number __Channel __

Local switch name: _____

Remote switch name: _____

Local switch WWN _____

Remote switch WWN _____

Zone 1: Name _____ Port number _____ Local switch domain ID _____

Zone 2: Name _____ Port number _____

Zone 3: Name _____ Port number _____ Remote switch domain ID _____

Zone 4: Name _____ Port number _____

Zone 5: Name _____ Port number _____

ISL 1 cabled to switch port number _____

ISL 2 cabled to switch port number _____

Remote node name: _____

VI cluster adapter is cabled to switch port number (must match local node connection) _____

FC HBA adapters are installed in slot numbers _____ (Adapter 1) and _____ (Adapter 2)

Adapter 1, port 1 is cabled to _____

Adapter 2, port 1 is cabled to _____

Adapter 1, port 2 is cabled to _____

Adapter 2, port 2 is cabled to _____

First switch port, port number ____ is cabled to

Second switch port, port number ____ is cabled to

FC-AL adapter in slot __port number__ or

FC-AL adapter in slot __port number__ or

Disk shelf ID number __Channel __

Disk shelf ID number __Channel __

Third switch port, port number ____ is cabled to

Fourth switch port, port number ____ is cabled to

FC-AL adapter in slot __port number__ or

FC-AL adapter in slot __port number__ or

Disk shelf ID number __Channel __

Disk shelf ID number __Channel __

Step 2: Installing and preparing the local node and disk shelves

Installing the local node and disk shelves

To install the local node and accompanying disk shelves, complete the following steps.

Step	Action
1	Install the switch in either an equipment rack or system cabinet, as described in the appropriate Brocade switch hardware reference manual.
2	Connect the switch to a system console and, using hyperterm, establish communication with the switch, as described in the appropriate Brocade switch hardware reference manual.
3	Install the local node in either an equipment rack or system cabinet, as described in the <i>F800 Series Hardware Guide</i>. Caution— Do not boot the node.
4	Install the disk shelves into the equipment rack, as described in the <i>DS14 Hardware Guide</i>. Note— Each disk shelf is terminated in a MetroCluster. Caution— Do not turn on the disk shelves.
5	Set a unique disk shelf ID, from 1 to 7, for each disk shelf in the cluster. Caution— You cannot have duplicate disk shelf IDs anywhere within the cluster, even if the disk shelves are in different loops.
6	Go to “ Step 3: Configuring the switch ” on page 119.

Step 3: Configuring the switch

Tasks for configuring the switch

To configure the Brocade 3800 or 3200 switch for clustering, you must complete the following tasks:

- ◆ Configure the switch
- ◆ Connect the switch
- ◆ Enable and configure the qloop feature, set up zoning, and create the configuration

Configuring the switch

To configure a switch for a MetroCluster, follow the procedures for switch configuration on the NOW site at <http://now.netapp.com/>.

Connecting the ISLs

To connect the switch to the two ISL (Inter Switch Link) cables, complete the following steps.

Step	Action
1	Connect the long distance SFPs to the switch, making sure that the key is on top, pushing firmly until it is seated and the latch clicks. Example: The first ISL is cabled into Port 14 on the switch.
2	Connect one of the LC/LC ISL cables to the SFP in a selected port. Note the port, because the other end of the ISL cable must be connected to the corresponding port on the partner's switch.
3	Connect another LC/LC ISL cable to an SFP, to ensure redundancy. Example: The second ISL is cabled into Port 15 on the switch.


Step	Action
4	Check the switch by entering the following command on the console: switchshow Result: The output shows the current switch configuration for the ISL connection, as shown in the following example: <pre>admin> switchshow switchName: Metro1_switch switchType: 16.2 switchState: Online switchMode: Native switchRole: Principal switchDomain: 111 switchId: fffcde switchWwn: 10:00:00:60:69:00:60:11 switchBeacon: OFF Zoning: OFF</pre>
5	Disable the switch by entering the following command on the console: switchdisable
6	Set the distance and speeds for the ISL ports. See the NOW Web site at http://now.netapp.com for the specific configuration procedures.

Cabling a MetroClusters with QuickLoop

Step 4: Cabling the node, disk shelves, and switch

Cabling the node,
disk shelves, and
switch

To cable the node, disk shelves, and switch, complete the following steps, using the illustration for reference.



Step	Action
1	Cable one end of the SC/LC VI cluster adapter cable to the VI cluster adapter, then cable the other end of the cable to the switch port. In the example, the port was port 5 on the switch.

Step	Action
2	Connect the local disk, pool 0 of the mirror, by completing the following steps: 1. Connect the Channel A loop by cabling one end of an SC/LC FC-AL to disk shelf cable to Port 1 on the FC HBA adapter in slot 5, and the other end of the cable to the Channel A LRC of the local disk shelf. In the example, the local disk shelf is identified as M1_local. 2. Connect the Channel B loop by cabling one end of an SC/LC disk shelf-to-switch cable to the Channel B In port on the disk shelf and the other end to the switch port for the local disk. The example shows switch port Port 1. Note _____ You should see green LEDs at the cable connections. If they do not light, reseal the cable and SFPs. If they still do not light, replace the cable or SFP.
3	Connect the partner's remote disk, pool 1 of the mirror, by completing the following steps: 1. Connect the Channel A loop by cabling an SC/LC disk shelf-to-switch cable to the Channel A port of the partner disk shelf and the other end to the switch port for the partner disk. The example shows switch port 4. 2. Connect the Channel B loop by connecting an SC/LC FC-AL to disk shelf cable to Port 1 on the FC HBA adapter in slot 7 and the other end to the In port of the Channel B LRC of the partner disk.

Step	Action
4	Connect the FC-AL adapters to the switch by completing the following steps: 1. Connect the switch to the FC-AL adapter by cabling one end of an LC/LC FC-AL to switch cable to Port 2 on the FC HBA adapter in slot 5, and the other end to the port on the switch reserved for the adapter. The example shows switch port 3. 2. Connect the switch to the FC-AL adapter by cabling one end of an LC/LC FC-AL to switch cable to Port 2 on the FC HBA adapter in slot 7 and the other end of the cable to the port on the switch reserved for the adapter. The example shows switch port 2.
5	Make sure that you terminate the loops in each disk shelf.

Cabling a MetroClusters with QuickLoop

Step 5: Configuring the local node

Configuring the node

To configure the local node, complete the following steps.

Step	Action
1	Boot the node and disk shelves, then configure the local plexes, as described in the <i>Data Protection Guide</i> .
2	Enable clustering and other licenses on the node by entering the following command from the node console: <pre>license add xxxxx</pre> xxxxx is the license code you received for the cluster feature. Note——— If you are using a version of Data ONTAP earlier than 6.4, the command is license option=code. —————
3	View the options set on the node by entering the following command at the node console: <pre>options</pre> Result: The current option settings on the node are listed. You should see output similar to the following: <pre>autosupport.doit DONT autosupport.enable off . . dns.domainname filer1 (might be overwritten in takeover) dns.enable off (might be overwritten in takeover) ftpd.enable off . .</pre>

Step	Action
4	Set the node options by entering the following command on the node console: options <i>option_name</i> <i>option_value</i> Note_____ Make sure to set the <code>cf.quickloop.enable</code> variable to True when setting variables on the node. _____


Step	Action
5	Check the configuration by entering the following command on the node console: sysconfig -v Example: <pre>metro1> sysconfig -v NetApp Release Data ONTAP 6.4.1: Wed Apr 30 13:36:54 PST 2002 System ID: 0033586696 (metro1); partner ID: 0033582111 (metro2) System Serial Number: 6112 (metro1) System Rev: A0 . . . slot 5: FC Host Adapter 5a (Dual-channel, QLogic 2312 rev. 2, 64-bit, L-port, <UP>)Firmware rev:2.1.32 . . . 61: NETAPP X235_SCHT5073F10 NA04 68.0GB (140174232520B/sect) slot 7: FC Host Adapter 7a (Dual-channel, QLogic 2312 rev. 2, 64-bit, L-port, <OFFLINE (hard)>) Firmware rev: 3.1.15 Host Loop Id: 0 FC Node Name: 2:000:00e08b:07a1a8 Cacheline size: 8 FC Packet size: 2048 SRAM parity: Yes External GBIC: No Link Data Rate: 1 Gbit . . . slot 11: NVRAM (NetApp NVRAM III) . . .</pre> Result: The configuration is correct if you see the following information: ◆ Local node and local disks ◆ Partner node disks ◆ Any empty adapters

Step	Action
6	Check the switch configuration by entering the following command on the switch console: switchshow Example: The example display shows the output for the command for the only cabled components of the MetroCluster: <pre>metro1:admin> switchshow switchName: metro1 switchType: 16.2 switchState: Online switchMode: Native switchRole: Principal switchDomain: 111 switchId: fffcde switchWwn: 10:00:00:60:69:00:60:11 switchBeacon: OFF Zoning: ON (metro1_metro2) port 0: id N2 L-Port No Module port 1: id N1 Online L-Port 12 private port 2: id 1G Online L-Port 1 private port 3: id 1G Online L-Port 1 private port 4: id N1 Online L-Port 12 private port 5: id N1 Online L-Port 1 private port 14: id N1 Online E-Port 1 10:00:00:60:69:00:60:11 "Metro1" port 15: id N1 Online E-Port 1 10:00:00:60:69:00:60:12 "Metro2" (upstream)</pre> Result: The configuration is correct if you see the following information: ◆ ISL ports, labeled E-Port, showing the switch address ◆ Two or more disk ports, labeled L-Port, followed by the number of disks in the disk shelf attached to the port ◆ VI cluster adapter port, labeled L-Port
7	Shut down the node to the OK prompt from the node console.
8	Go to the remote site and complete “ Step 6: Setting up and configuring the remote part of the cluster ” on page 128.

Step 6: Setting up and configuring the remote part of the cluster

Setting up the remote node

To set up the remote portion of the MetroCluster, complete the following steps, using the illustration for reference.



Step	Action
1	Install the node, disks, and switch in the equipment rack or system cabinet at the remote site, as described in “ Step 2: Installing and preparing the local node and disk shelves ” on page 118.
2	Configure the switch, as described in “ Step 3: Configuring the switch ” on page 119.

Step	Action
3	Connect the SFPs to the ports on the target ports on the switch, then connect the ISL cables to the same ports as used in the local switch. Example: If <i>Metro1_switch</i> uses ports 14 and 15, <i>Metro2_switch</i> must use ports 14 and 15 for the ISL connection. Note _____ Zone information is transferred from the local switch to the remote switch after communication between the two switches is established.
4	Check to make sure that there is a single fabric between the switches by entering the following command: fabricshow Example: <pre>tort1_8:admin> fabricshow Switch ID Worldwide Name Enet IP Addr FC IP Addr Name ----- 111: fffc6f10:00:00:60:69:00:60:11 172.27.100.85 0.0.0.0 "metro1_switch" 222: fffcde10:00:00:60:69:00:60:12 172.27.100.86 0.0.0.0 "metro2_switch"</pre>
5	Enable the ports on the switch by entering the following command: portenable # # is the number of the port that you want enabled.
6	Save the configuration by entering the following command: cfgsave
7	Check the configuration by entering the following command: cfgshow
8	Confirm and enable the configuration by entering the following command: cfgenable

Step	Action
9	Connect the local disk, pool 0 of the mirror, by completing the following steps: 1. Connect the Channel A loop by cabling one end of an SC/LC FC-AL to disk shelf cable to Port 2 on the FC HBA adapter in slot 5, and the other end of the cable to the Channel A LRC of the local disk shelf. In the example, the local disk shelf is identified as M2_local. 2. Connect the Channel B loop by cabling one end of an SC/LC disk shelf-to-switch cable to the Channel B In port on the disk shelf and the other end to the switch port for the local disk. The example shows switch port Port 3.
10	Connect the partner's remote disk, pool 1 of the mirror, by completing the following steps: 1. Connect the Channel A loop by cabling an SC/LC disk shelf-to-switch cable to the Channel A port of the partner disk shelf and the other end to the switch port for the partner disk. The example shows switch port 2. 2. Connect the Channel B loop by connecting an SC/LC FC-AL to disk shelf cable to Port 2 on the FC HBA adapter in slot 7 and the other end to the In port of the Channel B LRC of the partner disk.
11	Connect the FC-AL adapters to the switch by completing the following steps: 1. Connect the switch to the FC-AL adapter by cabling one end of an LC/LC FC-AL to switch cable to Port 1 on the FC HBA adapter in slot 5, and the other end to the port on the switch reserved for the adapter. The example shows switch port 1. 2. Connect the switch to the FC-AL adapter by cabling one end of an LC/LC FC-AL to switch cable to Port 1 on the FC HBA adapter in slot 7, and the other end to the port on the switch reserved for the adapter. The example shows switch port 4.
12	Make sure that you terminate each desk shelf on all the loops.

Converting an existing cluster to a Fabric Attached MetroCluster

Converting a cluster

To convert an existing cluster into a Fabric Attached MetroCluster, complete the following steps.

Note_____

If you want to convert a switched MetroCluster to a Fabric Attached MetroCluster, contact Customer Support for assistance.

Step	Action
1	Update Data ONTAP, filer firmware, and disk firmware, as described in “Updating your software” on page 35. Make sure to shut down the nodes to the OK prompt. Note_____ Make sure that you halt the nodes and power cycle the disk shelves after installing updated disk firmware.
2	Update your hardware, as described in “Updating your hardware” on page 31. Note_____ You must have FC HBA for Disk adapters installed in each node and each node must have VI-MC cluster adapters.
3	Determine your switch and cluster configuration by completing the planning worksheet in “Step 1: Planning the Fabric Attached MetroCluster” on page 100.
4	Set up and configure the local switches, as described in “Step 2: Configuring the local switches” on page 101.
5	Cable the local node, as described in “Step 3: Cabling the local node” on page 104.

Step	Action	
6	If you are converting a...	Then...
	Standard cluster	1. Set up mirroring and configure the local node, as described in “Step 3: Cabling the local node” on page 104. 2. Go to Step 7.
	Non-switched MetroCluster	Go to Step 7.
7	Transport the partner node, remote node’s Pool 0 and the local node’s Pool 1 disk shelves, and the third and fourth switches to the remote location.	
8	Set up the remote node, disk shelves, and switch, as described in “ Step 4: Cabling the remote switches and node ” on page 107.	
9	Configure the cluster as described in “ Step 5: Configuring the nodes ” on page 111.	

Replacing MetroCluster components

Replacing a node or node component

To replace a node or node component, complete the following steps.

Step	Action	
1	Identify the failed node or component, then have its partner take it over, by entering the following command on the partner node console: cf takeover	
2	Shut down the node and disconnect it from the power sources.	
3	Disconnect the cabling to the node, marking each cable you disconnect.	
4	If the replacement is...	Then...
	The node	1. Replace the node, as described in the <i>F800 Series Field Service Guide</i> . 2. Check configuration, making sure that all licenses are activated and all options are set correctly. 3. Go to Step 4.
	A node component	1. Replace the node component, as described in the <i>F800 Series Field Service Guide</i> . 2. Go to Step 4.
5	Reconnect the cables you removed in Step 2 and apply power to the node.	
6	Reboot the node and enable clustering, if it was disabled.	

Interpreting hot swap disk loop error messages

About the error messages

A Fabric Attached MetroCluster generates error messages when a disk drive is removed from or added to a disk loop. These messages are informational and require no action from the user.

Messages and descriptions

The following table lists and describes informational error messages you might encounter when removing or inserting a disk into a loop.

Error message	Description
[node1: ispf_timeout_1:warning]: node2:11.34 (NxN2N2Nbae) (NxN359da8N,Nx28:NN6Ne3dN:N1NN,N/N,4951/N/N,29927/N): command timeout, quiescing drive to allow outstanding I/O to complete.	When the disk drive is pulled from the bay, data flow to all drives in the loop is interrupted. The error message indicates that the data flow was interrupted and the system is trying to reestablish contact with that loop.
[node1: ispf_timeout_1:error]: node2:11.34 (NxN2N2Nbae): global device timer timeout, initiating device recovery.	The system is attempting to recover from the data flow interruption.
[node1: scsi.cmd.abortedByHost:error]: Device node2:11.34: Command aborted by host adapter: HA status Nx4: cdb Nx28:NN6Ne3dN:N1NN.	The current command to the drive loop is aborted. The system is continuing to recover the drive loop.
[node1: scsi.cmd.retrySuccess:info]: Device deadhead:11.39: request successful after retry #1: cdb Nx28:N162416N:N1NN.	Contact has been reestablished with the disk loop.
[node1: ses_admin:info]: initiating rescan on channel deadhead:11	The system is scanning the disk loops to determine what has been added or removed.
[node1: raid.disk.inserted:info]: Disk deadhead:11.72 (S/N 3HZNL8HLNNNN7323DX8A) has been inserted into the system	The system found a new disk drive added to the loop.
[node1: dyn_dev_qual_admin:info]: Firmware is up-to-date on all disk drives	The system checked the firmware revisions on all disks and found that it is up to date.

About this appendix This appendix provides you with some examples of alternative FC-AL slot configuration possibilities for your F800 clusters. See the NOW Web page at <http://now.netapp.com> for additional FC-AL combinations and specific slot assignments for the FC-AL adapter.

This appendix also provides you with a series of figures illustrating how to complete the Channel B loops for clustered filers that have more than two Fibre Channel interfaces.

Topics in this appendix


The following topics are discussed in this appendix:

- ◆ “[FC-AL adapter combinations](#)” on page 136
- ◆ “[Channel B loop cabling alternatives](#)” on page 138

FC-AL adapter combinations

About this section The following section gives you examples of how you can configure either your FC-AL or dual Fibre Channel adapters in your F800 series cluster. The examples in this section are of F840 clusters. However, the F840 and F880 filers have specific slots that can be used, that are not available on other F800 series filers.

Example FC-AL adapter diagram The following illustration shows four FC-AL adapters in each filer. It is an example of one way four FC-AL adapters can be arranged in your cluster.



Additional FC-AL combination examples The following table lists other possible FC-AL combinations with two, three, or four FC-AL adapters in the local and partner nodes. See the NOW Web page at <http://now.netapp.com> if the FC-AL configuration you are considering is not in this table.

local node	partner node
One Channel A adapter and one Channel B adapter	One Channel A adapter and one Channel B adapter
Two Channel A adapters and one Channel B adapter	One Channel A adapter and two Channel B adapters
Two Channel A adapters and two Channel B adapters	Two Channel A adapters and two Channel B adapters
Three Channel A and one Channel B adapters (AAAB)	One Channel A and 3 Channel B adapters (ABBB)

local node	partner node
Three Channel A and one Channel B adapters (ABAA)	One Channel A and 3 Channel B adapters (BABB)
Three Channel A and one Channel B adapters (AABA)	One Channel A and 3 Channel B adapters (BBAB)
Three Channel A and one Channel B adapters (AAAB)	One Channel A and 3 Channel B adapters (BBBA)


Channel B loop cabling alternatives

About this section


The following section gives you an example of how you can cable a cluster that has three or four disk shelf loops. The example in this section is of an F840 cluster with DS14 disk shelves and FC-AL adapters. You can cable an F810, F820, F825, or F880 cluster in a similar fashion. You can also cable any one of these clusters with FC7, FC8, or FC9 disk shelves in a similar fashion.

F880 with three DS14 disk shelf loops diagram


The dashed line in this figure between the disk shelf on the upper right and the partner node on the left shows the first completed Channel B loop.



The dashed line in this figure between the disk shelf on the lower right and the partner node on the left shows the second completed Channel B loop in this cluster.




The dashed line in this figure between the disk shelf on the lower left and the local node on the right shows the third completed Channel B loop in this cluster.




F880 with four DS14 disk shelf loops diagram


The dashed line in this figure between the disk shelf on the lower right and the local node on the left shows the first completed Channel B loop.



The dashed line in this figure between the disk shelf on the upper right and the local node on the left shows the second completed Channel B loop in this cluster.



The dashed line in this figure between the disk shelf on the lower left and the partner node on the right shows the third completed Channel B loop in this cluster.



About this appendix This appendix lists the regulatory notices you need to be aware of when installing and operating your Network Appliance equipment.

Regulatory notices

FCC notices (U.S. only)

Network Appliance devices are designed for a CFR 47 (Code Federal Regulations) Part 15 Class A environment.

The FCC and Network Appliance guarantee the user's rights to operate this equipment only if the user complies with the following rules and regulations:

- ◆ Install and operate this equipment in accordance with the specifications and instructions in this guide.
- ◆ Modify this equipment only in the ways specified by Network Appliance.
- ◆ Use shielded cables with metallic RFI/EMI connector hoods to maintain compliance with applicable emissions standards.
- ◆ If the system has nine or more fibre channel disk shelves, install the system in two or three NetApp System Cabinets to maintain performance within Part 15 of CFR 47 regulations.

Compliance with Part 15 of CFR 47

This equipment has been tested and found compliant with Part 15 of the CFR 47 rules for Class A digital devices. These rules are designed to provide reasonable protection from interference to electronics equipment operated in a commercial environment.

Operation of this device is subject to the following two conditions:

- ◆ This device cannot cause harmful interference.
- ◆ This device must accept any interference received, including interference that may cause undesired operation.

Compliance with ICES-003

This Class A digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe A conforme à la norme NMB-003 du Canada.


Compliance with EN regulations

Marking by the symbol  indicates compliance of this Network Appliance device to the EMC Directive and the Low Voltage Directive of the European Union. Such marking is indicative that this Network Appliance device meets the technical standards listed in “Declaration of Conformity,” later in this appendix.

Caution

This is a Class A product. In a domestic environment this product may cause radio interference, in which case the user may be required to take adequate measures.


Bureau of Standards, Metrology, and Inspections notice (BSMI, Taiwan only)



Translation of the BSMI notice:

Warning: This is a Class A product. In a domestic environment this product may cause radio interference, in which case the user may be required to take adequate measures.

Voluntary Control Council for Interference by Information Technology Equipment (VCCI, Japan)



Translation of the VCCI-A notice:

This is a Class A product based on the standard of the Voluntary Control Council for Interference by Information Technology Equipment (VCCI). If this equipment is used in a domestic environment, radio disturbance may arise. If such trouble occurs, the user may be required to take corrective actions.

Declaration of Conformity

**Network Appliance, Inc.,
495 East Java Drive
Sunnyvale, California, 94089, U.S.A.,**

declare under our sole responsibility that the products

Type of equipment	Description	Model number	Year of manufacture
Network File Server	NetApp F810 Filer	F810	2001
Network File Server	NetApp F820 Filer	F820	2000
Network File Server	NetApp F825 Filer	F825	2002
Network File Server	NetApp F840 Filer	F840	2000
Network File Server	NetApp F880 Filer	F880	2001

to which this declaration relates conform to the following standards:

EN 60950:1992, with A1-A4, A11, Information Technology Equipment (Safety)

EN 55022:1998, Emissions Requirements for Information Technology Equipment

EN 50024:1998, Immunity Requirements for Information Technology Equipment

EN 60825-1, Safety of Laser/LED Equipment

IEC 1000-3-2:1995 Limits for Harmonic Current Emissions

IEC 1000-3-3:1994 Limitation of Voltage Fluctuations and Flicker in Low-Voltage Supply Systems

following the provisions of the directives listed below:

73/23/EEC, Low Voltage Directive (Product Safety)

89/336/EEC, Electromagnetic Compatibility Directive

Part number: 210-00155

Index

A

adapters

- adding or replacing in clusters 54
- See* cluster interconnect adapter
- See* FC-AL adapter

audience vii

B

booting a new cluster 40

C

cabinets, system

- completing setup 15, 86
- required cables 5, 80
- required tools and manuals 5
- setup 7
- single cabinet 3

cables

- types for MetroCluster 94
- types of 25

Channel A loops, existing in conversion 34

cluster conversion requirements 27, 35

- MetroCluster 131
- mirrored 75

clusters

- adding or replacing adapters 54
- configuration examples 135
- disabling functionality 66
- installing FC-AL adapters 55
- licensing 41

commands

- cf enable (enables clustering) 16, 87
- cf giveback (enables giveback) 18, 89
- cf nfo disable 58
- cf takeover (enables takeover) 18, 88, 133
- cfgenable 129
- cfgsave 129
- cfgshow 129
- configshow 103
- configure 102

fabricshow 129

ifconfig 49

license add xxxxx 41, 45, 111, 124

options 124

portenable 129

setup 48

switchdisable 102, 120

switchenable 103

switchname 103

switchshow 120, 127

sysconfig 126

version 101

configuration

- of local node in a switched MetroCluster 124
- using the Cluster Configuration Checker 16, 43

conventions

- command viii
- formatting viii
- keyboard viii

conversions

- determining firmware versions 29
- required hardware 23

D

dedicated interfaces

- configuring 52
- network interfaces 47
- using ifconfig 49

definitions

- disk shelf vii
- Fabric Attached MetroCluster vii
- FC HBA vii
- FC HBA for Disk vii
- filer vii
- MetroCluster vii
- node vii

disabling negotiated failover 58

disaster, recovering from a 90

disk ownership, description 98

disk shelves

- adding a DS14 58

- determining firmware version 29
- DS14-HA 61
- mixed environments 31
- dual Fibre Channel adapter
 - maximum needed for mirrored cluster 74
 - minimum needed for mirrored cluster 74
- dual MAC
 - configuring 51
 - configuring with ifconfig 49
 - defined 46

E

- equipment racks
 - basic configuration 4
 - preparing equipment 7

F

- Fabric Attached MetroCluster
 - cabling the local node and disk shelves 104
 - cabling the remot node and disk shelves 107
 - configuring the switches in 101
 - disk ownership described 98
 - planning worksheet example 100
 - software requirements 27
- FC-AL adapter
 - maximum needed for mirrored cluster 74
 - minimum needed for mirrored cluster 74
- fibre channel loops, existing 34
- filers, determining firmware version 29

G

- GBIC
 - removing 68
- giveback testing 15

H

- hardware
 - FC-AL adapter slot assignments 31
 - installing MetroCluster 118
 - installing new in the filer 32
 - required for MetroCluster 94
 - requirements for conversion 23

- verifying firmware versions 29
- hot-standby interfaces, configuring 52

I

- installation
 - of additional FC-AL adapters 32
 - of additional LRC adapters 32
 - of MetroCluster hardware 118
- interfaces 46
 - configuring 48
 - dedicated 47, 49
 - dual MAC 46, 49
 - modes 47
 - network IP addresses 46
 - shared 47, 49
 - standby 47, 49
 - VIF 28, 46
- ISL
 - cabling switches together with the 119
 - setting distance for ports on switch 120

L

- licenses
 - enabling 41
 - enabling CIFS 41, 124
 - enabling HTTP 41, 124
 - enabling NFS 41, 124
 - required for MetroCluster 96

M

- MetroClusters
 - configuring the local node in 124
 - configuring the switches in 119
 - converting to a switched 131
 - definition of Fibre Channel switched 93
 - definition of nonswitched 93
 - planning worksheet 116
 - planning worksheet example 115
 - required cables and hardware for 94
 - setting up the remote node in 128
 - software requirements 27
 - switch configuration tasks in 119
- mirrored clusters

- bringing up 86
- cabling the Channel A loops of 81
- cabling the Channel B loops of 84
- required tools and manuals for 80
- software requirements for 27

mixed disk shelf environment, configuring 31

N

network interfaces

- configuring 48
- dedicated 46, 47, 49, 52
- dual MAC 46, 49
- hot-standby 49
- IP addresses 46
- modes 47
- requirements 46
- shared 47, 49
- standby 46, 47
- types 46

NICs, types in clustering 46

R

racks, equipment

- basic configuration 4
- preparing 7

required tools and manuals

- for MetroCluster 96
- MetroCluster 94
- mirrored clusters 80
- system cabinets 5

requirements

- system cabinet cables 5, 80
- takeover 44, 48
- tools and manuals for conversions 30

S

shared network interfaces 47, 49

single NIC, defined 46

special messages ix

standby setup using ifconfig 49

switches

- cabling the local node 121
- cabling the local node and disk shelves 104
- cabling the local remote node and disk shelves 107
- checking firmware revision 101
- configuration tasks for MetroCluster 119
- configuring for Fabric Attached MetroCluster 101
- configuring for MetroCluster 119
- disabling 102
- setting the parameters 102

system cabinets

- completing setup of 15, 86
- handling single 3
- required cables for 5, 80
- required tools and manuals for 5
- setup 7

T

takeover

- configuring network interfaces 48
- requirements 44
- testing 15

terminology vii

U

updates

- of hardware 31
- of software and firmware 35

V

VI cluster adapter

- cable requirements 5, 23, 80
- removing 68
- software requirements 27

VIF 28, 46

