

FAS900 Series Appliance **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 210-00580
December 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 United States, Canada, the European Union, and Japan. 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. DataFabric, FilerView, NearStore, 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, gFiler, MultiStore, 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	v
	Safety Information (Sicherheitshinweise)	ix
Chapter 1	Preparing for the Installation.	1
	Required documentation, tools, and equipment	2
	Clustering concepts	5
	Standard clusters	6
	Mirrored clusters and MetroClusters	7
	Fabric Attached MetroCluster	9
Chapter 2	Cabling a Standard or Mirrored Cluster	13
	Cluster installation types	14
	Cabling the cluster	15
	Step 1: Preparing your equipment	17
	Step 2: Cabling Channel A	19
	Step 3: Cabling the cluster adapters	23
	Step 4: Cabling Channel B	24
Chapter 3	Cabling a MetroCluster.	29
	Cabling a new Fabric Attached MetroCluster	30
	Step 1: Planning the Fabric Attached MetroCluster	31
	Step 2: Configuring the local switches	32
	Step 3: Cabling the local node	35
	Step 4: Cabling the remote switches and node	38
	Step 5: Configuring the nodes	42
	Cabling a nonswitched MetroCluster.	44
	Converting an existing cluster to a Fabric Attached MetroCluster.	45
Chapter 4	Configuring the Cluster.	47
	Step 1: Bringing up the cluster for the first time	48
	Step 2: Enabling licenses	49

Step 3: Setting options and parameters 51

Step 4: Configuring networking interfaces 54

Step 5: Checking and testing cluster configuration. 57

Chapter 5

Basic Troubleshooting. 59

Takeover or giveback failures 60

Hot-swap disk loop error messages on Fabric Attached MetroClusters 62

Recovering from a disaster 63

Out-of-date software and firmware. 65

 Updating Data ONTAP 66

 Updating node, disk, or shelf firmware only. 67

Not enough storage 72

LRC needs upgrading to ESHs or LRC/ESH fails 77

Failed adapters 81

Removing clustering 83

Appendix A

Communications Regulations. 87

Regulatory notices 88

Declaration of Conformity 90

Index 91

Preface

About this guide

This guide describes the following tasks and topics:

- ◆ Preparing to install an FAS900 series cluster
- ◆ How to cable a pair of FAS900 series appliances into a cluster
- ◆ How to configure the cluster when it is first brought up
- ◆ How to perform basic troubleshooting on the cluster

This guide does not cover basic system or network administration, which is covered in the *Data ONTAP System Administration Guide*.

Audience

This guide is for qualified system administrators and service personnel who are familiar with Network Appliance™ equipment and who need to

- ◆ Cable two systems into a standard or mirrored cluster system
- ◆ Convert stand-alone systems into a standard or mirrored cluster system

Terminology

This guide uses the following terms:

- ◆ *Mirrored cluster* refers to a cluster where the data of the local node contains an exact copy of the partner node's data, and vice versa.
- ◆ *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 up to 30 kilometers.
- ◆ *Node* refers to either the local or partner filer.
- ◆ *Standard cluster* refers to a cluster that is in the same building or has the local and partner nodes directly connected to each other.
- ◆ *FC HBA for Disk or FC HBA* refers to the Fibre Channel host bus adapter that connects the node to the switch or to the disks.

Command conventions

You can enter commands on the system console or from any client that has access to the system 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. 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.
type	Used to mean pressing one or more keys on the keyboard.
enter	Used to mean pressing one or more keys and then pressing the Enter 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.
Monospaced font	◆ Command and daemon names. ◆ Information displayed on the system console or other computer monitors. ◆ Contents of files.

Formatting convention	Type of information
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


All products are Class 1 laser devices. 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.

- ◆ **Switzerland only—for FAS900, GF900, R200, and C6200 systems:**
This equipment relies on fuses/circuit breakers in the building installation for overcurrent protection. Each power supply must receive power from a separately dedicated outlet with a 10A fuse/circuit breaker.
- ◆ 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 has 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


Alle Produkte sind Lasergeräte der Klasse 1. Die folgenden Sicherheitshinweise sind beim Betreiben des Geräts unbedingt zu beachten:

Vorsicht

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

- ◆ **Nur für die Schweiz - Systeme FAS900, GF900, R200 und C6200:**
Diese Geräte erfordern den Festeinbau von Sicherungen zum Überstromschutz. Jeder Netzanschluss muss mit Strom aus getrennten, speziell für diesen Zweck vorgesehenen Steckdosen versorgt werden, die jeweils mit einer 10A-Sicherung geschützt sind.
- ◆ 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 Abzwegleitungen 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 how to plan and prepare for the installation of an FAS900 series clustered system.

Topics in this chapter This chapter contains the following information:

- ◆ “[Required documentation, tools, and equipment](#)” on page 2
- ◆ “[Clustering concepts](#)” on page 5

Required documentation, tools, and equipment

Required documentation

The following table lists and briefly describes the documentation you might need to reference when preparing a new cluster, or converting two stand-alone systems into a clustered system.

Manual name	Description
<i>System Cabinet Guide</i>	This guide describes how to install NetApp equipment into a system cabinet.
<i>Site Requirements Guide</i>	This guide describes the physical requirements your site must meet in order to install NetApp equipment.
The appropriate disk shelf Guide	These guides describe how to cable a disk shelf to NetApp appliances.
<i>FAS900 Series Hardware and Service Guide</i>	This guide describes how to install an FAS900 Series appliance, connect it to a network, and bring it up for the first time.
<i>NetApp Diagnostics Guide</i>	This guide describes the diagnostics tests that can be run on the appliance.
<i>Upgrade Guide</i>	This guide describes how to upgrade appliance and disk firmware, and how to upgrade appliance software.
<i>Data Protection Guide</i>	This guide describes, among other topics, what a synchronous mirror is, how to enable it, and how to maintain it.
<i>System Administrator's Guide</i>	This guide describes how to maintain your appliance and disk drive software.
<i>Software Setup Guide</i>	This guide describes how to configure the software of a new appliance for the first time.
The appropriate Brocade® manuals	These guides describe how to configure and maintain Brocade switches.

Required tools

The following list specifies the tools you need to install the cluster:

- ◆ #1 and #2 Phillips screwdrivers
- ◆ Hand level
- ◆ Marker

Required equipment

When you receive your cluster, you should receive the equipment listed in the following table. See the System Configuration guide to confirm this equipment.

Required equipment	Standard cluster	Nonswitched MetroCluster	Fabric Attached MetroCluster
Appliance type	Two of the same type of FAS900 appliance: FAS940, FAS960, or FAS980		
Storage	Up to the maximum for mirrored clusters. See the System Configuration Guide for storage maximums.		Up to 32 TB of storage for DS14 and DS14mk2 FC disk shelves.
Cluster interconnect adapter	IB cluster adapter (The NVRAM5 functions as the cluster interconnect adapter on the FAS980.)		VI-MC cluster adapter
FC-AL adapter or FC HBA (FC HBA for Disk)	Minimum of two FC-AL adapters	Two FC HBA adapters Note _____ 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. _____	
Fibre Channel switches	N/A	N/A	Two pairs of Brocade 3800 or 3200 switches, the same type per node.
SFP modules	N/A	N/A	Two per ISL, long distance. One per switch port used, short distance.

Required equipment	Standard cluster	Nonswitched MetroCluster	Fabric Attached MetroCluster
NVRAM5 cluster adapter media converter	Only if using fiber cabling.	Only if using fiber cabling.	N/A
Cables	◆ One optical cable per port on the dual Fibre Channel adapter or One HSSDC to DB-9 node-to-shelf cable per disk loop ◆ Multiple DB-9 shelf-to-shelf cables ◆ Two optical cables with media converters for the IB cluster adapter Note_____ Clusters that have FC9 disk shelves require one MIA (media interface adapter) for each input LRC on the first disk shelves only. _____	◆ Four SC/LC FC-AL to disk shelf cables ◆ Two SC/LC IB cluster adapter-to-switch cables ◆ Four SC/LC disk shelf-to-switch cables	◆ 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 cables

Clustering concepts

About cluster concepts

There are two types of clusters you can install in your site. The following table lists the cluster types and gives a brief description.

Cluster type	Description
Standard cluster	Appliances are physically close and directly connected to each other. There is no mirroring of the data in this type of cluster.
Mirrored cluster and nonswitched MetroCluster	Mirrored clusters have appliances separated by up to 30 meters. Nonswitched MetroClusters have appliances and disk shelves separated by more than 30 meters and up to 500 meters away from each other. Mirrored and nonswitched MetroClusters require different cables from each other and standard clusters. Data on the local node is mirrored (synchronized copy) on the partner node, and vice versa. The local and partner nodes are directly connected to each other.
Fabric Attached MetroCluster	Appliances and disk shelves are cabled directly to switches and not to each other. Unlike a traditional cluster, each node and set of associated disk shelves are directly connected to a set of switches, and not to each other. Disk ownership is determined by where the FC HBA (FC HBA for Disk) is cabled to the switch and where the disks are cabled. Data on the cluster is mirrored. This cluster type only supports the DS-14 series disk shelves.

For detailed information

For details about these cluster types, see the following sections:

- ◆ [“Standard clusters”](#) on page 6
- ◆ [“Mirrored clusters and MetroClusters”](#) on page 7
- ◆ [“Fabric Attached MetroCluster”](#) on page 9

Standard clusters

What standard clusters contain

A standard cluster contains two sets of Fibre Channel disk shelves, a local node, and a partner node. The nodes are connected to each other through an IB cluster adapter or NVRAM5 adapter (FAS980), which allows a takeover node to serve data to the disks of a failed partner node. Each node continually monitors its partner, mirroring the data for each other's NVRAM. The two nodes in a cluster can be separated from each other by as much as 500 meters, requiring them to be connected with the LC-LC 50/125 cable and connected to their partner disk shelves with SC-SC 50/125 cables.

Advantages of standard clusters

The advantages of having a standard cluster are as follows:

- ◆ Data continues to be processed if one of the nodes fails and cannot reboot.
- ◆ Data continues to be processed if an adapter on a node fails.
- ◆ System maintenance can be performed without interrupting service.

In any of these situations, the takeover node continues to serve data until the problem is corrected.

How Data ONTAP works with standard clusters

In a clustered system, Data ONTAP functions so that each node monitors the functioning of its partner through a heartbeat signal sent between the nodes. Data from the NVRAM of one node is mirrored by its partner and can take over the partner's disks, if the partner should fail. Also, the nodes synchronize each other's time. Any spare disks belonging to one node are not shared by the partner node. See the *Data ONTAP System Administrator's Guide* for more information.

Standard cluster setup requirements and restrictions

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

- ◆ IB cluster adapters and cables must be installed.
- ◆ Nodes must be attached to the same network and the NICs must be configured correctly.
- ◆ Cluster software must be licensed and enabled on both nodes.
- ◆ The same system software must be licensed and enabled on both nodes.

Mirrored clusters and MetroClusters

About mirroring

Mirrored clusters and MetroClusters contain two complete copies of the specified 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 data

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

- ◆ Mirrored data on your cluster survives a double disk error in one of your disk shelves.
- ◆ Mirrored data on your cluster survives if one part of the cluster loses connectivity to your disks.
- ◆ Mirrored data on your cluster survives if your cluster loses connectivity to remote storage.
- ◆ Mirrored 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 MetroClusters

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.

Setup requirements and restrictions

The following list specifies the requirements and restrictions you must 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 node slots 2 through 5, if available.
- ◆ Disk pool 1 volumes must be cabled to FC-AL or dual Fibre Channel adapters in node slots 8 through 11 (not applicable in a switched MetroCluster).
- ◆ The two halves of the mirrored disk should have no common components.
- ◆ Local SyncMirror must be licensed and enabled on the local and partner nodes (not applicable in a switched MetroCluster).
- ◆ MetroClusters must have cluster_remote licensed and enabled on both nodes.
- ◆ There must be sufficient spares in each pool to account for a disk failure.
- ◆ Cluster nodes that are separated by more than 30 meters, and less than 500 meters, must use an LC-LC 50/125 cable to connect their IB cluster adapters or NVRAM5 adapters.

Fabric Attached MetroCluster

What a Fabric Attached MetroCluster is

A Fabric Attached MetroCluster is a clustered system where the local and partner nodes are separated by more than 500 meters and up to 30 kilometers. Each node and associated disk shelves are connected directly to a pair of Fibre Channel switches, then across some distance, to the partner node and disk shelves. Disk shelf and node ownership are determined by where they are plugged into the switch and combined with the partner switch, create a fabric.

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 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.



How disk ownership is determined

Disk ownership by a node is determined by where the FC HBA for Disk connects into the switch and by where the disks connect to the switch banks. FC HBAs connect to one switch 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.

Advantages of Fabric Attached MetroClusters

The advantages of Fabric Attached MetroClusters are the same as those of standard or mirrored clusters. See “[Standard clusters](#)” on page 6 and “[Mirrored clusters and MetroClusters](#)” on page 7 for information about the advantages of each cluster type.

In addition, Fabric Attached MetroClusters have the following advantages:

- ◆ Expanded distance between halves of the cluster beyond 500 meters.
- ◆ Disk shelves and nodes are not connected directly to each other, but are connected to a fabric, with multiple data routes.

How Data ONTAP works with these MetroClusters

Data ONTAP functions the same as it does in a nonswitched MetroCluster.

Setup requirements and restrictions

Cabling a Fabric Attached MetroCluster is the same as for standard or mirrored clusters, with the following exceptions:

- ◆ Each clustered pair of nodes must use two switches, per cluster side, to communicate with the clients.
- ◆ Switches must be either Brocade 3200 or 3800 and must be the same type on each side of the MetroCluster.
- ◆ Each node must have the HBA kit installed and running.
- ◆ Each node must have the VI-MC cluster adapter installed; slot position is dependent on the type of filer.
- ◆ syncmirror_local and cluster_remote must be licensed and enabled on the local and partner nodes.
- ◆ Mixed switch types can be used, but must be the same type on each side of the cluster.

About this chapter

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

See Chapter 3, “[Cabling a MetroCluster](#),” on page 29 for information about cabling a MetroCluster.

Topics in this chapter

The following topics are covered in this chapter:

- ◆ “[Cluster installation types](#)” on page 14
- ◆ “[Cabling the cluster](#)” on page 15

Cluster installation types

Types of cluster installations

There are two types of cluster installations:

- ◆ “[Clusters in a system cabinet](#)” on page 14
- ◆ “[Clusters in an equipment rack](#)” on page 14

Clusters in a system cabinet

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

Single system cabinet cluster: If the ordered cluster has six or fewer disk shelves, it arrives in a single system cabinet. This system cabinet has both the Channel A and Channel B disk shelves cabled, and also has the cluster adapters cabled.

Multiple system cabinet cluster: If the ordered cluster has more than six disk shelves, the cluster arrives in two or more system cabinets. You must complete the cabling by cabling 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 the nodes together by cabling the cluster adapters. The system cabinets might also need to be connected to each other. See the *System Cabinet Guide* for information about connecting your system cabinets together for clustering.

Clusters in an equipment rack

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 *FAS900 Series Hardware and Service Guide*.

Cabling the cluster

Type of FC-AL adapter used in the cluster

The dual-port Fibre Channel (dual-port FC-AL) adapters have two ports per adapter. These ports support both the Channel A and Channel B disk shelf loops, with the top port supporting the Channel A loop and the bottom port supporting the Channel B loop. The dual-port Fibre Channel adapter allows the FAS900 series appliance to support more storage than a single-port FC-AL adapter.

Minimum FC-AL adapters needed for mirrored clusters

If the cluster is mirrored, each node must have a minimum of two dual-port FC-AL adapters per node. The number of dual-port FC-AL adapters depends on the number of disk shelves you have in each loop.

Recommended dual-port FC-AL adapter locations

The following table lists the most common recommended locations for the dual-port FC-AL adapters in a single node of the cluster. Check the System Configuration Guide on the NOW Web page at <http://now.netapp.com> to confirm the slot assignment information.

Cluster type	Node slot number on back of the node
Standard cluster—one dual-port FC-AL adapter	Slot 2
Mirrored cluster—one dual-port FC-AL adapter per node	N/A
Standard cluster—Two dual-port FC-AL adapters per node	Slot 2, Slot 3
Mirrored cluster—Two dual-port FC-AL adapters per node	Slot 8 (pool 1), slot 5 (pool 0)
Standard cluster—Four dual-port FC-AL adapters per node	Slot 2, slot 3, slot 4, slot 8
Mirrored cluster—Four dual-port FC-AL adapters per node	Slot 8 (pool 1), slot 5 (pool 0), slot 10 (pool 1), slot 3 (pool 0)

**For detailed
information**

For detailed information about the cabling process, see the following topics:

- ◆ [“Step 1: Preparing your equipment”](#) on page 17
- ◆ [“Step 2: Cabling Channel A”](#) on page 19
- ◆ [“Step 3: Cabling the cluster adapters”](#) on page 23
- ◆ [“Step 4: Cabling Channel B”](#) on page 24

Step 1: Preparing your equipment

Where to install a cluster

A cluster can be installed in a system cabinet, multiple system cabinets, or in equipment racks. You must install the equipment as described in the following sections.

- ◆ “[Preparing the system cabinet](#)” on page 17
- ◆ “[Preparing the equipment racks](#)” on page 18

Preparing the system cabinet

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

Step	Action	
1	Install the system cabinet or cabinets, as described in the <i>42U System Cabinet Guide</i> .	
2	If you have a cluster in...	Then...
	A single system cabinet	1. Connect the nodes to the network. 2. Connect the system cabinet to an appropriate power source. 3. Apply power to the system cabinet and configure the cluster, as described in Chapter 4, “Configuring the Cluster,” on page 47.
	Multiple system cabinets	1. Remove the front and rear doors and any side panels that need to be removed to connect the system cabinets together. 2. Connect the system cabinets together, as described in the <i>42U System Cabinet Guide</i>. 3. Connect the nodes to the network. 4. Go to “Step 2: Cabling Channel A” on page 19.

Preparing the equipment racks

To prepare your nodes for clustering in equipment racks, complete the following steps.

Step	Action
1	Install the nodes in the equipment racks, as described in the <i>FAS900 Series Hardware and Service Guide</i> .
2	Install the disk shelves in the equipment rack, as described in the appropriate Fibre Channel disk shelf guide.
3	Label the dual-port FC-AL interfaces, where appropriate.
4	Connect the nodes to the network, as described in the <i>FAS900 Series Hardware and Service Guide</i> .
5	Go to “ Step 2: Cabling Channel A ” on page 19.

Step 2: Cabling Channel A


Channel A loop cabling differences

The procedure for cabling the Channel A loops for a standard cluster and mirrored cluster are different. These differences are described in the following sections:

- ◆ “Cabling the Channel A loops in a standard cluster” on page 19
- ◆ “Cabling the Channel A loops in a mirrored cluster” on page 21

Cabling the Channel A loops in a standard cluster


To cable the Channel A loops in a standard cluster, complete the following steps, using the diagram for reference.



Step	Action
1	Cable the dual-port FC-AL, 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. Label this cable with a solid-colored label.
3	Repeat Step 2, connecting output to input, for the remaining disk shelves in the loop.
4	Set the terminate switches on the LRC/ESH modules to Off on all but the last disk shelf in the loop, then set the LRC/ESH terminate switches on the last disk shelf to On.
5	Repeat Steps 1 through 5 for the partner node.
6	Go to “ Step 3: Cabling the cluster adapters ” on page 23.

Cabling the Channel A loops in a mirrored cluster

To cable the Channel A loops in a mirrored cluster, complete the following steps, using the diagram for reference.




Step	Action
1	Cable the dual-port FC-AL in slot 8, Port 0 to the disk shelf 1 LRC/ESH Input port of disk pool 1, then cable the dual-port FC-AL in Slot 3, Port 0 to the disk shelf 1 LRC/ESH Input port of disk pool 0.
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 this cable with a solid-colored label.
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	Repeat Steps 1 through 4 for the partner node.
6	Go to “ Step 3: Cabling the cluster adapters ” on page 23.

Step 3: Cabling the cluster adapters

Cabling the cluster adapters

To cable the cluster adapters, complete the following steps, using the diagram for reference. The FAS940 and FAS960 use the IB cluster adapter. For the FAS980, your NVRAM5 adapter in slot 11 is used as the cluster interconnect adapter.



Step	Action
1	Plug one end of the optical cable into one of the local node’s cluster adapter ports, then plug the other end into the corresponding partner node’s cluster adapter port.
2	Repeat Step 1 for the remaining port on the cluster adapter.
3	Gently pull on the cable ends to ensure that they are properly seated.
4	Go to “ Step 4: Cabling Channel B ” on page 24.

Step 4: Cabling Channel B


Channel B loop cabling differences

The procedure for cabling the Channel B loops for a standard cluster and mirrored cluster are different. These differences are described in the following sections.

- ◆ “Cabling the Channel B loops in a standard cluster” on page 24
- ◆ “Cabling the Channel B loops in a mirrored cluster” on page 26

Cabling the Channel B loops in a standard cluster


To cable the Channel B loops in a standard cluster, complete the following steps, using the diagram for reference.



Step	Action
1	Cable the dual-port FC-AL adapter in slot 2, 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 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 Chapter 4, " Configuring the Cluster ," on page 47.

Cabling the Channel B loops in a mirrored cluster

To cable the Channel B loops in a mirrored cluster, complete the following steps, using the diagram for reference.



Step	Action
1	Cable the dual-port FC-AL adapter in slot 8, Port 1 of the local node, to the partner node disk shelf 1 LRC/ESH Input port of disk pool 1, then cable the dual-port FC-AL adapter in slot 5, Port 1 of the local node to the partner node disk shelf 1 LRC/ESH Input port of 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. Label this cable with a striped label.

Step	Action
3	Repeat Step 2 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 Chapter 4, “ Configuring the Cluster ,” on page 47.

About this chapter

This chapter contains instructions for cabling a new Fabric Attached or nonswitched MetroCluster. This chapter also describes how to convert a cluster into a Fabric Attached MetroCluster.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Cabling a new Fabric Attached MetroCluster”](#) on page 30
- ◆ [“Cabling a nonswitched MetroCluster”](#) on page 44
- ◆ [“Converting an existing cluster to a Fabric Attached MetroCluster”](#) on page 45

Cabling a new Fabric Attached MetroCluster

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 should also be familiar with the characteristics of Fabric Attached MetroClusters, as described in [“Fabric Attached MetroCluster”](#) on page 9. 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 details about setting up and cabling your Fabric Attached MetroCluster, see the following sections:

- ◆ [“Step 1: Planning the Fabric Attached MetroCluster”](#) on page 31
- ◆ [“Step 2: Configuring the local switches”](#) on page 32
- ◆ [“Step 3: Cabling the local node”](#) on page 35
- ◆ [“Step 4: Cabling the remote switches and node”](#) on page 38
- ◆ [“Step 5: Configuring the nodes”](#) on page 42

Step 1: Planning the Fabric Attached MetroCluster

Fabric Attached MetroCluster worksheet

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

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 is cabled to Switch 1 Port__

Switch 1 name: _____

Switch 1 IP address:_____

FC HBA adapter 1 port 1 is cabled to Switch 1, Port__

FC HBA adapter 1 port 2 is 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 adapter 1 port 1 is cabled to Switch 3, Port__

FC HBA adapter 1 port 2 is cabled to Switch 4, Port__

Pool 0 disk shelves:

Disk 1 connected to Port ____

Disk 2 connected to Port ____

Shelf 1 ESH A IN connected to ____

Shelf 1 ESH B IN connected to ____

Cluster adapter Port 2 is cabled to Switch 2 Port __

Switch 2 name: _____

Switch 2 IP address:_____

FC HBA adapter 2, port 1 is cabled to Switch 1, Port _

FC HBA adapter 2, port 2 is 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 adapter 2 port 1 is cabled to Switch 3, Port _

FC HBA adapter 2, port 2 is cabled to Switch 4, Port __

Pool 1 disk shelves:

Disk 3 connected to Port ____

Disk 4 connected to Port ____

Shelf 2 ESH A IN connected to ____

Shelf 2 ESH B IN connected to ____

Cabling a new Fabric Attached MetroCluster

Step 2: Configuring the local switches

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	Disable the configuration by entering the following command: cfgdisable
7	Save the configuration by entering the following command: cfgsave
8	Ensure that the switch is configured with default settings by entering the following command: configdefault
9	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
10	Exit the configuration utility by entering the following command: ctrl-d
11	Reboot the switch so that the default settings take effect by entering the following command: fastboot
12	Log in to the switch and disable the switch by entering the following command: switchdisable

Step	Action	
13	Set the switch port speeds to half duplex by entering the following command for each switch port into which a disk is plugged: portcfglport <i>port#</i> 0,0,1 Note _____ You only need to do this command on the ports to which disks are attached.	
14	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-<i>port#</i>, "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-<i>port#</i>, "L1"
15	Enable the switch by entering the following command: switchenable	
16	Set the switch name by entering the following command: switchname <i>Metro1_switch1</i> <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.	
17	Check the switch settings by entering the following command: configshow	
18	Repeat Steps 1 through 16 on the second switch on the local node.	
19	Go to “Step 3: Cabling the local node” on page 35	


Step 3: Cabling the local node

Cabling the local node and disks

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 the 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. 4. Set the terminate switch on the last disk shelf in the loop to Off. 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 the port in Pool 1 in Bank 1 on Switch 2. 2. Connect the Input port of LRC/ESH B to the 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. 4. Set the terminate switch on the last disk shelf in the loop to Off. 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 have dual paths 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	Boot the node.																																																
9	Check to make sure that the node can see all disk shelves by entering the following command on the console: vol status -r Result: All disks are displayed.																																																
10	Go to the remote site and follow the instructions in “ Step 4: Cabling the remote switches and node ” on page 38.																																																

Step 4: Cabling the remote switches and node

About remote switch cabling

The remote switches need to be cabled so that they can work with the local switches.


Cabling the remote switches

To cable 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 32.
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 3 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 3 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 the port in Pool 0 in Bank 1 on Switch 3. 2. Connect the Input port of LRC/ESH B to the 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. 4. Set the terminate switch on the last disk shelf in the loop to Off. 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 the port in Pool 1 in Bank 2 on Switch 3. 2. Connect the Input port of LRC/ESH B to the 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. 4. Set the terminate switch on the last disk shelf in the loop to Off. 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 have dual paths 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	Boot the node.																																																
9	Check to make sure that the node can see all disk shelves by entering the following command on the console: vol status -r																																																
10	Go to “ Step 5: Configuring the nodes ” on page 42.																																																

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 run 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: <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 <code>license option=code</code>. _____
3	Reboot the local node by entering the following command: <pre>reboot</pre>
4	Repeat Steps 1 through 3 for the remote node.
5	License cluster_remote on both nodes.
6	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 <code>vol mirror vol-name</code> command. _____

Step	Action
7	Enable clustering by entering the following command on the console: <pre>cf enable</pre>
8	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> Result: 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.
9	Configure both nodes, as described in Chapter 4, “Configuring the Cluster,” on page 47. Note_____ You do not have to reenable clustering licenses. _____

Cabling a nonswitched MetroCluster

Cabling a nonswitched MetroCluster

To cable a nonswitched MetroCluster, complete the following steps.

Step	Action
1	Prepare your sites, as described in “Step 1: Preparing your equipment” on page 17.
2	Cable the Channel A loops of the MetroCluster, as described in “Cabling the Channel A loops in a mirrored cluster” on page 21.
3	Cable the nodes together, as described in “Cabling the cluster adapters” on page 23.
4	Cable the Channel B loops of the MetroCluster, as described in “Cabling the Channel B loops in a mirrored cluster” on page 26.
5	Configure the cluster, as described in Chapter 4, “Configuring the Cluster,” on page 47. Make sure that you enable the syncmirror_local and cluster_remote licenses on both nodes.

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.

Step	Action
1	Update Data ONTAP, filer firmware, and disk firmware, as described in “Out-of-date software and firmware” on page 65, making 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 by replacing components, as described in “Failed adapters” on page 81. 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 31.
4	Set up and configure the local switches, as described in “Step 2: Configuring the local switches” on page 32.
5	Cable the local node, as described in “Step 3: Cabling the local node” on page 35.

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 35. 2. Go to Step 7.
	Nonswitched MetroCluster	Go to Step 7.
7	Transport the partner node, disk shelves, and 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 38.	
9	Configure the cluster, as described in “ Step 5: Configuring the nodes ” on page 42.	

About this chapter

This chapter describes how to bring up a new clustered system for the first time. It describes how to enable licenses, set cluster options, configure networking for the cluster, and test the configuration.

Topics in this chapter

This chapter addresses the following topic:

- ◆ [“Step 1: Bringing up the cluster for the first time”](#) on page 48
- ◆ [“Step 2: Enabling licenses”](#) on page 49
- ◆ [“Step 3: Setting options and parameters”](#) on page 51
- ◆ [“Step 4: Configuring networking interfaces”](#) on page 54
- ◆ [“Step 5: Checking and testing cluster configuration”](#) on page 57

Step 1: Bringing up the cluster for the first time

What to do to bring up the cluster

After cabling the nodes and storage together, you must connect the cluster to the network, apply power to the system, and configure the nodes of the cluster.

Bringing up the new cluster

To bring up the cluster for the first time, complete the following steps.

Step	Action
1	Connect both nodes to the network, as described in the <i>Hardware and Service Guide</i> .
2	Turn on the power to the disk shelves.
3	Turn on the power to both nodes and wait for the system prompt.
4	Run the diagnostics program by entering the following command: sys-diag Result: The diagnostic program loads.
5	Select the diagnostics tests that check all dual-port FC-AL adapters and the cluster adapter. For a description of these tests, see the Diagnostics Guide on the NOW Web site at http://now.netapp.com .
6	Boot both nodes and observe the console output, by entering the following command from each node's console: boot
7	Check the status of the cluster by entering the following command on each node's console: cf status Result: The following message is displayed on the console, indicating that clustering is not activated: cluster monitor not initialized
8	Go to “ Step 2: Enabling licenses ” on page 49.

Step 2: Enabling licenses

- About licenses on the cluster

License activation enables the cluster to run software that is installed on the cluster. For clustering, you must enable the cluster license at bring up.
- Licenses for mirrored clusters

If you are setting up a mirrored cluster, you must also enable the Local SyncMirror license for synchronous mirroring. If you are setting up a cluster with nodes in separate locations, you must also enable the Cluster Remote license.
- Other licenses

There are other licenses that you can set, depending on what licensed software is installed on the cluster. The process for enabling each license is the same. See the *System Administrator's Guide* for additional licenses you can set.

Enabling licenses on the cluster

To enable a license on the cluster, complete the following steps.

Step	Action
1	Enter the cluster license and code on both node consoles, as shown in the following example. <code>license cluster = C1234586</code> Note If you are setting up a mirrored cluster, you should enable Local SyncMirror during this step. Note If you are using a version of Data ONTAP earlier than 6.4, the command is <code>license option=code</code>.
2	Reboot both nodes using the following command: <code>reboot</code>

Step	Action
3	Enable clustering on each node by entering the following command on the local node console: cf enable
4	Check to see that clustering is enabled by entering the following command on each node console: cf status Cluster enabled, filer 2 is up.
5	Repeat Step 1 for other licenses, using the license type and code for each licensed product installed on the cluster.
6	Go to “ Step 3: Setting options and parameters ” on page 51.

Step 3: Setting options and parameters

What the options do Options help you maintain various functions of your node, such as security, file access, and network communication. During takeover, the value of an option might be changed by the node doing the takeover. This can cause unexpected behavior during a takeover. To avoid unexpected behavior, specific option values must be the same on both the local and partner node.

Each option performs a specific task or has a specific function. See the *System Administrator’s Guide* for a description of available options and the function each performs.

Options that must be the same The options listed in the following table must be the same for a takeover to be successfully completed.

Option name	Option name
cifs.show_snapshot	nfs.v2.df.3gb.lim
dns.domainname	nis.domainname
dns.enable	nis.enable
http.admin.enable	pcnfsd.enable
httpd.timeout	raid.reconstruct_speed
httpd.timewait.enable	raid.timeout
ip.path_mtu_discovery.enable	wafl.maxdirsize
nfs.per_client_stats.enable	wafl.root_only_chown
nfs.v2.df.2gb.lim	vol.vopy.throttle

What the parameters do The parameters help keep the individual synchronized and help maintain routing of processes throughout the system.

Parameters that must be the same

The parameters listed in the following table must be the same so that takeover is smooth and data is transferred between the nodes correctly.

Parameter	Parameter
date	date, rdate
NDMP (on or off)	ndmp (on or off)
route table published	route
routed enabled	routed (on or off)
Time zone	timezone

Manually setting options

To manually set options, complete the following steps.

Step	Action
1	View and note the values of the options on the local and partner nodes, using the following command on each console: options Result: The current option settings for the node are displayed on the console. Output similar to the following is displayed: <pre>autosupport.doit DONT autosupport.enable off . . .</pre>
2	Verify that the options that have comments in parentheses are set to the same value for both nodes. The comments are as follows: ◆ Value might be overwritten in takeover ◆ Same value required in local+partner ◆ Same value in local+partner recommended
3	Correct any mismatched options using the following command: options option_name option_value

Manually setting parameters

To manually set parameters, complete the following steps.

Step	Action
1	View the individual parameter settings on both the local and partner nodes, similar to the following example: routed status Result: The default gateway list status is displayed.
2	Verify that the parameters on both nodes are set to the same value.
3	Correct any mismatched values.
4	Go to “ Step 4: Configuring networking interfaces ” on page 54.

Setting options and parameters with 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 both 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 nodes.
4	Examine the script output and correct any errors or differences.
5	Go to “ Step 4: Configuring networking interfaces ” on page 54.

Step 4: Configuring networking interfaces

What the networking interfaces do

When a node in a cluster fails, the surviving node must be able to assume the identity of the failed node on the network. Networking interfaces allow individual nodes in the cluster to maintain communication with the network if the partner fails.

Network connection requirements

Both nodes in the cluster must meet the following network interface connection characteristics:

- ◆ Both nodes must have interfaces that access the same collection of networks and subnetworks.
- ◆ Both nodes must have interfaces that can take over the IP addresses of each other.
- ◆ Both nodes must have interfaces that can take over the MAC (media access control) addresses of each other.

Types of takeover configurations

The following table lists the three types of interface configurations that can be enabled in a cluster.

Interface type	Description
Shared	This type of interface supports both the local and partner nodes. It contains both the local node and partner node's IP addresses. During takeover, it supports the identity of both nodes.
Dedicated	This type of interface only supports the node in which it is installed. It contains the local node's IP address only and does not participate in network communication beyond local node support during takeover.
Standby	This type of interface is on the local node, but only contains the IP address of the partner node. During takeover, this interface supports a virtual filer identity. See the <i>System Administrator's Guide</i> for more information about virtual filers.

Information needed for configuration

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

Configuring networking for the first time

To configure the networking interfaces on a cluster for the first time, complete the following steps.

Step	Action
1	Enter the following command at the local node console: setup
2	Answer the questions displayed on the screen. Please enter IP address for Network Interface ns0 [<i>n.n.n.n</i>): 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: [<i>m.m.m.m</i>):
3	Enter the following command to halt the node: halt
4	Enter the following command to reboot the node: reboot
5	Repeat Steps 1 through 4 for the partner node.
6	Go to “ Step 5: Checking and testing cluster configuration ” on page 57.

Configuring interfaces

To configure an interface as a dedicated, standby, or shared interface, complete the following steps.

Step	Action	
1	Check the current settings of the local node by entering the following command on the console: 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 fffffff0 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	If you are setting the interface as a...	Then...
	Dedicated interface	1. Enter the following command at the prompt, or directly into the /etc/rc file of the local node: ifconfig e0 local_node_ip_address 2. Repeat Step 1 for the partner node.
	Shared interface	1. Enter the following command at the prompt, or directly into the /etc/rc file of the local node: ifconfig NIC_name ip_address partner \ ip_address Example: ifconfig e0 199.9.204.255 \ partner 199.9.204.254 2. Enter the following command at the prompt, or directly into the /etc/rc file of the partner node: ifconfig NIC_name ip_address partner \ ip_address Note_____ You do not have to enter the MAC address. _____
	Standby interface	1. Enter the following command at the prompt, or directly into the /etc/rc file of the local node: ifconfig e1 local_node_ip_address 2. Repeat Step 1 for the partner node.

Step 5: Checking and testing cluster configuration

Reason to check and test the cluster

After configuring networking and setting licenses and options, you should check your configuration with the cluster configuration utility. This utility runs under UNIX or as a CGI script from a UNIX Web server and test for the following configuration problems:

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

Checking configuration

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

Step	Action
1	Download the cluster configuration checker (cf-config-check.cgi) file from the NOW Web site at http://now.netapp.com/NOW/tools , then unzip the file.
2	Collect the IP addresses and login information for the local and partner nodes.
3	Run the script, as described in the online instructions on the script Web page.
4	Examine the script output and correct any errors or differences indicated in system settings, such as mismatched licenses.
5	Go to “ Testing takeover and giveback ” on page 57.

Testing takeover and giveback

To test takeover and giveback, complete the following steps.

Step	Action
1	Check the cabling on the cluster interconnect cables to make sure that they are secure.

Step	Action	
2	Verify that you can create and retrieve files on both nodes for each licensed protocol.	
3	From the local node, enter the following command and observe the output: cf enable Result: Takeover of the partner node is enabled on the local node.	
4	Enter the following command from the local node console: cf takeover Result: The local node takes over the partner node and gives the following output: takeover completed	
5	Make sure that the local node can access both the local and partner disks.	
6	Give back the partner node by entering the following command: cf giveback Result: The local node releases the partner node, which reboots and resumes normal operation. The following message is displayed on the console when the process completes: giveback completed	
7	If takeover and giveback...	Then...
	Is completed successfully	Repeat Steps 2 through 6 on the partner node.
	Fails	Go to Chapter 5, “ Basic Troubleshooting ,” on page 59.
8	Complete setup as described in the <i>Software Setup Guide</i> .	

About this chapter

This chapter describes some common problems you might encounter configuring or operating your cluster and provides steps you can take to correct these problems. This chapter also describes how to address full storage in the cluster and how to convert the cluster into two, stand-alone nodes.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Takeover or giveback failures”](#) on page 60
- ◆ [“Hot-swap disk loop error messages on Fabric Attached MetroClusters”](#) on page 62
- ◆ [“Recovering from a disaster”](#) on page 63
- ◆ [“Out-of-date software and firmware”](#) on page 65
- ◆ [“Not enough storage”](#) on page 72
- ◆ [“LRC needs upgrading to ESHs or LRC/ESH fails”](#) on page 77
- ◆ [“Failed adapters”](#) on page 81
- ◆ [“Removing clustering”](#) on page 83

Takeover or giveback failures

Correcting takeover or giveback failure

The following table describes tests you can run to determine why takeover or giveback failed and also provides information to correct the problem.

Note—
You must reboot the cluster after completing each test.

To correct a takeover or giveback failure, complete the following steps.

Step	Action	
1	Check communication between the local and partner nodes by entering the following command and observing the messages: cf status Result: If there is a problem with cabling or the adapter, the appropriate message is displayed. Recable the cluster adapter or replace the adapter, as needed.	
2	If the error message indicates...	Then...
	A cluster adapter error	Check the cluster adapter cabling. Make sure that the cabling is correct and properly seated at both ends of the cable.
	That the NVRAM5 adapter card is in the wrong slot number	Check the NVRAM slot number. Move it to the correct slot if it is in the wrong one.
	A Channel B cabling error	Check the cabling of the Channel B disk shelf loops and reseal and tighten any loose cables.
	A networking error	Check for network connectivity. See the <i>System Administrator's Guide</i> for more information.

Step	Action
3	Run the cluster configuration checker script, as described in “ Setting options and parameters with the cluster configuration checker ” on page 53. Correct any errors or differences displayed in the output.
4	Reboot the cluster and rerun the takeover and giveback tests.
5	If you still do not have clustering enabled, call Customer Support.

Hot-swap disk loop error messages on Fabric Attached MetroClusters

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]: Met1_sw1:11.34 (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 Met1_sw1:11.34: 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 Met1_sw1:11.39: request successful after retry #1: cdb Nx28:N162416N:N1NN.	Contact was reestablished with the disk loop.
[node1: ses_admin:info]: initiating rescan on channel Met1_sw1:11	The system is scanning the disk loops to determine what has been added or removed.
[node1: raid.disk.inserted:info]: Disk Met1_sw1:11.72 (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 has found that it is up to date.

Recovering from a disaster

Assumptions about your cluster

The following section assumes that you created a mirrored cluster with the following configuration:

- ◆ The local and partner nodes are in separate physical locations.
- ◆ Clustering and Local SyncMirror are licensed on both nodes.
- ◆ Remote SyncMirror is licensed and running on both nodes, prior to the disaster.
- ◆ The root volume is mirrored with one plex at each site and data volumes are mirrored with one plex at each site.

Types of disasters

You might encounter the following types of disasters:

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

Recovering from a disaster

To recover from a disaster, complete the following steps.

Note—
You need the *Data Protection Guide* to help you complete this procedure.

Step	Action
1	Isolate the partner node (node at disaster site) by powering it off. If you cannot physically turn off that node, isolate it by manually isolating (fencing) it, as described in the <i>Data Protection Guide</i> .

Step	Action	
2	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 using the following command: sysconfig -v 3. Confirm connection between the local and partner node, using the following command: cf monitor 4. Complete the disaster recovery procedures, as described in the <i>Data Protection Guide</i>.
	The remote node and storage are physically damaged or failed	1. Force the takeover of the partner node by using the following command: cf forcetakeover -d 2. Replace equipment and cabling, as needed. Contact Customer Support for assistance. 3. Complete the disaster recovery procedures, as described in the <i>Data Protection Guide</i>.

Out-of-date software and firmware

Updating outdated software and firmware

Updating your software can be a four-part process:

- ◆ Update Data ONTAP
- ◆ Update only the node firmware
- ◆ Update only disk firmware
- ◆ Update only disk shelf firmware

Note

If you are updating Data ONTAP to a minor release (a dot release), node and disk shelf firmware are automatically updated. If you are updating to a major release of Data ONTAP, you must manually update node and disk shelf firmware.

For detailed information

For details about updating Data ONTAP software, filer firmware, and disk shelf firmware, see the following sections.

- ◆ [“Updating Data ONTAP with node and disk firmware”](#) on page 66
- ◆ [“Updating node firmware only”](#) on page 68
- ◆ [“Updating disk firmware on mirrored clusters and nonswitched MetroClusters”](#) on page 69
- ◆ [“Updating disk shelf firmware only”](#) on page 71

Updating Data ONTAP

Updating Data ONTAP with node and disk firmware

To update Data ONTAP, node firmware, disk firmware, and shelf firmware on each node, complete the following steps.

Caution

If you are converting a cluster to a Fabric Attached MetroCluster, do not recable the system before updating Data ONTAP and firmware.

Step	Action
1	Back up the data on both nodes. See the <i>System Administrator's Guide</i> for backup instructions.
2	Verify that both nodes have network connections and that <code>httpd.admin.enable</code> is set to On for both nodes.
3	Download and untar the most current version of Data ONTAP to both nodes, from http://now.netapp.com/NOW/cgi-bin/software/ on the NOW Web site.
4	Download and run the Data ONTAP Upgrade Utility from http://now.netapp.com/NOW/download/tools/upgrade_ontap on the NOW Web site. Follow the directions to run this script.

Updating node, disk, or shelf firmware only

Topics in this section

This section gives you information about updating node, disk, and disk shelf firmware for all cluster types.

Determining the firmware level on the nodes

Before updating firmware on the cluster nodes, you must determine the current version running on the nodes. Both nodes must run the same version of firmware to avoid potential conflicts during takeover.

To determine the firmware revision numbers, complete the following step.

Step	Action	
1	If the cluster is...	Then...
	Running	Enter the following command from each node console and note the version: sysconfig -v
	Rebooting	Note the firmware revisions during boot.
	At the ok prompt	Enter the following command from the node console and note the version: .version

Where to get firmware

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.

Firmware type	Version
Filer	See the NOW site at http://now.netapp.com/NOW/tools/releasetable.shtml for current filer firmware.

Firmware type	Version
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.

Updating node firmware only

To update node firmware, complete the following steps.

Note

This procedure will not work if you are migrating to version 6.5 from an earlier version. See the *Data ONTAP Upgrade Guide* for more information.

Step	Action
1	Disable the cluster by entering the following command on the local node: cf disable
2	Enter the following command from the local node console to copy the executables to the boot blocks of the node's hard disk: download Result: The node firmware is downloaded to the boot disk. Note This process can take several minutes.
3	Shut down the local node by entering the following command: halt
4	Turn off, then turn on the local node.
5	Enter the following command after the node boots: update-flash Result: The system prompts you to continue, then the firmware is copied to the CompactFlash™ card in the node.

Step	Action
6	Reboot the system by entering the following command: bye
7	Repeat Steps 2 through 6 for the partner node.
8	Reenable clustering by entering the following command from the local node's console: cf enable
9	Check the configuration, as described in “ Step 5: Checking and testing cluster configuration ” on page 57.

Updating disk firmware on mirrored clusters and nonswitched MetroClusters

To update disk firmware on a standard, mirrored or nonswitched MetroCluster, 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	Download the most current version disk firmware from http://now.netapp.com/NOW/download/tools/diskfw/ from the NOW Web site. Place the file in /etc/disk_fw on both nodes.
2	Make sure clustering is enabled by entering the following command: cf status Result: If clustering is enabled, the following message is displayed: Clustering enabled Note_____ If clustering is disabled, enable it with the cf enable command.

Step	Action
3	Install the new disk firmware on the local node by entering the following command: <pre>disk_fw_update</pre> Caution————— You must power cycle the disk shelves after the firmware is updated.
4	After the firmware update is complete on the local node, repeat Step 3 on the partner node.

Updating disk firmware on Fabric Attached MetroClusters

To update disk firmware on a Fabric Attached MetroCluster, complete the following steps.

Note—————
This procedure is for spare disks only.

Step	Action	
1	Download the most current version of disk firmware from http://now.netapp.com/NOW/download/tools/diskfw/ .	
2	Place the file in /etc/disk_fw on both nodes.	
3	Install the new disk firmware on the local node by entering the following command: <pre>disk_fw_update</pre>	
4	If updating firmware on...	Then...
5	A few spare disks only	- Partially remove the disk from the disk shelf. - Wait 10 seconds and reinsert the disk into the shelf.
6	An entire disk shelf	Power cycle the disk shelf.

Updating disk shelf firmware only

To update disk shelf firmware, complete the following steps.

Step	Action
1	Download the most current version of disk shelf firmware from http://now.netapp.com/NOW/download/tools/diskshelf/ .
2	Place the file in /etc/shelf_fw on both nodes.
3	Make sure clustering is enabled by entering the following command: cf status Result: If clustering is enabled, the following message is displayed: Clustering enabled Note_____ If clustering is disabled, enable it with the cf enable command. _____
4	Install the new shelf firmware on the local node by entering the following commands: priv set advanced storage download shelf
5	After the firmware update is complete on the local node, repeat Step 3 on the partner node. Note_____ Disk shelf firmware is automatically updated on the next reboot after it is installed on each disk shelf. _____

Not enough storage

What to do when storage is full

You can add a DS14/DS14mk2 FC disk shelf to a cluster that is made up of DS14/DS14mk2 FC disk shelves without powering down your cluster. This enables you to add disk shelves without a service interruption. However, you cannot add more disk shelves to your cluster than your cluster can support. See the NOW Web page at <http://now.netapp.com> for the amount of storage your cluster can support.

Caution——
Hot-adding a disk shelf is different from hot-swapping a disk shelf. Hot-swapping a disk shelf, which means removing an existing disk shelf and installing a new one in its place, is not supported. Your system supports hot-adding of disk shelves only.

Considerations for adding storage to a cluster

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...	Then...
FC9 disk shelves	adding 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...	Then...
DS14/DS14mk2 FC 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 loop speed switch must be set for the appropriate installation.

Note

See the System Configuration guide on the NOW Web page at <http://now.netapp.com> for additional restrictions.

Adding storage to a cluster

To add storage to a cluster, complete the following steps.

Step	Action
1	Determine whether shelf counting is enabled by entering the following command: cf nfo status

Step	Action	
2	If disk shelf counting is...	Then...
	Enabled	1. Disable negotiated failover in your cluster by entering the following command from the node console: cf nfo disable disk_shelf 2. Go to Step 3.
	Not enabled	Go to Step 3 .
3	Install the new disk shelf in your equipment rack, as described in the <i>DiskShelf14 and DiskShelf14mk2 FC Hardware Guide</i> .	
4	Ground the new disk shelf, then give the disk shelf a unique shelf ID.	
5	Verify that the new disk shelf ID is unique by entering the following command: fcstat device_map Result: Shelf IDs for all Fibre Channel devices are listed.	

Step	Action	
6	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 1 Gb. 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. _____
7	Plug in the new disk shelf and turn it on.	
8	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.	
9	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.	
10	Set both terminate switches on the new disk shelf to On.	
11	Set both terminate switches on the previous disk shelf to Off.	

Step	Action	
12	If disk shelf counting was...	Then...
	Enabled	1. Reenable negotiated failover in your cluster by entering the following command from the node console: cf nfo enable disk_shelf 2. Use the cluster.
	Not enabled	Use the cluster.

LRC needs upgrading to ESHs or LRC/ESH fails

About this section

This section describes how to perform the following tasks:

- ◆ “Upgrading LRC modules to ESH modules” on page 77
- ◆ “Hot-swapping a single LRC/ESH module” on page 79
- ◆ “Testing the LRC/ESH modules” on page 80

Reasons for upgrading an LRC

If your cluster has disk shelves that have LRC modules and you add a DS14/DS14mk2 FC disk shelf that has ESH modules, you need to replace the LRC modules with ESH modules in the existing DS14 disk shelves.

This section describes how to upgrade your LRC modules to ESH modules. It also describes how to hot-swap a failed LRC or ESH module.

Upgrading LRC modules to ESH modules

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

Step	Action	
1	If the cluster is a...	Then...
	Standard cluster, mirrored cluster, nonswitched 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 terminate 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 a cluster takeover...	Then...
	Was performed in Step 1	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.
	Was not performed in Step 1	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 " Step 5: Checking and testing cluster configuration " on page 57.	

Hot-swapping a single LRC/ESH module

You can hot-swap an LRC or ESH module in a cluster. To hot-swap a single LRC/ESH module that belongs to the partner node, complete the following steps.

Step	Action	
1	If your system is...	Then...
	Setup as a fabric MetroCluster	Go to Step 2 .
	Not setup as a fabric MetroCluster	Take over the partner filer's disk shelves by entering the following command on the local filer: cf takeover
2	Note the terminate switch position of the LRC, or terminate 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	If cf takeover was...	Then...
7	Performed in Step 1 .	Return control of the partner filer's disk shelves by entering the following command: cf giveback
8	Not performed in Step 1 .	Test the replacement LRC/ESH module, as described in " Testing the LRC/ESH modules " on page 80.

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 that 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.

Failed adapters

Tasks for replacing an adapter

Replacing an adapter consists of the following tasks:

- ◆ “[Removing an adapter](#)” on page 81
- ◆ “[Installing an adapter](#)” on page 81

Removing an adapter

To remove an adapter from a node, complete the following steps.

Step	Action
1	Ground yourself using an antistatic wrist strap.
2	Take over the node that has the adapter you want to replace by entering the following command from its partner’s console: cf takeover Result: The node takes over the partner node. You see a message similar to the following: takeover completed
3	Turn off the power to the target node, then unplug it from the power source.
4	Remove the appropriate cabling to the adapter, then remove the adapter, as described in the <i>Hardware and Service Guide</i> .
5	Go to “ Installing an adapter ” on page 81.

Installing an adapter

To install an adapter in a node, complete the following steps.

Step	Action
1	While grounded, install the new adapter, as described in the <i>Hardware and Service Guide</i> .
2	Close the node and reconnect the appropriate cables to the adapter.

Step	Action
3	Plug the node back into the power source, and then apply power.
4	After the node boots and is at the 'Waiting for Giveback' prompt, run giveback 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

Removing clustering

Tasks for removing clustering

- Removing clustering is a two-step process:
- ◆ Disable clustering software, such as the cluster license
 - ◆ Remove the clustering hardware

Disabling clustering software

To disable the clustering software, complete the following steps.

Step	Action
1	Enter the following command on either node console, as shown in the following example: cf disable
2	Disable the cluster license using the following command, as shown in the following example: license delete cluster Note _____ If you are using a version of Data ONTAP earlier than 6.4, the command is <code>license cluster=disable</code> . _____
3	Open the /etc/rc file with a text editor and remove references to the partner node in the ifconfig entries, as shown in the following example: Original entry: <code>ifconfig e0 199.9.204.254 partner 199.9.204.255</code> Edited entry: <code>ifconfig e0 199.9.204.254</code>
4	Repeat Steps 1 through 3 on the partner node.
5	Go to “Removing clustering hardware” on page 84.

Removing clustering hardware

To remove the clustering hardware and recable the node's disk shelves, complete the following steps.

Step	Action	
1	Halt both nodes by entering the following command on each console: halt	
2	Turn off the power to each node, then turn off the power to the disk shelves and unplug them from the power source.	
3	Ground yourself, then remove the cluster interconnect cables from both nodes.	
4	In a FAS940 or FAS960, open both nodes and remove the cluster adapters, and then close the nodes. In a FAS980, if a regular cluster is being converted to a stand-alone filer, the NVRAM5 adapter must be removed from slot 11 and installed in slot 10. If a metrocluster is being converted to a stand-alone filer, the VI cluster adapter must be removed from slot 11. See the <i>Hardware and Service Guide</i> for additional information.	
5	Unplug the local node from the partner node's Channel B disk shelf and do the same for the partner node.	
6	If you are converting to a...	Then...
	System with nonmirrored disks	1. Disconnect all cabling from the Channel B loop on the local node. 2. Repeat Step 1 for the partner node.
	System with mirrored disks or redundant Channel B loop	1. Connect the local node to the open Channel B loop in its local disk shelves, as described in the appropriate disk shelf guide. 2. Repeat Step 1 for the partner node.
7	Power on the disk shelves, then the individual nodes, monitoring the system console for error messages during the boot process.	

Step	Action
8	Run all system diagnostics at the boot prompt by entering the following command on the system console: sysdiag
9	Unset the partner sysid by entering the following command at the prompt: unsetenv partner-sysid
10	Boot the node by entering the following command: boot
11	Check clustering status by entering the following command: cf status Result: If clustering is disabled, you see the following output: Cluster monitor not initialized
12	Repeat Step 1 through Step 11 for the partner node.

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 appliance	NetApp FAS900 Series appliances	FAS940 FAS960 FAS980	2004

to which this declaration relates conform to the following standards:

EN 60950:2002, 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: 1994+All, Safety of Laser/LED Equipment

EN 61000-3-2:2000 Limits for Harmonic Current Emissions

EN 61000-3-3:1995/A1:2001 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: 210-00583

Index

A

adapters

- cabling cluster adapters 23
- dual-port FC-AL 15
- installing 81
- NVRAM5 6, 23
- removing 81

adding

- disk shelf to mirrored cluster 72
- disk shelf to standard cluster 73

audience v

B

bring up 51

- configuring interfaces 54
- enabling licenses 49
- first time 48
- manually setting options 52
- manually setting parameters 53

C

cabling

- Channel A loop in mirrored cluster 21
- Channel A loop in standard cluster 19
- Channel B loop in mirrored cluster 26
- Channel B loop in standard cluster 24
- dual-port FC-AL 15
- FC-ALs in mirrored cluster 15
- location of dual-port FC-ALs 15
- preparing equipment racks for 18
- preparing the system cabinet for 17
- required 3

cluster adapter, cabling 23

cluster conversion requirements

- MetroCluster 45

cluster types

- Fabric Attached MetroCluster overview 9
- Fabric Attached MetroCluster, description of 5
- installed in equipment racks 14
- installed in system cabinets 14

mirrored cluster, description of 5

mirrored cluster, overview of 7

standard cluster overview 6

standard cluster, description of 5

clusters

- adding a disk shelf to a mirrored 72
- adding a disk shelf to a standard 73
- checking configuration with the configurator script 53
- configuring interfaces 54
- disabling 83
- first-time bring up of 48
- removing hardware from 84
- setting options in 51
- setting parameters in 52
- testing 57

commands

- configshow 34
- configure 33
- license add xxxxx 42
- switchdisable 33
- switchenable 34
- switchname 34
- version 32

concepts, cluster types 5

configuring

- dedicated interfaces 56
- interfaces 55
- shared interfaces 56
- standby interfaces 56

conventions

- command v
- formatting vi
- keyboard vi

D

Data ONTAP

- behavior in a Fabric Attached MetroCluster 10
- behavior in a mirrored cluster 7
- behavior in a standard cluster 6
- updating 66

Disk ownership, description 10

E

equipment racks, installing in 14

F

Fabric Attached MetroCluster

- advantages of 10
- behavior of Data ONTAP 10
- cabling the local node and disk shelves 35
- cabling the remot node and disk shelves 39
- configuring the switches in 32
- contents 9
- description 5
- disk ownership described 10
- planning worksheet example 31
- restrictions 11

firmware

- updating disk 69
- updating disk shelf 71
- updating node 68

H

hardware

- installing an adapter 81
- removing an adapter 81
- removing clustering 84

I

installing, adapters 81

interfaces

- dedicated description 54
- shared description 54
- standby description 54

L

licenses

- disabling the cluster 83
- enabling 49

M

manuals, required 2

MetroCluster, converting to a Fabric Attached 45

mirrored cluster

- adding a disk shelf 72
- advantages of 7
- behavior of Data ONTAP 7
- cabling the Channel A loop 21
- cabling the Channel B loop 26
- contents 7
- description 5
- FC-ALs in 15
- restrictions 8

O

options, manually setting 52

P

parameters, manually setting 53

preparing equipment racks 18

R

required

- cables 3
- manuals 2
- tools 3

restrictions

- in Fabric Attached MetroCluster 11
- in mirrored clusters 8
- in standard clusters 6

S

setting options 51

software

- disabling clustering 83
- updating Data ONTAP 66

special messages vii

standard cluster

- adding a disk shelf 73
- advantages of 6
- behavior of Data ONTAP 6

- cabling the Channel A loop 19
- cabling the Channel B loop 24
- contents 6
- description 5
- restrictions 6
- switches
 - cabling the local node and disk shelves 35
 - cabling the local remote node and disk shelves 39
 - checking firmware revision 32
 - configuring for Fabric Attached MetroCluster 32
 - disabling 33
 - setting the parameters 33
- system cabinets
 - installing in 14
 - preparing for cabling 17

T

- terminology v
- tools, required 3
- troubleshooting
 - disaster recovery 63
 - giveback failure 60
 - takeover failure 60

U

- updating
 - Data ONTAP 66
 - disk firmware 69
 - disk shelf firmware 71
 - node firmware 68

