

NetApp® Hardware **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 210-00704
September 2004

Copyright and trademark information

Copyright information

Copyright © 1994–2004 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 or purchase of this product does 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, the Network Appliance design, NetApp—the Network Appliance Company, FAServer, NearStore, NetCache, and WAFL are registered trademarks of Network Appliance, Inc. in the United States and other countries. 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. SnapCopy is a registered trademark of Network Appliance, Inc. in the European Union and a trademark of Network Appliance, Inc. in the United States. DataFabric, FilerView, SecureShare, SnapManager, SnapMirror, and SnapRestore are registered trademarks of Network Appliance, Inc. in the United States. Data ONTAP and Snapshot are trademarks 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, SnapMover, SnapVault, vFiler, and Web Filer are trademarks of Network Appliance, Inc. in the United States. Spinnaker Networks, the Spinnaker Networks logo, SpinAccess, SpinCluster, SpinFS, SpinHA, SpinMove, SpinServer, and SpinStor are registered trademarks of Spinnaker Networks, LLC in the United States, other countries, or both. SpinAV, SpinManager, SpinMirror, and SpinShot are trademarks of Spinnaker Networks, LLC in the United States, other countries, or both.

Apple is a registered trademark and QuickTime is a trademark of Apple Computer, Inc. in the United States and/or other countries.

Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries.

RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

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

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

Network Appliance NetCache is certified RealSystem compatible.

Table of Contents

	Preface	vii
	Safety Information (Sicherheitshinweise)	xi
Chapter 1	Preparing for the Installation.	1
	Clustered failover and single-point-of-failure analysis.	2
	Clustering concepts	6
	Standard clusters	7
	Mirrored clusters and nonswitched MetroClusters.	9
	Fabric Attached MetroCluster	11
	Required documentation, tools, and equipment	14
Chapter 2	Cabling a Standard or Mirrored Cluster	19
	Cluster installation types	20
	Preparing your equipment	21
	Cabling a FAS900 series cluster	23
	Cabling Channel A	25
	Cabling Channel B	29
	Cabling the cluster adapters	33
	Cabling an F800 series cluster	34
	Cabling Channel A	35
	Cabling Channel B	38
	Cabling the cluster adapters	41
Chapter 3	Cabling a MetroCluster.	43
	Cabling a new Fabric Attached MetroCluster	44
	Stage 1: Planning the Fabric Attached MetroCluster	45
	Stage 2: Configuring the local switches	46
	Stage 3: Cabling the local node	49
	Stage 4: Cabling the switches and nodes at the partner site	53
	Stage 5: Configuring the nodes	58
	Cabling a nonswitched MetroCluster.	60
	Converting a cluster into a Fabric Attached MetroCluster	61

Chapter 4	Configuring the Cluster	63
	Stage 1: Bringing up the cluster for the first time	64
	Stage 2: Enabling licenses	66
	Stage 3: Setting options and parameters	68
	Stage 4: Configuring networking interfaces	72
	Stage 5: Checking and testing cluster configuration	76
	Testing F800 series clusters	77
	Testing FAS900 series clusters.	80
 Chapter 5	 Basic Troubleshooting	 83
	Takeover or giveback failures	84
	Hot-swap disk loop error messages on Fabric Attached MetroClusters	87
	Recovering from a disaster	88
	Not enough storage	90
	LRC needs upgrading or LRC, ESH, or ESH2 fails	95
	Upgrading LRC modules to ESH or ESH2 modules.	96
	Hot-swapping an LRC, ESH, or ESH2.	98
	Hot-upgrading or hot-downgrading a disk shelf	103
	Failed adapters	107
	Removing clustering	109
	Removing F800 series clustering hardware	110
	Removing FAS900 series clustering hardware	112
	Failover event cause-and-effect table.	114
 Chapter 6	 Adding a Disk Shelf to an Existing Cluster	 121
	Installing a disk shelf in an existing mixed-shelf loop	122
	Hot-adding a DS14/DS14mk2 FC to an existing loop	126
 Appendix A	 Communications Regulations	 131
	Regulatory notices	132
	Declaration of Conformity	134
 Appendix B	 Feature Update Record	 135

Index 137

Preface

About this guide

This guide describes the following tasks and topics:

- ◆ Preparing to install an F800 or a FAS900 series cluster
- ◆ Cabling a pair of F800 or FAS900 series appliances into a cluster
- ◆ Configuring the cluster when you first bring it up
- ◆ Performing 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*. This is primarily a cluster *hardware* guide covering hardware clustering issues. Some basic software commands are also covered in this guide. See the *Data Protection Online Backup and Recovery Guide* for more information about clustering software and commands.

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
- ◆ Cable a new Fabric Attached or nonswitched MetroCluster.

Terminology

This guide uses the following terms:

- ◆ *FAS900* refers to NetApp appliances belonging to that series, such as the FAS980c.
- ◆ *F800* refers to NetApp filers belonging to that series, such as the F880c.
- ◆ *Mirrored cluster* refers to a cluster where there are 2 copies, or plexes, of the data.
- ◆ *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 standard cluster fashion, and where the nodes are separated by more than 500 meters and up to 30 kilometers.
- ◆ *Nonswitched MetroCluster* refers to a cluster running syncmirror_local and cluster_remote, where the nodes are separated without a switch by more than 30 meters and up to 500 meters.
- ◆ *Node* refers to either the local or partner appliance.

- ◆ *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.
- ◆ *Appliance* refers to any F800 series filer or FAS900 series appliance.
- ◆ *Disk shelf* refers to a Fibre Channel StorageShelf FC7, Fibre Channel StorageShelf FC8, Fibre Channel StorageShelf FC9, or Fibre Channel DiskShelf14/DS14mk2.
- ◆ The LRC (*Loop Resiliency Circuit*) keeps the FC-AL loop intact during the addition and removal of disks within a disk shelf. It also contains the enclosure services processor, which communicates the environmental data of the disk shelf. The LRC reconditions the signal so that there is no accumulated error in the data signals. As the signal moves, waveforms may distort slightly. Over many hops these distortions would otherwise accumulate, causing high error rates. The LRC prevents this.
- ◆ The ESH (*Embedded Switching Hub*) is a device that provides a means of managing a FC-AL loop in an intelligent manner such that a single drive failure will not take down the loop. It also contains the enclosure services processor, which communicates the environmental data of the disk shelf.
- ◆ ESH2 refers to a second-generation ESH module.

Command conventions

You can enter commands on the system console or from any client that can obtain access to the filer using a Telnet session. In examples that illustrate commands executed on a UNIX® workstation, the command syntax and output might differ, depending on your version of UNIX.

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.

What is in this guide...	What it means...
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.
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 Abzweigleitungen angeschlossen sein.
- ◆ Zum Schutz vor Körperverletzung oder Sachschäden am Gerät lassen Sie die inneren Bauteile stets vor dem Berühren abkühlen. Sorgen Sie dafür, dass das Gerät richtig abgestützt ist oder fest aufrecht steht, bevor Sie neues Zubehör einbauen.
- ◆ Dieses Gerät ist für die Einspeisung aus einer geerdeten Netzverbindung ausgelegt. Der Netzstecker mit Erdungsvorrichtung ist ein wichtiger Sicherheitsschutz. Zum Schutz vor elektrischem Schlag oder Sachschäden am Gerät die Erdung nicht abschalten.
- ◆ Das Gerät ist mit einer oder mehreren auswechselbaren Batterien ausgestattet. Bei unsachgemäßem Auswechseln der Batterie besteht Explosionsgefahr. Batterien nur mit dem vom Hersteller empfohlenen Typ oder entsprechenden Typen ersetzen. Gebrauchte Batterien sind gemäß den Anweisungen des Herstellers zu entsorgen.

Warnhinweis für Geräte mit mehrfachen Netzananschlussleitungen



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

ACHTUNG

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

About this chapter This chapter describes how to prepare for the installation of a clustered system.

Topics in this chapter This chapter contains the following information:

- ◆ “[Clustered failover and single-point-of-failure analysis](#)” on page 2
- ◆ “[Clustering concepts](#)” on page 6
- ◆ “[Required documentation, tools, and equipment](#)” on page 14

Clustered failover and single-point-of-failure analysis

Overview

Clustered failover, a high availability data clustering solution, can be used to further increase the uptime of filers. It is primarily designed to protect against filer head failure by transferring the data service from the failed filer to its partner filer. Besides filer head failure, clustered failover can also protect against other hardware failures such as a network interface card, FC-AL loop, and a LRC. Clustered failover is also an effective tool for reducing planned downtime of one of the filers.

In this section, all hardware components in a single filer and in clustered filers are reviewed to determine if they represent a single-point-of-failure (SPOF). An examination of how clustered failover can protect against SPOF and maintain data access in the unlikely event of hardware problems is also presented.

Single-point-of-failure definition

SPOF represents the failure of a single hardware component which can lead to loss of data access or potential loss of data. SPOF does not include multiple/rolling hardware errors such as double disk failure, dual LRC failure, etc.

All hardware components have demonstrated very good reliability with low failure rates. In the event of a hardware component failure, such as filer head failure or adapter failure, clustered failover can be utilized to provide continuous data availability and preserve data integrity for client applications and users.

Best practices when deploying clustered failover

There are many things that can be done to ensure clustered filers achieve maximum uptime. They include:

- ◆ Make sure the clustered filers and shelves are on different power supplies/grids such that a single power outage will not affect both filers and shelves.
- ◆ Use VIFs (virtual interfaces) to provide redundancy and improve availability of network communication.
- ◆ Follow NetApp documented procedures in the *Upgrade Guide* when upgrading clustered filers.
- ◆ Maintain consistent configuration between clustered filers. Inconsistent clustered filers configuration is often a culprit for failover problems.

- ◆ Test the failover capability routinely (e.g. during planned maintenance) to ensure proper clustered filers configuration.
- ◆ Make sure clustered filers have sufficient resources to adequately support workload during takeover mode.

Eliminating filer SPOFs with clustered failover

The following table shows which hardware components are SPOFs, and how clustered failover (CF) can eliminate these SPOFs to improve data availability.

Hardware components	SPOF		How clustered failover eliminates SPOF
	Single filer	Cluster filer	
Filer head	Yes	No	If a filer head fails, CF will automatically failover to its partner filer and serve data from the takeover filer.
NVRAM	Yes	No	If an NVRAM card fails, CF will automatically failover to its partner filer and serve data from the takeover filer.
CPU fan	Yes	No	If the CPU fan fails, the filer head will gracefully shut down. CF will automatically failover to its partner filer and serve data from the takeover filer.
Multiple NIC cards with VIFS	No	No	If one of the networking links fails, the networking traffic will be automatically sent over the remaining networking links on the filer. No failover is needed in this situation. If all the NIC cards fail, the operator can initiate a failover to its partner filer and serve data from the takeover filer. Note_____ Customers should always use multiple NIC cards with VIFS to improve networking availability for both single filer and cluster filers. _____

Hardware components	SPOF		How clustered failover eliminates SPOF
	Single filer	Cluster filer	
Single NIC card	Yes	No	If a NIC card fails, the operator can initiate a failover to its partner filer and serve data from the takeover filer.
FC-AL card	Yes	No	If a FC-AL card for the A-loop fails, the filer will automatically reboot itself. At boot time, the clustered filers will invoke the negotiated failover feature to determine which filer is best for serving data based on the disk shelf count. If the FC-AL card for the B-loop fails, the failover capability will be disabled, but both filers will continue to serve data to their respective applications and users with no impact or delay.
FC-AL cable (filer-to-shelf, shelf-to-shelf)	Yes	No	If a FC-AL loop breaks, it could lead to a filer reboot, depending on the shelf type. At boot time, the cluster filers will invoke the negotiated failover feature to determine which filer is best for serving data based on the disk shelf count.
LRC	Yes	No	If a LRC fails, it could lead to a filer reboot. At boot time, the cluster filers will invoke the negotiated failover feature to determine which filer is best for serving data based on the disk shelf count.
Hard drive	No	No	If a disk fails, the filer can re-construct data from the RAID 4 parity disk. No failover is needed in this situation.

Hardware components	SPOF		How clustered failover eliminates SPOF
	Single filer	Cluster filer	
Disk shelf (including backplane)	No	No	Disk shelf is a passive backplane with dual power supplies, dual fans, dual LRCs, and dual FC-AL loops. It is the single most reliable component in a filer with a MTBF > 2 million hours (i.e. 228 years).
Power supply	No	No	Both the filer head and disk shelf have dual power supplies. If one power supply fails, the second power supply will automatically kick in. No failover is needed in this situation.
Fan (filer or disk shelf)	No	No	Both the filer head and disk shelf have multiple fans. If one fan fails, the second fan will automatically provide cooling. No failover is needed in this situation.
Cluster adapter card	N/A	No	If a cluster adapter card fails, the failover capability will be disabled, but both filers will continue to serve data to their respective applications and users.
Cluster interconnect cable	N/A	No	The cluster adapter card supports dual cluster interconnect cables. If one cable fails, the cluster traffic (heartbeat and NVRAM data) will be automatically sent over the second cable with no delay or interruption. If both cables fail, the failover capability will be disabled, but both filers will continue to serve data to their respective applications and users.

Clustering concepts

Types of clusters

The following sections provide details about the types of clusters you can install at your site:

- ◆ [“Standard clusters”](#) on page 7
- ◆ [“Mirrored clusters and nonswitched MetroClusters”](#) on page 9
- ◆ [“Fabric Attached MetroCluster”](#) on page 11

Standard clusters

What standard clusters contain

A standard cluster contains two sets of Fibre Channel disk shelves, a local node, and a partner node. Both nodes must be identical (i.e. you cannot cluster a FAS920 with a FAS980).

The nodes are connected to each other through an IB cluster adapter or NVRAM5 adapter, which allows one node to serve data to the disks of its 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.

Note

DS14mk2 disk shelves using ESH or ESH2 modules operating at 2-Gb Fibre Channel speeds can attain a maximum node separation distance of only 300 meters, not 500.

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. See the *Data ONTAP System Administration 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 or NVRAM5 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 nonswitched MetroClusters

About mirroring

Mirrored clusters and nonswitched 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 multiple disk failure within your volume.
- ◆ Mirrored data on your cluster survives a double disk error in a raid group. This would cause a loss of the volume on nonmirrored volumes.
- ◆ 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 ONTAP Data Protection Online Backup and Recovery Guide* for detailed information about disaster protection through mirroring.

How Data ONTAP works with nonswitched 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 nonswitched MetroCluster, or converting an existing cluster to a mirrored cluster or nonswitched MetroCluster:

- ◆ Disks in the same volume must be part of the same plex, with its mirror in a different plex.
- ◆ Disk pool 1 volumes must be cabled to FC-AL or dual Fibre Channel adapters in node slots 8 through 11.
- ◆ The two halves of the mirrored disk should have no common components.
- ◆ Local SyncMirror® software must be licensed and enabled on the local and partner nodes.
- ◆ MetroClusters must have cluster_remote licensed 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, but less than 500 meters, must use an LC-LC 50/125 cable to connect their IB cluster adapters or NVRAM5 adapters. Visit <http://now.netapp.com> for more information.

Fabric Attached MetroCluster

Definition of a switch fabric

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

For example, the first fabric in a MetroCluster begins from Switch 1 on the local side and is completed by connecting the ISL to the first switch on the remote side—Switch 3. The second fabric is created using Switch 2 on the local side, connected through a second ISL cable, to the second switch on the remote 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, and quadrants 2 and 4 belong to Pool 1. Each switch bank contains a Pool 0 and a Pool 1.

See the *Data ONTAP Data Protection Online Backup and Recovery Guide* for a detailed discussion of disk pools.

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



Advantages of Fabric Attached MetroClusters

The advantages of Fabric Attached MetroClusters are the same as those of both standard and mirrored clusters. See [“Standard clusters”](#) on page 7 and [“Mirrored clusters and nonswitched MetroClusters”](#) on page 9 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, ensuring no single point of failure.

How Data ONTAP works with Fabric Attached MetroClusters

Data ONTAP functions the same as it does in a nonswitched MetroCluster. For more information, see [“How Data ONTAP works with nonswitched MetroClusters”](#) on page 9.

Setup requirements and restrictions

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

- ◆ Each failover 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 appliance.
- ◆ 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.

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>42U 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>F800 Series Hardware Guide</i>	This guide describes how to install an F800 series filer, connect it to a network, and bring it up for the first time.
<i>FAS900 Series Hardware and Service Guide</i>	This guide describes how to install a 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 you can 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 Online Backup and Recovery Guide</i>	This guide describes, among other topics, what a synchronous mirror is, how to enable it, and how to maintain it.
<i>System Administration Guide</i>	This guide describes filer administration and clustering software commands.

Manual name	Description
<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* at <http://now.netapp.com> to confirm your appliance type, storage capacity, etc.

Required equipment	Standard or mirrored cluster	Nonswitched MetroCluster	Fabric Attached MetroCluster
Appliance type	Two of the same type of FAS900 series appliance or two of the same type of F800 series filer.		
Storage	See the <i>System Configuration Guide</i> .		
Cluster interconnect adapter	IB cluster adapter (The NVRAM5 functions as the cluster interconnect adapter on FAS900 series filers).		VI-MC cluster adapter
FC-AL or FC HBA (FC HBA for Disk) adapters	Minimum of two FC-AL adapters.	Two Fibre Channel HBA adapters. Note _____ The ports on the Fibre Channel HBA adapters are labeled 1 and 2. However, the software refers to them as a and b. You see these labeling conventions in the user interface and system messages displayed on the console. _____	

Required equipment	Standard or mirrored cluster	Nonswitched MetroCluster	Fabric Attached MetroCluster
Fibre Channel switches	N/A	N/A	Two pairs of Brocade 3800 or 3200 switches, the same type per node.
SFP (Small Form Pluggable) modules	N/A	N/A	Two per ISL (Inter Switch Link)—4 total, long distance. One per switch port used, short distance.
NVRAM5 cluster adapter media converter	Only if using fiber cabling.	Only if using fiber cabling.	N/A

Required equipment	Standard or mirrored cluster	Nonswitched MetroCluster	Fabric Attached MetroCluster
Fibre Channel switches	N/A	N/A	Two pairs of Brocade 3800 or 3200 switches, the same type per node.
SFP (Small Form Pluggable) modules	N/A	N/A	Two per ISL (Inter Switch Link)—4 total, long distance. One per switch port used, short distance.
NVRAM5 cluster adapter media converter	Only if using fiber cabling.	Only if using fiber cabling.	N/A

Required equipment	Standard or mirrored cluster	Nonswitched MetroCluster	Fabric Attached MetroCluster
Cables (provided with shipment unless otherwise noted)	◆ 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 ◆ For systems using NVRAM5 - two 4xIB copper cables, or two 4xIB optical cables Note_____ Longer optical cables must be purchased separately for cabling distances greater than 30 meters. ◆ 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 cables ◆ Four SC/LC or LC/LC 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 or DB-9 to DB-9 shelf-to-shelf cables for DS14 disk shelves with ESH or ESH2 modules

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 43 for information about cabling a MetroCluster.

Topics in this chapter

The following topics are covered in this chapter:

- ◆ “[Cluster installation types](#)” on page 20
- ◆ “[Preparing your equipment](#)” on page 21
- ◆ “[Cabling a FAS900 series cluster](#)” on page 23
- ◆ “[Cabling an F800 series cluster](#)” on page 34

Cluster installation types

Types of cluster installations

There are two types of cluster installations:

- ◆ Clusters in a system cabinet
- ◆ Clusters in an equipment rack

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 you 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 cluster you ordered 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 cluster you ordered 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 *42U 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, hardware guide, or Quick Start guide that came with your equipment.

Preparing your equipment

Where to install a cluster

You can install a cluster in one or more system cabinets or in equipment racks. Install the equipment as described in the appropriate section:

- ◆ “[Preparing the system cabinets](#)”
- ◆ “[Preparing the equipment racks](#)” on page 22

Preparing the system cabinets

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	a. Connect the nodes to the network. b. Connect the system cabinet to an appropriate power source. c. Apply power to the system cabinet and configure the cluster, as described in Chapter 4, “Configuring the Cluster,” on page 63.
	Multiple system cabinets	a. Remove the front and rear doors and any side panels that need to be removed to connect the system cabinets together. b. Connect the system cabinets together, as described in the <i>42U System Cabinet Guide</i>. c. Connect the nodes to the network.

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 Fibre Channel disk shelf guide, hardware guide, or Quick Start guide that came with your equipment.
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> or the <i>F800 Series Hardware Guide</i> .

Cabling a FAS900 series cluster

Type of 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. The dual-port FC-AL adapter allows your appliance to support more storage than a single-port FC-AL adapter.

Minimum number of 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 specific number of dual-port FC-AL adapters you need depends on the number of disk shelves you have in each loop. See the *System Configuration Guide* at <http://now.netapp.com> for more information about your hardware.

Recommended slot locations

The following table lists recommended slot locations for the dual-port FC-AL adapters in a single node of a FAS900 series cluster. Check the *System Configuration Guide* at <http://now.netapp.com> to confirm the slot assignment information.

Cluster type	Slot number on the back of the node
Standard cluster—One dual-port FC-AL adapter	Slot 2
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:

- ◆ [“Cabling Channel A”](#) on page 25
- ◆ [“Cabling Channel B”](#) on page 29
- ◆ [“Cabling the cluster adapters”](#) on page 33


Channel A loop
cabling differences

The procedure for cabling the Channel A loops for a standard cluster and mirrored cluster are different. Follow the directions in the appropriate section:

- ◆ “Cabling the Channel A loops in a standard cluster”
- ◆ “Cabling the Channel A loops in a mirrored cluster” on page 27

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. The circled numbers in the diagram correspond to the step numbers in the table.




Step	Action
1	Cable the dual-port FC-AL adapter in slot 2, Top port to the disk shelf 1 Channel A LRC, ESH, or ESH2 Input port.

Step	Action
2	Cable the disk shelf 1 Channel A LRC, ESH, or ESH2 Output port to the disk shelf 2 Channel A LRC, ESH, or ESH2 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 or ESH modules to Off on all but the last disk shelf in the loop, then set the LRC or ESH terminate switches on the last disk shelf to On. Note_____ ESH2 modules are self-terminating, and therefore do not have a terminate switch. _____
5	Repeat Step 1 through Step 4 for the partner node.
6	Go to “ Cabling Channel B ” on page 29.

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 adapter in slot 8, Top port of the local node to the disk shelf 1 Channel A LRC, ESH, or ESH2 Input port of disk pool 1, then cable the dual-port FC-AL in slot 5, Top port of the local node to the disk shelf 3 LRC, ESH, or ESH2 Input port of disk pool 0.
2	Cable the local node's disk shelf 1 Channel A LRC, ESH, or ESH2 Output port to the disk shelf 2 Channel A LRC, ESH, or ESH2 Input port for both disk pools. Label this cable with a solid-colored label.

Step	Action
3	Repeat Step 2 connecting Channel A output to input, for any remaining disk shelves in the disk pool. Note_____ The illustration shows only two disk shelves per disk pool. The number of disk shelves per pool might be different for your appliance. _____
4	Set the terminate switches on the LRC or ESH modules to Off on all but the last disk shelf in the disk pool, then set the LRC or ESH terminate switches on the last disk shelf to On. Note_____ ESH2 modules are self-terminating, and therefore do not have a terminate switch. _____
5	Repeat Step 1 through Step 4 for the partner node.
6	Go to “ Cabling Channel B ” on page 29.

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. Follow the directions in the appropriate section:

- ◆ “Cabling the Channel B loops in a standard cluster”
- ◆ “Cabling the Channel B loops in a mirrored cluster” on page 31

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, Bottom port of the local node, to the partner node disk shelf 1 Channel B LRC, ESH, or ESH2 Input port.
2	Cable the local node's disk shelf 1 Channel B LRC, ESH, or ESH2 Output port to the local node's disk shelf 2 Channel B LRC, ESH, or ESH2 Input port. Label this cable with a striped label.
3	Repeat Step 2 for the remaining disk shelves in the disk pool.
4	Set the terminate switches on the LRC or ESH modules to Off on all but the last disk shelf in the disk pool, then set the LRC or ESH terminate switches on the last disk shelf to On. Note_____ ESH2 modules are self-terminating, and therefore do not have a terminate switch. _____
5	Repeat Step 1 through Step 4 for the partner node.
6	Go to “ Cabling the cluster adapters ” on page 33.

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, Bottom port of the local node to the partner node disk shelf 1 Channel B LRC, ESH, or ESH2 Input port of disk pool 1, then cable the dual-port FC-AL adapter in slot 5, Bottom port of the local node to the partner node disk shelf 1 Channel B LRC, ESH, or ESH2 Input port of disk pool 0.
2	Cable the local node's disk shelf 1 Channel B LRC, ESH, or ESH2 Output port to the local node's disk shelf 2 Channel B LRC, ESH, or ESH2 Input port of disk pool 1. Label this cable with a striped label.


Step	Action
3	Repeat Step 2 connecting Channel B output to input, for any remaining disk shelves in the disk pool. Note_____ The illustration shows only two disk shelves per disk pool. The number of disk shelves per pool might be different for your appliance. _____
4	Set the terminate switches on the LRC or ESH modules to Off on all but the last disk shelf in the disk pool, then set the LRC or ESH terminate switches on the last disk shelf to On. Note_____ ESH2 modules are self-terminating, and therefore do not have a terminate switch. _____
5	Repeat Step 1 through Step 4 for the partner node.
6	Go to “ Cabling the cluster adapters ” on page 33.

Cabling a FAS900 series cluster

Cabling the cluster adapters

Cabling the cluster adapters

To cable the cluster adapters, complete the following steps, using the diagram for reference. The IB cluster adapters are located in slot 6. NVRAM5 adapters, which function as the cluster interconnect adapters, are located in slot 11.



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 Chapter 4, “ Configuring the Cluster ,” on page 63.

Cabling an F800 series cluster

Recommended slot locations

The following table lists recommended slot locations for F800 series clusters with FC-AL adapters or with dual Fibre Channel adapters. Check <http://now.netapp.com> to confirm the slot assignment information.

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

For detailed information

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

- ◆ “[Cabling Channel A](#)” on page 35
- ◆ “[Cabling Channel B](#)” on page 38
- ◆ “[Cabling the cluster adapters](#)” on page 41

Cabling Channel A

Channel A loop cabling differences

Generally, the Channel A loop in a cluster is cabled by connecting the FC-AL adapter in the node to the LRC, ESH, or ESH2 Input port of the first disk shelf in the loop, followed by cabling the disk shelves together. This process differs only by the type of Fibre Channel adapter you have in your cluster and the type of disk shelves you have.

Note—————
If you are cabling a cluster in a system cabinet, skip Channel A cabling and go to “[Cabling the cluster adapters](#)” on page 41. The Channel A loop is already cabled for you.


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

If you have...	Then...
Nodes using dual Fibre Channel adapters with any disk shelf type	◆ You use optical cables to connect the nodes to the first disk shelf in the Channel A loop. ◆ The first port (labeled Port 1 on the adapter) of the dual Fibre Channel adapter is reserved for the Channel A loop.
DS14 or DS14mk2 FC disk shelves with nodes using copper FC-AL adapters	You use the HSSDC to DB-9 cable to connect the nodes to the first disk shelf in the Channel A loop.
DS14 or DS14mk2 FC disk shelves with ESH or ESH2 modules	All disk shelves must have ESH or ESH2 modules installed in them and you must not have mixed shelf-type loops. Use the supplied cables for your FC-AL adapters.
FC9 disk shelves with nodes using dual Fibre Channel adapters	Your first disk shelf LRC requires an MIA (media interface adapter) plugged into the Input port.

Cabling the Channel A loops

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

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



Step	Action
1	Cable the FC-AL adapter in Slot 8, Top port of the local node, to the disk shelf 1 Channel A LRC, ESH, or ESH2 Input port.
2	Cable the disk shelf 1 Channel A LRC, ESH, or ESH2 Output port to the disk shelf 2 Channel A LRC, ESH, or ESH2 Input port. Apply solid-colored labels to the cables.
3	Repeat Step 2 , connecting output to input, for the remaining disk shelves in the loop.

Step	Action	
4	Set the terminate switches on the LRC or ESH modules to Off on all but the last disk shelf in the disk pool, then set the LRC or ESH terminate switches on the last disk shelf to On.	
	Note _____ ESH2 modules are self-terminating, and therefore do not have a terminate switch.	
5	If you have...	Then...
	A mix of DS14 and DS14mk2 FC disk shelves with LRC modules	Set the 1/2 Gb switch to 1 Gb.
	A mix of DS14 and DS14mk2 FC disk shelves with ESH or ESH2 modules	Set the 1/2 Gb switch to 1 Gb.
	Only DS14mk2 FC disk shelves with ESH or ESH2 modules	Set the 1/2 Gb switch to either 1 or 2 Gb.
6	Repeat Step 1 through Step 5 for the partner node.	
7	Go to “ Cabling Channel B ” on page 38.	

Cabling Channel B

Channel B loop cabling differences

Generally, the Channel B loop in a cluster is cabled by connecting the FC-AL adapter in the local node to the LRC, ESH, or ESH2 Input port of the first disk shelf in the partner node's Channel B loop, followed by cabling the disk shelves together. This process differs only by the type of Fibre Channel adapter you have in your cluster and the type of disk shelves you have.

Note


If you are cabling a cluster in a system cabinet, skip Channel B cabling and go to [“Cabling the cluster adapters”](#) on page 41. The Channel B loop is already cabled for you.

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

If you have...	Then...
Nodes using dual Fibre Channel adapters with any disk shelf type	◆ You use optical cables to connect the nodes to the first disk shelf in the Channel A loop. ◆ The second port (labeled Port 1 on the adapter) of the dual Fibre Channel adapter is reserved for the Channel B loop.
DS14 or DS14mk2 FC disk shelves with nodes using copper FC-AL adapters	You use the HSSDC to DB-9 cable to connect the nodes to the first disk shelf in the Channel A loop.
DS14 or DS14mk2 FC disk shelves with ESH or ESH2 modules	All disk shelves must have ESH or ESH2 modules installed in them and you must not have mixed shelf-type loops. Use the supplied cables for your FC-AL adapters.
FC9 disk shelves with nodes using dual Fibre Channel adapters	Your first disk shelf LRC requires an MIA (media interface adapter) plugged into the Input port.

Cabling the Channel B loops

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



Step	Action
1	Cable the FC-AL adapter in Slot 8, Bottom port of the local node, to the partner node disk shelf 1 channel B LRC, ESH, or ESH2 Input port.
2	Cable the local node's disk shelf 1 Channel B LRC, ESH, or ESH2 Output port to the local node's disk shelf 2 Channel B LRC, ESH, or ESH2 Input port. Label this cable with a striped label.
3	Repeat Step 2 , connecting output to input, for the remaining disk shelves in the loop.

Step	Action
4	Set the terminate switches on the LRC or ESH modules to Off on all but the last disk shelf in the disk pool, then set the LRC or ESH terminate switches on the last disk shelf to On. Note_____ ESH2 modules are self-terminating, and therefore do not have a terminate switch. _____
5	Repeat Step 1 through Step 4 for the partner node.
6	Go to “ Cabling the cluster adapters ” on page 41.

Cabling the cluster adapters

Cluster cabling differences

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

See the *System Configuration Guide* at <http://now.netapp.com> for current information about supported cluster adapters for specific appliances.

Cluster adapter type	General characteristics
Cluster interconnect adapter	◆ No longer shipped with new systems. ◆ Does not support mirrored clusters.
IB cluster adapter	◆ Uses an optical cable up to 500 meters long. ◆ Requires two SC/SC cables and GBICs for distances over 30 meters. These cables are not included.
VI cluster adapter (Rev B)	◆ Uses an optical cable up to 500 meters long. ◆ Requires an LC/LC cable for distances over 30 meters (purchased separately).
VI-MC cluster adapter	◆ Connects directly to Fibre Channel switches. ◆ Requires two SC/SC cables. These cables are not included.


Cabling the IB cluster adapter

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

Step	Action
1	Plug one end of the optical cable into one of the local node's IB cluster adapter ports, then plug the other end into the partner node's corresponding IB cluster adapter port. Repeat for the second cable.
2	Gently pull on the cable ends to ensure that they are properly seated.
3	Go to Chapter 4, " Configuring the Cluster ," on page 63.

Cabling the VI cluster adapter

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



Step	Action
1	Plug one end of the optical cable into the GBIC of the local node’s VI cluster adapter, then plug the other end into the partner node’s VI cluster adapter. Note————— You must not cross-cable the cluster interconnect adapter. Cable the local node ports only to the identical ports on the partner node. —————
2	Go to Chapter 4, “Configuring the Cluster,” on page 63.

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 44
- ◆ [“Cabling a nonswitched MetroCluster”](#) on page 60
- ◆ [“Converting a cluster into a Fabric Attached MetroCluster”](#) on page 61

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 11. You must also have the following information:

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

Note

The switches ship with a default IP address (10.77.77.77) which can be used if the switches won't be attached to a network.

- ◆ Ethernet subnetmask
- ◆ Gateway address

For detailed information

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

- ◆ “[Stage 1: Planning the Fabric Attached MetroCluster](#)” on page 45
- ◆ “[Stage 2: Configuring the local switches](#)” on page 46
- ◆ “[Stage 3: Cabling the local node](#)” on page 49
- ◆ “[Stage 4: Cabling the switches and nodes at the partner site](#)” on page 53
- ◆ “[Stage 5: Configuring the nodes](#)” on page 58

Stage 1: Planning the Fabric Attached MetroCluster

Fabric Attached MetroCluster worksheet

The Fabric Attached MetroCluster worksheet is a planning aid to record specific cabling information about your Fabric Attached MetroCluster. You must identify several pieces of information that you use during configuration procedures. Recording this information can reduce configuration errors.

Local node name: _____

Cluster adapter Port 1 is cabled to Switch 1 Port__

Switch 1 name: _____

Switch 1 IP address:_____

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

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

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

Pool 0 disk shelves:

Disk 1 connected to Port ____

Disk 2 connected to Port ____

Cluster adapter Port 2 is cabled to Switch 2 Port __

Switch 2 name: _____

Switch 2 IP address:_____

FC HBA 2, port 1 is cabled to Switch 1, Port __

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

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

Pool 1 disk shelves:

Disk 3 connected to Port ____

Disk 4 connected to Port ____

Stage 2: Configuring the local switches

Configuring the switch

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

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 http://now.netapp.com for the current version of switch firmware.	
3	If the switch firmware is...	Then...
	Out of date	a. Download the switch firmware from http://now.netapp.com and install it, as described in the appropriate Brocade switch hardware reference manual. b. Reboot the switch.
	Current	Go to Step 4 .
4	Disable the switch by entering the following command: switchdisable	
5	Clear any existing configuration on the switch by entering the following command: cfgclear	

Step	Action	
6	If you are configuring for...	Then...
	FAS900 series	a. Disable the configuration by entering the following command: cfgdisable b. Save the configuration by entering the following command: cfgsave
	F800 series	Proceed to Step 7 .
7	Ensure that the switch is configured with default settings by entering the following command: configdefault	
8	Set the switch parameters by entering the following command: configure You should set the following parameters only: ◆ Fabric parameters = y ◆ domain_id = XXX (See your system administrator for this value.) ◆ Disable device probing = 1	
9	Exit the configuration utility by entering the following command: ctrl-d	
10	Reboot the switch so that the default settings take effect by entering the following command: fastboot	
11	Log in to the switch and disable the switch by entering the following command: switchdisable	

Step	Action	
12	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.	
13	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"
14	Save the new switch configuration by entering the following command: cfgsave	
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> Where <i>Metro1_switch1</i> is the name given to the current switch. Use a name that is meaningful, such as node name and switch.	
17	Check the switch settings by entering the following command: configshow	
18	Repeat Step 1 through Step 17 for the second switch on the local node.	
19	Go to " Stage 3: Cabling the local node " on page 49.	

Cabling a new Fabric Attached MetroCluster


Stage 3: Cabling the local node

Cabling the local node and disk shelves


To cable the node, complete the following steps, using the appropriate illustration for reference.

Note

You can only cable a maximum of two disk shelves per pool.



F800 series



Step	Action
1	Cable the first Fibre Channel HBA in slot 5 of the local node, to Switch 1 and Switch 2. In the example, Port A of the Fibre Channel HBA is cabled to Port 0 of Switch 1 and Port B is cabled to Port 0 of Switch 2.
2	Cable the second Fibre Channel HBA in slot 8 of the local node, to Port 1 on both switches, with Port A on Switch 1 and Port B on Switch 2.

Step	Action
3	In the example, the Pool 0 disk shelves are connected to switch port 9 on both Switch 1 and 2. However, any open set of switch ports in this quadrant can be used (limited to ports 8-11 in this configuration). Connect the local node Pool 0 disks to the switches by completing the following steps: a. Connect the Input port of the LRC, ESH, or ESH2 channel A on disk shelf 1 to port 9 on Switch 2. b. Connect the Input port of LRC, ESH, or ESH2 B on disk shelf 1 to port 9 on Switch 1. c. Connect disk shelf 1 to disk shelf 2 by connecting the Output ports of the LRC, ESH, or ESH2 of disk shelf 1 to the Input ports of the corresponding LRC, ESH, or ESH2 modules of the next disk shelf. d. Set the terminate switch (except for ESH2) on the last disk shelf in the loop to On, and set the switch to Off for all other disk shelves in the loop.
4	In the example, the Pool 1 disk shelves are connected to switch port 5 on both Switch 1 and 2. However, any open set of switch ports in this quadrant can be used (limited to ports 4-7 in this configuration). Connect the remote node Pool 1 disks to the switches by completing the following steps: a. Connect the Input port of LRC, ESH, or ESH2 A on disk shelf 3 to port 5 on Switch 2. b. Connect the Input port of B on disk shelf 3 to port 5 on Switch 1. c. Connect disk shelf 3 to disk shelf 4 by connecting the Output ports of the LRC, ESH, or ESH2 of disk shelf 3 to the Input ports of the corresponding LRC, ESH, or ESH2 modules of the next disk shelf. d. Set the terminate switch (except for ESH2) on the last disk shelf in the loop to On, and set the switch to Off for all other disk shelves in the loop.

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, or ESH2 they belong, as shown in the following example: Example: ha15*> storage show disk -p <table><tr><th>PRIMARY</th><th>PORT</th><th>SECONDARY</th><th>PORT</th><th>SHELF</th><th>BAY</th></tr><tr><td>-----</td><td>----</td><td>-----</td><td>----</td><td>-----</td><td>----</td></tr><tr><td>switch1:4.40</td><td>A</td><td>switch2:4.40</td><td>B</td><td>5</td><td>0</td></tr><tr><td>switch2:4.41</td><td>B</td><td>switch1:4.41</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>switch1:8.52</td><td>B</td><td>switch2:8.52</td><td>A</td><td>6</td><td>4</td></tr></table>	PRIMARY	PORT	SECONDARY	PORT	SHELF	BAY	-----	----	-----	----	-----	----	switch1:4.40	A	switch2:4.40	B	5	0	switch2:4.41	B	switch1:4.41	A	5	1	.						.						.						switch1:8.52	B	switch2:8.52	A	6	4
PRIMARY	PORT	SECONDARY	PORT	SHELF	BAY																																												
-----	----	-----	----	-----	----																																												
switch1:4.40	A	switch2:4.40	B	5	0																																												
switch2:4.41	B	switch1:4.41	A	5	1																																												
.																																																	
.																																																	
.																																																	
switch1:8.52	B	switch2:8.52	A	6	4																																												
6	Connect one ISL cable to any port on each switch, making sure that the ports have the same number. In the example, Port 13 is used on both switches.																																																
7	Connect the VI-MC cluster adapter to the switches, with the top port connecting to Switch 1 and the bottom port to Switch 2. Note_____ There should be one VI-MC cluster adapter connection per switch. The switch port connections do not have to be identical. In the example, the VI-MC cluster adapter connects to switch ports 7 and 12, but you can use any available switch ports. _____																																																
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 “Stage 4: Cabling the switches and nodes at the partner site” on page 53.																																																

Stage 4: Cabling the switches and nodes at the partner site

About partner site switch cabling

You need to cable the switches at the partner site so that they can work with the local switches.


Cabling switches at the partner site

To cable the partner site switches, complete the following steps.


Step	Action
1	Log into Switch 3 and set it up as described in “ Stage 2: Configuring the local switches ” on page 46.
2	Repeat Step 1 for Switch 4.

Cabling the remote node and disk shelves at the partner site

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



F800 series



Step	Action
1	Cable the first Fibre Channel HBA in slot 5 of the remote node to Switch 3 and Switch 4. In the example, the top port of the Fibre Channel HBA is cabled to Port 8 of Switch 3 and the bottom port is cabled to Port 8 of Switch 4.
2	Cable the second Fibre Channel HBA in slot 8 of the remote node, to Switch 3 and Switch 4. In the example, the top port of the Fibre Channel HBA is cabled to Port 9 of Switch 3 and the bottom port is cabled to Port 9 of Switch 4.

Step	Action
3	In the example, the Pool 0 disk shelves are connected to switch port 0 on both Switch 3 and 4. However, any open set of switch ports in this quadrant can be used (limited to ports 0-3 in this configuration). Connect the local node Pool 0 disks to the switches by completing the following steps: a. Connect the Input port of LRC, ESH, or ESH2 Channel A of disk shelf 3, to port 0 on Switch 4. b. Connect the Input port of LRC, ESH, or ESH2 Channel B of disk shelf 3, to port 0 on Switch 3. c. Connect disk shelf 3 to disk shelf 4 by connecting the Output ports of the LRC, ESH, or ESH2 of disk shelf 3 to the Input ports of the corresponding LRC, ESH, or ESH2 modules of disk shelf 4. d. Set the terminate switch (except for ESH2) on the last disk shelf in the loop to On, and set the switch to Off for all other disk shelves in the loop.
4	In the example, the Pool 1 disk shelves are connected to switch port 12 on both Switch 3 and 4. However, any open set of switch ports in this quadrant can be used (limited to ports 12-15 in this configuration). Connect the remote node Pool 1 disks to the switches by completing the following steps: a. Connect the Input port of LRC, ESH, or ESH2 Channel A on disk shelf 1 to port 12 on Switch 4. b. Connect the Input port of LRC, ESH, or ESH2 Channel B on disk shelf 1 to port 12 on Switch 3. c. Connect disk shelf 1 to disk shelf 2 by connecting the Output ports of the LRC, ESH, or ESH2 of disk shelf 1 to the Input ports of the corresponding LRC, ESH, or ESH2 modules of disk shelf 2. d. Set the terminate switch (except for ESH2) on the last disk shelf in the loop to On, and set the switch to Off for all other disk shelves in the loop.

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, or ESH2 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>switch3:0.40</td><td>A</td><td>switch4:0.40</td><td>B</td><td>5</td><td>0</td></tr><tr><td>switch4:0.41</td><td>B</td><td>switch3:0.41</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>switch3:12.52</td><td>B</td><td>switch4:12.52</td><td>A</td><td>6</td><td>4</td></tr></tbody></table>	PRIMARY	PORT	SECONDARY	PORT	SHELF	BAY	-----	----	-----	----	-----	--	switch3:0.40	A	switch4:0.40	B	5	0	switch4:0.41	B	switch3:0.41	A	5	1	.						.						.						switch3:12.52	B	switch4:12.52	A	6	4
PRIMARY	PORT	SECONDARY	PORT	SHELF	BAY																																												
-----	----	-----	----	-----	--																																												
switch3:0.40	A	switch4:0.40	B	5	0																																												
switch4:0.41	B	switch3:0.41	A	5	1																																												
.																																																	
.																																																	
.																																																	
switch3:12.52	B	switch4:12.52	A	6	4																																												
6	Connect the ISL from Switch 1 to Switch 3 and connect the ISL from Switch 2 to Switch 4.																																																
7	Connect the VI-MC cluster adapter to the switches, with the top port connecting to Switch 3 and the bottom port connecting to Switch 4. Note_____ The example shows the VI-MC cluster adapter connecting to switch ports 3 and 13, but you can use any available switch port. _____																																																
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 “ Stage 5: Configuring the nodes ” on page 58.																																																

Cabling a new Fabric Attached MetroCluster

Stage 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 nodes, complete the following steps.

Step	Action
1	Install the new Data ONTAP release on the local node and run setup.
2	License cluster and syncmirror_local on the local node by entering the following command: <pre>license add xxxxx yyyy</pre> xxxxx is the license code you received for the cluster feature and yyyy is the syncmirror_local license. 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 Step 1 through Step 3 for the remote node.
5	License cluster_remote on both nodes.
6	Set up your new mirrored volumes. See the <i>System Administration Guide</i> or the <i>Data Protection Online Backup and Recovery Guide</i> for more information.
7	Enable clustering by entering the following command on the console: <pre>cf enable</pre>

Step	Action
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 remote disks should be displayed. The local volumes should be mirrored and there should be a list 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 63. 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 “Preparing your equipment” on page 21.	
2	If you are cabling for a...	Then...
	FAS900 series	a. Cable the Channel A loops of the MetroCluster, as described in “Cabling the Channel A loops in a mirrored cluster” on page 27. b. Cable the Channel B loops of the MetroCluster, as described in “Cabling the Channel B loops in a mirrored cluster” on page 31. c. Cable the nodes together, as described in “Cabling the cluster adapters” on page 33.
	F800 series	a. Cable the Channel A loops, as described in “Cabling the Channel A loops” on page 36. b. Cable the Channel B loops, as described in “Cabling the Channel B loops” on page 39. c. Cable the nodes together, as described in “Cabling the cluster adapters” on page 41.
3	Enable the syncmirror_local and cluster_remote licenses on both nodes.	
4	Configure the cluster, as described in Chapter 4, “Configuring the Cluster,” on page 63.	

Converting a cluster into 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, appliance firmware, and disk firmware, as described in the <i>Data ONTAP Upgrade Guide</i> , making sure to shut down the nodes to the OK prompt.	
2	If you are converting...	Then...
	a FAS980 system	Move the NVRAM5 adapter from slot 11 to slot 10, and add the VI-MC cluster adapter to slot 11. For information about removing adapter cards, see “ Failed adapters ” on page 107.
	all other platforms	Update your hardware by replacing components, as described in “ Failed adapters ” on page 107. 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 “ Stage 1: Planning the Fabric Attached MetroCluster ” on page 45.	
4	Set up and configure the local switches, as described in “ Stage 2: Configuring the local switches ” on page 46.	
5	Cable the local node, as described in “ Stage 3: Cabling the local node ” on page 49.	

Step	Action	
6	If you are converting a...	Then...
	Nonswitched MetroCluster	Proceed to Step 7 .
	Standard cluster	Set up mirroring and configure the local node, as described in “ Stage 3: Cabling the local node ” on page 49.
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 “ Stage 4: Cabling the switches and nodes at the partner site ” on page 53.	
9	Configure the cluster, as described in “ Stage 5: Configuring the nodes ” on page 58.	

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 topics:

- ◆ [“Stage 1: Bringing up the cluster for the first time”](#) on page 64
- ◆ [“Stage 2: Enabling licenses”](#) on page 66
- ◆ [“Stage 3: Setting options and parameters”](#) on page 68
- ◆ [“Stage 4: Configuring networking interfaces”](#) on page 72
- ◆ [“Stage 5: Checking and testing cluster configuration”](#) on page 76

Stage 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 and wait 30 seconds for the electronics to complete initialization.
3	Turn on the power to both nodes and wait for the system prompt. Caution _____ Press the Delete key to stop booting if the system starts to boot. _____

Step	Action	
4	If the nodes of your cluster...	Then...
	Use a CompactFlash™ card, and are CFE (Common Firmware Environment) based systems (e.g. FAS200 series)	Run the diagnostics program from the CompactFlash card by entering the following command at the prompt: boot_diags
	Use a CompactFlash card, and are not CFE based (e.g. F825, F880, NearStore series, FAS900 series)	Run the diagnostics program from the CompactFlash card by entering the following command at the prompt: boot_diags
	Do not use a CompactFlash card (e.g. F810, F820, F840, C3100, C6100)	Run the diagnostics program from the motherboard flash by entering the following command at the prompt: sys-diags
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 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 “ Stage 2: Enabling licenses ” on page 66.	

Stage 2: Enabling licenses

About licenses

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.

If you are setting up a mirrored cluster, you must also enable the syncmirror_local license for synchronous mirroring. If you are setting up a cluster with nodes in separate locations, you must also enable the cluster_remote license.

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 Administration Guides* 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: license add xxxxx xxxxx is the license code you received for the cluster feature. Note_____ If you are setting up a mirrored cluster, you should enable Syncmirror_local 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: reboot
3	Enable clustering on each node by entering the following command on the local node console: cf enable

Step	Action
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 “Stage 3: Setting options and parameters” on page 68.

Stage 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 Administration 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.copy.throttle

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 <i>option_name</i> <i>option_value</i>

What the parameters do

The parameters help keep each node 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

Parameter	Parameter
routed enabled	routed (on or off)
Time zone	timezone

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 by entering the following command on the console: 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 “ Stage 4: Configuring networking interfaces ” on page 72.

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 <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.

Step	Action
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 “ Stage 4: Configuring networking interfaces ” on page 72.

Stage 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. See the *System Administration Network Management Guide* for a description of available options and the function each performs

- 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.

Interface type	Description
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 Administration Guides</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 [n.n.n.n]: Please enter netmask [255.255.255.0]: Should interface ns0 take over a partner IP address during failover? [y]: Please enter the IP address to take over: [m.m.m.m]:
3	Enter the following command to reboot the node: reboot
4	Repeat Step 1 through Step 3 for the partner node.

**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 ffffffff00 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>

Step	Action	
2	If you are setting the interface as a...	Then...
	Dedicated interface	a. Enter the following command at the prompt, or directly into the /etc/rc file of the local node: <pre>ifconfig e0 local_node_ip_address</pre> b. Repeat for the partner node.
	Shared interface	a. Enter the following command at the prompt, or directly into the /etc/rc file of the local node: <pre>ifconfig NIC_name ip_address partner \ ip_address</pre> Example: <code>ifconfig e0 199.9.204.255 \ partner 199.9.204.254</code> b. Enter the following command at the prompt, or directly into the /etc/rc file of the partner node: <pre>ifconfig NIC_name ip_address partner \ ip_address</pre> Note _____ You do not have to enter the MAC address. _____
	Standby interface	a. Enter the following command at the prompt, or directly into the /etc/rc file of the local node: <pre>ifconfig e1 local_node_ip_address</pre> b. Repeat for the partner node.

Stage 5: Checking and testing cluster configuration

About this section

This section describes how to test your cluster configuration. The procedure for testing an F800 series cluster is different from that of the FAS900 series.

Topics in this section

This section contains the following information:

- ◆ [“Testing F800 series clusters”](#) on page 77
- ◆ [“Testing FAS900 series clusters”](#) on page 80

Testing F800 series clusters

Testing takeover and giveback

To test your configured cluster, complete the following steps.

Step	Action
1	Verify the network addresses by entering the following command from one node: ifconfig -a Result: The local node and partner node IP addresses are listed.
2	Verify that you can create and retrieve files on both nodes for each licensed protocol.
3	Enable the clustering licenses by entering the following command on the local and partner nodes: license add xxxxx xxxxx is the license code for your cluster. Note _____ If you are using a version of Data ONTAP earlier than 6.4, the command is <code>license option=code</code> . _____ Result: Clustering is enabled on your system.
4	From the same node, enter the following command: cf enable Result: Takeover of the partner node is enabled.
5	Run the cluster configuration checker (<code>cf-config-check.cgi</code>) script, as described in “ Setting options and parameters with the cluster configuration checker ” on page 70.

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

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

Testing FAS900 series clusters

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 tests 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 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 81.

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.	
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 is complete: giveback completed	
7	If takeover and giveback...	Then...
	Is completed successfully	Repeat Step 2 through Step 6 on the partner node.
	Fails	Go to Chapter 5, “ Basic Troubleshooting ,” on page 83.
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 add storage to 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 84
- ◆ [“Hot-swap disk loop error messages on Fabric Attached MetroClusters”](#) on page 87
- ◆ [“Recovering from a disaster”](#) on page 88
- ◆ [“Not enough storage”](#) on page 90
- ◆ [“LRC needs upgrading or LRC, ESH, or ESH2 fails”](#) on page 95
- ◆ [“Failed adapters”](#) on page 107
- ◆ [“Removing clustering”](#) on page 109
- ◆ [“Failover event cause-and-effect table”](#) on page 114

Takeover or giveback failures

About this section

The procedure for correcting a takeover or giveback failure on an F800 series cluster is different from that of a FAS900 series cluster. See one of the following sections according to your system hardware:

- ◆ [“Troubleshooting F800 series takeover or giveback failures”](#)
- ◆ [“Troubleshooting FAS900 series takeover or giveback failures”](#) on page 85

Troubleshooting F800 series takeover or giveback failures

If takeover or giveback fails on an F800 series cluster, complete the following steps.

Step	Action
1	Check the Channel B cabling on the disk shelves, and tighten any loose cables.
2	Check the cabling to your cluster adapters, making sure that the cabling is correct, that the cables are properly seated in the adapters, and that the cables are secured to the adapters.
3	Check communication between the nodes by entering the following command: cf status Result: If there is a cabling error or the interconnect link is down, the system displays an error message indicating the error.
4	Check the licenses on the local and partner nodes. Both nodes must have the cluster license enabled. Add missing licenses.
5	Check network connectivity. See the appropriate System Administration guide for more information.
6	Reboot the cluster and rerun the takeover and giveback tests.
7	If you still do not have clustering enabled, call Netapp technical support.

Troubleshooting FAS900 series takeover or giveback failures

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 or replace the cluster 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 is in the wrong slot number	Check the NVRAM5 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 Administration Guides</i> for more information.
3	Run the cluster configuration checker script, as described in “ Setting options and parameters with the cluster configuration checker ” on page 70. Correct any errors or differences displayed in the output.	

Step	Action
4	Reboot the cluster and rerun the takeover and giveback tests.
5	If you still do not have clustering enabled, call Netapp technical 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 Online Backup and Recovery Guide* to help you complete this procedure.

Step	Action
1	Isolate the node at the 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 Online Backup and Recovery 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	a. Restore access by replacing or reconnecting the damaged or missing cables. Caution _____ Do not force a takeover of the partner. b. Confirm access to the remote storage by entering the following command: sysconfig -v c. Confirm connection between the local and partner node by entering the following command: cf monitor d. Complete the disaster recovery procedures, as described in the <i>Data Protection Online Backup and Recovery Guide</i>.
	The remote node and storage are physically damaged or failed	a. Force the takeover of the partner node by entering the following command: cf forcetakeover -d b. Replace equipment and cabling, as needed. Contact Netapp technical support for assistance. c. Complete the disaster recovery procedures, as described in the <i>Data Protection Online Backup and Recovery Guide</i>.

Not enough storage

What to do when storage is full on FAS900 series systems

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 <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.

What to do when storage is full on F800 series systems

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

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

Considerations for adding storage to any 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...
FC7 or FC8 disk shelves	Adding FC9 disk shelves in the same loop	Disk shelf 0 must be an FC7 disk shelf.

If your cluster has...	And you are...	Then...
FC7, FC8, or FC9 disk shelves	Adding DS14/DS14mk2 FC disk shelves with LRCs	A DS14/DS14mk2 FC must be the last disk shelf in the loop.
DS14/DS14mk2 FC disk shelves with ESH or ESH2 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. ◆ A DS14/DS14mk2 FC must be the last disk shelf in the loop. ◆ DS14/DS14mk2 FC disk shelves with ESH or ESH2 modules can only be added to disk shelf loops with ESH or ESH2 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* 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 disk shelf counting is enabled by entering the following command: cf nfo status	
2	If disk shelf counting is...	Then...
	Enabled	Disable negotiated failover in your cluster by entering the following command from the node console: cf nfo disable disk_shelf
	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	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	
5	If disk shelves in the loop to which you are adding the new disk shelf have...	Then, for loops with...
	LRCs	◆ DS14 only: Go to Step 6. ◆ A mix of DS14 and DS14mk2 FC: Set the loop speed to 1 Gb. Caution_____ An incorrectly set loop speed results in an open loop condition. _____
	ESH or ESH2	◆ 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 will cause the appliance to panic. _____
6	Plug in the new disk shelf and turn it on.	
7	Cable the last disk shelf Channel A LRC, ESH, or ESH2 Output port to the new disk shelf Channel A LRC, ESH, or ESH2 Input port. Label the cable with a solid-colored label.	
8	Cable the last disk shelf Channel B LRC, ESH, or ESH2 Output port to the new disk shelf Channel B LRC, ESH, or ESH2 Input port. Label this cable with a striped label.	
9	Set both terminate switches (except for ESH2) on the new disk shelf to On.	

Step	Action	
10	Set both terminate switches (except for ESH2) on the previous disk shelf to Off.	
11	If disk shelf counting was...	Then...
	Enabled	a. Reenable negotiated failover in your cluster by entering the following command from the node console: cf nfo enable disk_shelf b. Use the cluster.
	Not enabled	Use the cluster.

LRC needs upgrading or LRC, ESH, or ESH2 fails

About the modules The LRC, ESH, or ESH2 module in a DS14/DS14mk2 FC includes a SCSI-3 Enclosure Services Processor. It maintains the integrity of the loop when disks are swapped and provides signal retiming for enhanced loop stability. There are two modules in the middle of the rear of the disk shelf, one for Channel A and one for Channel B.

Note_____

The Input and Output ports on module B on the DS14/DS14mk2 FC are inverted from module A.

Connectors in an LRC, ESH, or ESH2: The module has the following connectors.

Module connector	Function
Input	Provides the interface between the disk shelf and the storage appliance.
Output	Provides the interface between two disk shelves to create a loop of daisy-chained disk shelves.

For detailed information

This section provides information about the following topics:

- ◆ “[Upgrading LRC modules to ESH or ESH2 modules](#)” on page 96
- ◆ “[Hot-swapping an LRC, ESH, or ESH2](#)” on page 98

Upgrading LRC modules to ESH or ESH2 modules

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/ESH2 modules, you need to replace the LRC modules with ESH/ESH2 modules in the existing DS14 disk shelves.

This section describes how to upgrade your LRC modules to ESH/ESH2 modules. It also describes how to hot-swap a failed LRC or ESH/ESH2 module.

Upgrading a LRC module to an ESH/ESH2 module

To upgrade your DS14/DS14mk2 FC disk shelves with LRC modules into disk shelves with ESH/ESH2 modules, complete the following steps.

Step	Action	
1	If the cluster is a...	Then...
	Standard cluster, mirrored cluster, nonswitched MetroCluster, or switched MetroCluster	Take over the partner appliance's disk shelves by entering the following command on the local appliance: cf takeover
	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. Note _____ ESH2 is self-terminating and does not have a terminate switch. _____	
3	While grounded, unplug the cabling from the partner appliance's disk shelves. Note the cable locations, especially if your cluster is a Fabric Attached MetroCluster.	
4	Remove the Channel A loop LRCs of the partner's disk shelves and replace them with ESH/ESH2 modules.	

Step	Action	
5	Recable the disk shelves to their original locations before the ESH/ESH2 installation.	
6	Check the operation of the new ESH/ESH2 modules by entering the following command from the local node console: environ shelf Result: The local appliance reports the status of the modified shelves as good.	
7	If a cluster takeover...	Then...
	Was performed in Step 1	a. Return control of the partner appliance's disk shelves by entering the following command: cf giveback b. Repeat Step 2 through Step 6 for the Channel B loop.
	Was not performed in Step 1	Repeat Step 2 through Step 6 for the Channel B loop.
8	Repeat Step 1 through Step 7 to add storage to the partner appliance.	
9	Test the configuration, as described in " Stage 5: Checking and testing cluster configuration " on page 76.	

LRC needs upgrading or LRC, ESH, or ESH2 fails

Hot-swapping an LRC, ESH, or ESH2

Assumptions about this procedure This procedure assumes that the LRC, ESH, and ESH2 modules on the disk shelf are multipath connections to the appliance and that you are hot-swapping an LRC with another LRC, an ESH with another ESH, an ESH2 with another ESH2, or an ESH with an ESH2.

Note A hot-swap of an ESH with an ESH2 requires that you upgrade to Data ONTAP 6.4.4 or later and replace both ESH modules in the disk shelf.

Caution If you attempt to hot-swap the LRC, ESH, or ESH2 module on a disk shelf that does not have multipath connections, you lose all access to the drives on this disk shelf as well as others in the same disk pool.

Hot-swapping a module You can hot-swap an LRC, ESH, or ESH2 module in a cluster. To hot-swap a single module that belongs to the partner node, complete the following steps.

Step	Action
1	From the console of the appliance, enter the following command to disable the loop in which the failed module is a connection: <pre>cf takeover</pre> Note <code>cf takeover</code> is not needed in a MetroCluster configuration.
2	Verify that your appliance meets the minimum software requirements to support the disk shelf modules that you are hot-swapping. See the <i>DiskShelf14 and DiskShelf14mk2 FC Hardware Guide</i> for more information.
3	Put on the antistatic wrist strap and grounding leash.

Step	Action	
4	If you are hot-swapping...	Then...
	An LRC or ESH with like modules	Note the position of the terminate switch on the LRC or ESH you are removing and set the terminate switch on the replacement LRC or ESH to the same setting before inserting it into the disk shelf.
	An ESH with an ESH2 for a disk shelf that is either the first shelf in the loop or a shelf in the middle of the loop where all other disk shelves have ESHs	a. Place the terminate switch on all but the last disk shelf in the loop to the Off position on ESH A (up). b. Place the terminate switch on the last disk shelf in the loop to the On position on ESH A (down) and ESH B (up). Note _____ ESH2 is self-terminating and does not have a terminate switch. _____
	An ESH with an ESH2 for a disk shelf that is the last shelf in the loop, where all other disk shelves have ESHs	On all but the last disk shelf in the loop, place the terminate switch to the Off position on ESH A (up). Note _____ ESH2 modules are self-terminating, and therefore do not have a terminate switch. _____
	An ESH2 with an ESH2	Go to Step 5 .

Step	Action	
5	If you are hot-swapping...	Then...
	An LRC, ESH, or ESH2 with like modules	Disconnect the module that you are removing from the Fibre Channel cabling.
	An ESH with an ESH2	Disconnect the ESH from the Fibre Channel cabling and remove the cabling.
6	Using the thumb and index finger of both hands, press the levers on the CAM mechanism on the module to release it and pull it out of the disk shelf.	
7	Slide the replacement module into the slot at the rear of the disk shelf and push the levers of the CAM mechanism into place. Caution _____ Do not use excessive force when sliding the module into the disk shelf; you might damage the connector. _____	

Step	Action	
8	If you are hot-swapping...	Then...
	An LRC, ESH, or ESH2 with like modules	Reconnect the Fibre Channel cabling.
	An ESH with an ESH2 if it is the <i>first</i> shelf in the loop where all other disk shelves have ESHs	a. Reconnect the Fibre Channel cabling by using LC-to-LC cables when connecting to optical FC-AL adapters. b. Use HSSDC2-to-SFP cables for connections between the disk shelf with ESH2 and the disk shelf with ESH.
	An LRC, ESH, or ESH2 if it is a shelf in the <i>middle</i> of the loop, or the <i>last</i> shelf in the loop, where all other disk shelves have ESHs	Use HSSDC2-to-SFP cables for connections between a disk shelf with ESH and a disk shelf with ESH2.
9	From the console of the appliance, enter the following command to enable the loop in which the replacement module is a connection: cf giveback	
10	Repeat Step 3 through Step 9 for Loop B.	
11	Test the replacement module, as described in “ Testing the modules. ”	

Testing the modules

To test your LRC, ESH, or ESH2 modules, complete the following steps on both nodes.

Step	Action
1	Verify that all disk shelves with LRC, ESH, or ESH2 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.

LRC needs upgrading or LRC, ESH, or ESH2 fails

Hot-upgrading or hot-downgrading a disk shelf

When you can hot-upgrade or hot-downgrade

You can hot-upgrade a disk shelf with an LRC into a disk shelf with an ESH or ESH2, or hot-downgrade the disk shelf with an ESH or ESH2 into a disk shelf with an LRC, if your system incorporates redundancy through multipathing.

Note

MetroClusters are the only systems that use multipathing.

Caution

You must hot-upgrade or hot-downgrade all disk shelves in the loop.

Requirements for hot-upgrading or hot-downgrading

To hot-upgrade a disk shelf or hot-downgrade a disk shelf, the following requirements must be met.

Disk shelf loop speed	Hot-upgrade	Hot-downgrade
DS14 (only available at 1-Gb loop speed)	LRC to ESH or ESH2 hot-upgrade is supported.	ESH or ESH2 to LRC hot-downgrade is supported.
DS14mk2 FC set at 1-Gb loop speed	LRC to ESH or ESH2 hot-upgrade is supported.	ESH or ESH2 to LRC hot-downgrade is supported.
DS14mk2 FC set at 2-Gb loop speed	◆ Not applicable ◆ ESH to ESH2 is supported. See “Hot-swapping a module” on page 98.	◆ ESH or ESH2 to LRC hot-downgrade is not supported. ◆ ESH2 to ESH hot-downgrade is supported.

Hot-upgrading or hot-downgrading a disk shelf

To hot-upgrade or hot-downgrade a disk shelf, complete the following steps.

Step	Action	
1	Put on the antistatic wrist strap and grounding leash.	
2	If you are working with...	Then...
	A standard cluster system	From the console of the partner appliance, enter the following command to disable all the LRC, ESH, or ESH2 modules in Loop A: cf takeover
	A MetroCluster system	Go to Step 3 .
3	If you are...	Then...
	Hot-upgrading the disk shelf	Note the position of the terminate switch on the LRC you are removing and set the terminate switch on the ESH to the same setting before inserting it into the disk shelf. Note _____ ESH2 is self-terminating and does not have a terminate switch. _____
	Hot-downgrading the disk shelf from an ESH to an LRC	Note the position of the terminate switch on the ESH you are removing and set the terminate switch on the LRC to the same setting before inserting it into the disk shelf.
	Hot-downgrading the disk shelf from an ESH2 to an ESH or LRC	On all but the last disk shelf in the loop, place the terminate switch to the Off position on ESH/LRC A (up) and ESH/LRC B (down).
4	Disconnect the LRC, ESH, or ESH2 from the Fibre Channel cabling and remove the cabling.	

Step	Action	
5	Using the thumb and index finger of both hands, press the levers on the CAM mechanism on the module to release it and pull it out of the disk shelf.	
6	Slide the module into the slot at the rear of the disk shelf and push the levers of the CAM mechanism into place. Caution ————— Do not use excessive force when sliding the module into the disk shelf; you might damage the connector.	
7	If you are...	Then...
	Hot-upgrading a disk shelf from an LRC to an ESH	Reconnect the Fibre Channel cabling by using LC-to-LC cables when connecting to optical FC-AL adapters, and with HSSDC2-to-HSSDC2 cables for connections between shelves.
	Hot-upgrading a disk shelf from an LRC to an ESH2	Reconnect the Fibre Channel cabling by plugging an optical SFP into the ESH2 Input port and using an LC-to-LC cable when connecting to optical FC-AL adapters and with SFP-to-SFP cables for connections between shelves.
	Hot-downgrading a disk shelf from an ESH2/ESH to an LRC	Reconnect the Fibre Channel cabling by using HSSDC to DB-9 cables when connecting to copper FC-AL adapters or SC-to-LC cables when connecting to optical FC-AL adapters, and with DB-9 to DB-9 cables for connections between shelves.
	Hot-downgrading a disk shelf from an ESH2 to an ESH	Reconnect the Fibre Channel cabling by using LC-to-LC cables when connecting to optical FC-AL adapters, and with HSSDC2-to-HSSDC2 cables for connections between shelves.

Step	Action
8	Repeat Step 2 through Step 7 for the remaining LRC, ESH, or ESH2 modules in the loop.
9	From the console of the appliance, enter the following command to verify that the modules in Loop A report good status: environ shelf As long as only the modules in loop A are upgraded, you also receive the following error message: Element 1: component is from a different product family. This is an unsupported configuration. Note_____ Both modules in the disk shelf must be the same. A mixed module configuration is supported during the hot-swap procedure, but it is not a supported configuration for long-term operation. _____
10	From the console of the appliance, enter the following command to enable the Loop A: cf giveback
11	Repeat Step 1 through Step 10 for Loop B.

Failed adapters

Tasks for replacing an adapter

Replacing an adapter consists of the following tasks:

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

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> .

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.
3	Plug the node back into the power source, and then apply power.

Step	Action
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
 - ❖ “[Removing F800 series clustering hardware](#)” on page 110
 - ❖ “[Removing FAS900 series clustering hardware](#)” on page 112

Disabling clustering software

This procedure applies to F800 series and FAS900 series systems. To disable the clustering software, complete the following steps.

Step	Action
1	Enter the following command on either node console: cf disable
2	Disable the cluster license by entering the following command: license delete cluster Note _____ If you are using a version of Data ONTAP earlier than 6.4, the command is <code>license cluster=disable</code> . _____
3	Open the <code>/etc/rc</code> file with a text editor and remove references to the partner node in the <code>ifconfig</code> 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 Step 1 through Step 3 on the partner node.
5	Proceed to the appropriate following section based on your hardware.

Removing F800 series clustering hardware

Tasks for removing the hardware

Removing F800 series clustering hardware is a two-part process:

- ◆ Recabling your disk shelves
- ◆ Removing the hardware

Recabling Channel B for redundant disk loops

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

Note

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

Step	Action
1	Halt both filers by entering the following command at each console: halt
2	Ground yourself, then power off the nodes and disk shelves.
3	Unplug the local node from the partner node's Channel B disk shelf and do the same for the partner node.
4	Connect the local node to the open Channel B loop in its local disk shelves.
5	Repeat Step 4 for the partner node.
6	Go to " Removing the hardware " on page 111.

Removing the hardware

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

Step	Action
1	Ground yourself, using the antistatic wrist strap.
2	Remove the interconnect cables from the cluster adapter.
3	Open the back of the local node, as described in the <i>F800 Series Hardware Guide</i> .
4	If applicable, unlock and remove the GBIC from the adapter (typically used with VI cluster adapters).
5	Unscrew the adapter from the chassis.
6	Remove the adapter.
7	Replace the slot cover and secure it with the screw from the adapter.
8	Remove any extra FC-AL or dual Fibre Channel adapters.
9	Close and secure the back of the filer, as described in the <i>F800 Series Hardware Guide</i> .
10	Power on first the disk shelves and then the filer; press the Delete key during boot to get to the OK prompt.
11	Unset the partner system ID by entering the following command at the OK prompt: unsetenv partner-sysid
12	Boot the filer by entering the following command at the OK prompt: boot
13	Check the status of clustering on the filer by entering the following command on the console: cf status Result: Cluster monitoring is disabled, as indicated by the following message: Cluster monitor not initialized
14	Repeat Step 1 through Step 13 for the partner node.

Removing FAS900 series clustering hardware

Removing FAS900 series clustering hardware

To remove FAS900 series 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. See the <i>Hardware and Service Guide</i> for additional information.	
4	If you are converting a cluster to a stand-alone appliance for a...	Then...
	FAS920, 940, or 960 with IB adapter	Open both nodes, remove the cluster adapters, and then close the nodes.
	FAS900 series standard cluster with NVRAM5 adapter	Remove the NVRAM5 adapter from slot 11 and install it in slot 10.
	FAS980 Fabric Attached MetroCluster with NVRAM5 adapter	Completely remove the VI cluster adapter from slot 6.
5	Unplug the local node from the partner node’s Channel B disk shelf and do the same for the partner node.	

Step	Action	
6	If you are converting to a...	Then...
	System with nonmirrored disks	a. Disconnect all cabling from the Channel B loop on the local node. b. Repeat for the partner node.
	System with mirrored disks or a redundant Channel B loop	a. Connect the local node to the open Channel B loop in its local disk shelves, as described in the appropriate disk shelf guide. b. Repeat 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.	
8	Run all system diagnostics at the boot prompt by entering the following command on the system console: boot diags	
9	Unset the partner system ID 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.	

Failover event cause-and-effect table

When to use this table

Use this table to understand what happens when a failover event occurs, and how the various clustering configurations handle these events.

Event	Does the event trigger a failover?		Does the event prevent a future failover from occurring, or a failover from occurring successfully?		Is data still available on affected volume after event?		
	Cluster	Metro-Cluster	Cluster	Metro-Cluster	Single filer	Cluster	Metro-Cluster
Single disk failure	No	No	No	No	Yes	Yes	Yes
Double disk failure (2 disks fail in same RAID group)	Yes, unless you are using Sync-Mirror, then no.	No	Maybe. If root volume has double disk failure or if the mailbox disks are affected, then no CFO is possible.	No	No, unless you are using raid/dp or Sync-Mirror, then yes.	No, unless you are using raid/dp or Sync-Mirror, then yes.	Yes, with no CFO necessary.
Triple disk failure (3 disks fail in same RAID group)	Maybe. If Sync-Mirror is being used, there won't be a takeover. Otherwise, yes.	No	Maybe. If root volume has triple disk failure, no CFO is possible.	No	No	No	Yes, with no CFO necessary.

Event	Does the event trigger a failover?		Does the event prevent a future failover from occurring, or a failover from occurring successfully?		Is data still available on affected volume after event?		
	Cluster	Metro-Cluster	Cluster	Metro-Cluster	Single filer	Cluster	Metro-Cluster
Single HBA (initiator) failure, Loop A	Maybe. If Sync-Mirror is being used, then no. Otherwise, yes.	No	Maybe. If root volume has double disk failure, no CFO is possible.	No	Yes, if shelf is dual-pathed, or if Sync-Mirror is being used.	Yes, if Sync-Mirror is being used, or if CFO succeeds.	Yes, with no CFO necessary.
Single HBA (initiator) failure, Loop B	No	No	Yes, unless you are using Sync-Mirror and the mailbox disks aren't affected, then no.	No	Yes, if shelf is dual-pathed or Sync-Mirrored.	Yes, With no CFO necessary if Sync-Mirror was used.	Yes, with no CFO necessary.
Single HBA initiator failure, (both loops at same time)	Yes, unless the data is Sync-Mirrored on a different (up) loop, then no takeover needed.	No	Maybe. If the data is SyncMirrored and the mailbox disks aren't affected, then no. Otherwise, yes.	No	No, unless the data was Sync-Mirrored, then yes.	No CFO needed if data was Sync-Mirrored.	Yes, with no CFO necessary.

Event	Does the event trigger a failover?		Does the event prevent a future failover from occurring, or a failover from occurring successfully?		Is data still available on affected volume after event?		
	Cluster	Metro-Cluster	Cluster	Metro-Cluster	Single filer	Cluster	Metro-Cluster
LRC failure (Loop A)	Only if multi-disk volume failure or open loop condition occurs.	No	Maybe. If root volume has double disk failure, no CFO is possible since this impacts the mailbox disks.	No	Yes, if shelf is dual-pathed or Sync-Mirrored.	Yes, if CFO succeeds or if data is Sync-Mirrored.	Yes, with no CFO necessary.
LRC failure (Loop B)	No	No	Maybe. If the data is SyncMirrored and the mailbox disks aren't affected, then automatic CFO will happen.	No	Yes, if shelf is dual-pathed or the data is Sync-Mirrored.	Yes	Yes
ESH failure (Loop A)	Only if multi-disk volume failure or open loop condition occurs, and data isn't Sync-Mirrored.	No	Maybe. If root volume has double disk failure, no CFO is possible.	No	No	Yes, if CFO succeeds.	Yes, with no CFO necessary.

Event	Does the event trigger a failover?		Does the event prevent a future failover from occurring, or a failover from occurring successfully?		Is data still available on affected volume after event?		
	Cluster	Metro-Cluster	Cluster	Metro-Cluster	Single filer	Cluster	Metro-Cluster
ESH failure (Loop B)	No	No	Maybe. If SyncMirror is used, then no. Otherwise, yes.	No	Yes, if shelf is dual-pathed or data is Sync-Mirrored.	Yes	Yes
Shelf (backplane) failure	Only if multi-disk volume failure or open loop condition occurs, and data is Sync-Mirrored.	No	Maybe. If root volume has double disk failure or if the mailboxes are affected, no CFO is possible.	No	Maybe. If data is Sync-Mirrored, then yes. If not, then no.	Maybe. If data is Sync-Mirrored, then yes. If not, then no.	Yes, with no CFO necessary.
Shelf, single power failure	No	No	No	No	Yes	Yes	Yes

Event	Does the event trigger a failover?		Does the event prevent a future failover from occurring, or a failover from occurring successfully?		Is data still available on affected volume after event?		
	Cluster	Metro-Cluster	Cluster	Metro-Cluster	Single filer	Cluster	Metro-Cluster
Shelf, dual power failure	Only if multidisk volume failure or open loop condition occurs and data isn't Sync-Mirrored.	No	Maybe. If root volume has double disk failure or if the mailbox disks are affected, no CFO is possible.	No	Maybe. If data is Sync-Mirrored, then yes. Otherwise no.	Maybe. If data is Sync-Mirrored, then yes. Otherwise no.	Yes, with no CFO necessary.
Head, single power failure	No	No	No	No	Yes	Yes	Yes
Head, dual power failure	Yes	Yes	Yes, until power is restored.	Yes, until power is restored	No	Yes, if CFO succeeds.	Yes, if CFO is successful
Cluster interconnect failure (1 port)	No	No	No	No	n/a	Yes	Yes
Cluster interconnect failure (both ports)	No	No	Yes	No, CFOD is possible	n/a	Yes	Yes

Event	Does the event trigger a failover?		Does the event prevent a future failover from occurring, or a failover from occurring successfully?		Is data still available on affected volume after event?		
	Cluster	Metro-Cluster	Cluster	Metro-Cluster	Single filer	Cluster	Metro-Cluster
Ethernet interface failure (primary, no VIF)	No	No	No	No	No	No	No
Ethernet interface failure (primary, VIF)	No	No	No	No	Yes	Yes	Yes
Ethernet interface failure (secondary, VIF)	No	No	No	No	Yes	Yes	Yes
Ethernet interface failure (VIF, all ports)	No	No	No	No	No	No	No
Tape interface failure	No	No	No	No	Yes	Yes	Yes
Heat above +/- foo degrees	Yes	Yes	No	No	No	No	No

Event	Does the event trigger a failover?		Does the event prevent a future failover from occurring, or a failover from occurring successfully?		Is data still available on affected volume after event?		
	Cluster	Metro-Cluster	Cluster	Metro-Cluster	Single filer	Cluster	Metro-Cluster
Fan failures (disk shelves, filer, or storage engine)	No	No	No	No	Yes	Yes	Yes
Reboot/panic	No	No	No	No	Maybe. Depends on cause of panic.	Maybe. Depends on cause of panic.	Maybe. Depends on cause of panic.

About this chapter

This chapter discusses how to install a DS14/DS14mk2 FC disk shelf into an existing cluster system.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Installing a disk shelf in an existing mixed-shelf loop”](#) on page 122
- ◆ [“Hot-adding a DS14/DS14mk2 FC to an existing loop”](#) on page 126

Installing a disk shelf in an existing mixed-shelf loop

Assumptions about the existing mixed-shelf loop

This section is based on the assumption that the last disk shelf in the existing mixed-shelf loop connected to your appliance is not a DS14 or DS14mk2 FC disk shelf.

Note—————
If the last disk shelf in your storage appliance system is a DS14/DS14mk2 FC, you can add the disk shelf to the loop without shutting down the system. Go to [“Hot-adding a DS14/DS14mk2 FC to an existing loop”](#) on page 126.


Caution—————
Do not attempt to add a DS14/DS14mk2 FC disk shelf that has ESH or ESH2 modules to a mixed-shelf loop.

Installing the disk shelf

To install a DS14/DS14mk2 FC disk shelf into an existing mixed-shelf loop on an appliance, complete the following steps.

Caution—————
If you are installing the DS14/DS14mk2 FC disk shelf into a loop that contains FC7, FC8, or FC9 disk shelves, the DS14/DS14mk2 FC must be the last disk shelf in the loop and must have the highest disk shelf ID in the loop to avoid address conflicts.

Step	Action
1	Put on the antistatic wrist strap and grounding leash.
2	Shut down the appliance by entering the following command at the console: halt Caution ————— Always use the <code>halt</code> command to perform a clean shutdown.

Step	Action	
3	Verify that the LCD display at the front of your appliance displays the following message: Halted	
4	Turn off the power to your appliance and the disk shelves.	
5	Install the disk shelf, attach grounding cables, and connect the disk shelf to a power source. See the <i>DiskShelf14 and DiskShelf14mk2 FC Hardware Guide</i> for more information.	
6	If the last disk shelf in the loop is...	Then...
	An FC7 or FC8 disk shelf	Remove the terminator from the Output/Terminator connector and go to Step 7 .
	An FC9 disk shelf	Go to Step 7 .
7	Connect one end of the provided cable to the Output/Terminator connector of the LRC A on the last disk shelf in the loop.	
8	Connect the other end of the cable to the LRC A Input on the disk shelf that you are adding. The following illustration shows a DS14mk2 FC disk shelf installed above an FC9 disk shelf. 	

Step	Action	
9	If the DS14/ DS14mk2 FC...	Then...
	Is the only disk shelf being added to the loop	a. Connect one end of the provided cable to the LRC B Input on the disk shelf. b. Connect the other end of the cable to the LRC B Output on the disk shelf. c. Turn the terminate switches on LRCs A and B to the On position. The diagram shows a rectangular switch component. The bottom half of the rectangle is filled with black, and a downward-pointing arrow is positioned to the right of the rectangle, indicating the 'On' position. LRC A active terminate position Note Because LRC B on the disk shelf is inverted, the On termination position on LRC B is up instead of down.
	Is the first of additional DS14/DS14mk2 FC disk shelves being added to the loop	a. Turn the terminate switch on the first disk shelf in the mixed-shelf loop to the Off position for LRC A (up) and LRC B (down). b. Connect one end of the provided cable to the LRC B Output on the first disk shelf in the mixed-shelf loop. c. Connect the other end of the cable to the LRC B Input on the next disk shelf that you added to the mixed-shelf loop. d. Verify that the terminate switch on all but the last disk shelf is in the Off position for LRC A (up) and LRC B (down). The terminate switches on the last disk shelf should be in the On position.

Step	Action	
10	Give the disk shelf a unique disk shelf ID. See the <i>DiskShelf14</i> and <i>DiskShelf14mk2 FC Hardware Guide</i> for additional information. Note_____ A valid shelf ID is 1 through 7. Only use disk shelf ID 7 if it is the last ID available on this loop. A shelf ID that is not from 1 through 7 causes the drive addresses to default to those of a shelf with the ID switch set to 7, even though the Shelf ID indicator in the front operation panel displays a dash (-). _____	
11	If your disk shelf is a...	Then...
	DS14mk2 FC with an LRC module	Set the loop speed switch to 1 Gb. Caution_____ An incorrectly set loop speed results in an open loop condition. _____
	DS14	Go to Step 12 .
12	Turn on the power first to the disk shelves and then to the appliance.	
13	Reboot the appliance.	

Hot-adding a DS14/DS14mk2 FC to an existing loop

About this section This section provides information about how to hot-add a DS14 or DS14mk2 FC disk shelf to an existing loop. It also tabulates the error messages that appear on your appliance console if the attempt at hot-adding was unsuccessful.

Note_____The hot-added disk shelf must have new drives that will be used as spares only.

Requirements for hot-adding To hot-add a DS14/DS14mk2 FC to an existing loop, the following requirements must be met.

If the disk shelf being hot-added...	Then...
Has LRC modules	◆ The disk shelves in the existing loop must have a DS14/DS14mk2 FC disk shelf with LRCs as the last shelf. ◆ You must follow the procedure “Hot-adding a disk shelf to an existing loop” on page 127 exactly, if you are hot-adding to a mixed-shelf loop. Caution_____Failure to follow the procedure exactly causes the last FC9 in the loop to automatically terminate the loop. The appliance loses contact with the existing last DS14/DS14mk2 FC in the loop and crashes.
Has ESH or ESH2 modules	The disk shelves in the existing loop must have only DS14 or DS14mk2 FC disk shelves with ESH or ESH2 modules in the loop.

Hot-adding a disk shelf to an existing loop

To hot-add a disk shelf to an existing loop, complete the following steps.

Step	Action
1	Put on the antistatic wrist strap and grounding leash.
2	Install the disk shelf. See the <i>DiskShelf14 and DiskShelf14mk2 FC Hardware Guide</i> for more information.
3	Ground the disk shelf using the provided grounding cable. See the <i>DiskShelf14 and DiskShelf14mk2 FC Hardware Guide</i> for more information.
4	Connect the disk shelf to a power source. See the <i>DiskShelf14 and DiskShelf14mk2 FC Hardware Guide</i> for more information.
5	Give the disk shelf a unique disk shelf ID. See the <i>DiskShelf14 and DiskShelf14mk2 FC Hardware Guide</i> for more information. Note _____ Only use disk shelf ID 7 if it is the last ID available. _____
6	Verify that the disk shelf ID is not being used in the loop by entering the following command at the console: fcstat device_map adaptername Caution _____ An invalid disk shelf ID causes the appliance to panic. _____

Step	Action	
7	If disk shelves in the loop have...	Then, for loops with ...
	LRC modules	◆ DS14 only: Go to Step 8. ◆ 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.
	ESH or ESH2 modules	◆ 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 appliance to panic.
8	If the disk shelf you added has...	Then...
	ESH2 modules	Turn on the power to the disk shelf and wait 30 seconds for the shelf electronics to finish initializing.
	LRC or ESH modules	Turn on the power to the disk shelf. Caution _____ You must power up the new disk shelf before you connect it to the last disk shelf in the loop because there must be approximately a 120-second time difference between Step 8 (powering on the disk shelf) and Step 14 (turning off the terminate switch) of this procedure.

Step	Action	
9	Connect one end of the provided cable to the LRC, ESH, or ESH2 Channel A Output on the last disk shelf in the existing loop.	
10	Connect the other end of the cable to the LRC, ESH, or ESH2 Channel A Input on the new disk shelf.	
11	Connect one end of the provided cable to the LRC, ESH, or ESH2 Channel B Output on the last disk shelf in the existing loop.	
12	Connect the other end of the cable to the LRC, ESH, or ESH2 Channel B Input on the new disk shelf.	
13	Verify that all the cables are securely fastened. Caution _____ Poorly secured cables cause the appliance to panic over an open loop.	
14	If the disk shelf you added has...	Then...
	ESH2 modules	Result: In 60 seconds, the appliance recognizes the hot-added disk shelf.
	LRC or ESH modules	a. Turn the terminate switch of the newly added disk shelf to the On position for LRC or ESH Channel A (down) and LRC or ESH Channel B (up). Caution _____ Failure to place the terminate switches in the On position causes the appliance to panic over an open loop. b. Turn the terminate switch of the now second-to-last disk shelf to the Off position for LRC or ESH Channel A (up) and LRC or ESH Channel B (down). Result: In 60 seconds, the appliance recognizes the hot-added disk shelf.

Error messages

The following error messages appear on your appliance console if your attempt at hot-adding the DS14/DS14mk2 FC is unsuccessful.

Error message	Explanation
Open loop panic	One of three reasons cause this error message to appear: ◆ The terminate switch on LRC or ESH Channel A and LRC or ESH Channel B on the new disk shelf are not in the On position. ◆ The shelf-to-shelf cable between the now second-to-last disk shelf and the newly added disk shelf is defective or is not securely fastened. ◆ The speed of the newly added DS14mk2 FC disk shelf is incorrectly set.
Soft address panic	One of two reasons cause this error message to appear: ◆ There is an invalid disk shelf ID. ◆ The power was turned on before the disk shelf ID was changed and the disk shelf was not power-cycled after the disk shelf ID was changed.
Speed mismatch termination	The ESH modules on the disk shelf detected a speed mismatch between the preceding disk shelf and this disk shelf and is reporting them as automatic termination errors.

About this appendix This appendix lists the regulatory notices you need to be aware of when installing and operating your NetApp 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)

警告使用者：
這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

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)

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラスA情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

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 introduction
Network appliance	NetApp FAS900 series appliances	FAS920, FAS940, FAS960, FAS980	2003, 2004
	NetApp F810 filer	F810	2001
	NetApp F820 filer	F820	2000
	NetApp F825 filer	F825	2002
	NetApp F840 filer	F840	2000
	NetApp F880 filer	F880	2001

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

Feature update history

The following table lists and describes the history of changes made to this manual. When a change is implemented, it applies to the release in which it was implemented and all subsequent releases, unless otherwise specified.

Feature updates	Feature first implemented in	Feature release date
◆ Updates for FAS920 ◆ Update for NVRAM5 ◆ Illustration updates	Data ONTAP 6.5.1	May 2004
◆ Updates for NVRAM5 support in FAS900 series clusters, except for MetroCluster ◆ Failover event cause-and-effect table ◆ Declaration of Conformity update ◆ Addition of clustered failover and single-point-of-failure analysis	Data ONTAP 7.0	November 2004

Index

A

adapters

- cabling F800 series cluster 42
- cabling FAS900 series cluster 33
- dual-port FC-AL 23
- installing 107
- NVRAM5 7, 33
- removing 107

adding

- disk shelf to mirrored cluster 90
- disk shelf to standard cluster 92

audience vii

B

bring up 68

- configuring interfaces for 72
- enabling licenses for 66
- first time 64
- manually setting options for 69
- manually setting parameters for 70

C

cabling

- Channel A loop in mirrored FAS900 cluster 27
- Channel A loop in standard F800 cluster 36
- Channel A loop in standard FAS900 cluster 25
- Channel B loop in mirrored FAS900 cluster 31
- Channel B loop in standard F800 cluster 39
- Channel B loop in standard FAS900 cluster 29
- dual-port FC-AL 23
- FC-ALs in mirrored cluster 23
- location of FAS900 dual-port FC-ALs 23
- preparing equipment racks for 22
- preparing system cabinets for 21
- required 15

cluster adapter

- cabling for F800 series 42
- cabling for FAS900 series 33

cluster types

installed in equipment racks 20

installed in system cabinets 20

standard 7

clustered failover 2

clusters

- adding a disk shelf to mirrored 90
- adding a disk shelf to standard 92
- checking configuration with the configurator script 70
- configuring interfaces for 72
- disabling 109
- first-time bring up of 64
- installing FC-AL adapters in 90
- removing hardware from 110, 112
- setting options in 68
- setting parameters in 69
- testing 80

commands

- cf enable (enables clustering) 77, 81
- cf giveback (enables giveback) 79, 81
- cf takeover (enables takeover) 79, 81
- cfgdisable (disables configuration) 47
- cfgsave (saves configuration) 47, 48
- configshow (shows switch settings) 48
- configure (switch parameters) 47
- license add xxxxx (license cluster) 58, 66
- switchdisable (disables switch) 46
- switchenable (enables switch) 48
- switchname (sets switch name) 48
- version (displays switch firmware) 46

configuring

- dedicated interfaces 75
- interfaces 73
- shared interfaces 75
- standby interfaces 75
- using the Cluster Configuration Checker 77

conventions

- command viii
- formatting ix
- keyboard viii

D

Data ONTAP

- in a Fabric Attached MetroCluster 12
- in a mirrored cluster 9
- in a standard cluster 7

definitions

- appliance viii
- disk shelf viii
- F800 vii
- Fabric Attached MetroCluster vii
- FAS900 vii
- FC HBA viii
- FC HBA for Disk viii
- filer viii
- mirrored cluster vii
- nonswitched MetroCluster vii
- switch fabric 11

E

eliminating filer single-point-of-failure 3

equipment racks

- installing in 20
- preparation of 22

F

Fabric Attached MetroCluster

- advantages of 12
- behavior of Data ONTAP with 12
- cabling the local node and disk shelves in 49
- cabling the remote node and disk shelves in 54
- configuring the switches in 46
- planning worksheet example 45
- restrictions 13

failover cause-and-effect table 114

G

GBIC, removing 111

giveback

- testing F800 series 77
- testing FAS900 series 81

H

hardware

- installing an adapter 107
- removing an adapter 107
- removing F800 series clustering 110
- removing FAS900 series clustering 112

I

installing, adapters 107

interface configurations

- dedicated 72
- shared 72
- standby 73

L

licenses

- disabling cluster 109
- enabling 66

M

manuals, required 14

MetroCluster, converting to a Fabric Attached 61

mirrored clusters

- adding a disk shelf to 90
- advantages of 9
- behavior of Data ONTAP with 9
- cabling Channel A loop for FAS900 series in 27
- cabling Channel B loop for FAS900 series in 31
- contents 9
- FC-ALs in 23
- restrictions 10

O

options, manually setting 69

P

parameters, manually setting 70

preparing equipment racks 22

R

racks, installing in equipment 20

requirements

- cable 15

- equipment 15

- manual 14

- tool 15

restrictions

- Fabric Attached MetroCluster 13

- in mirrored clusters 10

- in standard clusters 8

S

setting options and parameters 68

single-point-of-failure 2

software

- disabling clustering 109

special messages ix

SPOF 2

standard cluster

- adding a disk shelf to 92

- advantages of 7

- behavior of Data ONTAP with 7

- cabling Channel A loop for F800 series in 36

- cabling Channel A loop for FAS900 series in 25

- cabling Channel B loop for F800 series in 39

- cabling Channel B loop for FAS900 series in 29

- contents 7

- restrictions 8

switch fabric 11

switches

- cabling the local node and disk shelves with 49

- cabling the remote node and disk shelves with 54

- checking firmware revision for 46

- configuring for Fabric Attached MetroCluster 46

- disabling 46

- setting the parameters 47

system cabinets

- installing in 20

- preparing for cabling 21

T

takeover

- testing F800 series 77

- testing FAS900 series 81

terminology vii

tools, required 15

troubleshooting

- disaster recovery 88

- giveback failure on F800 series 84

- giveback failure on FAS900 series 85

- takeover failure on F800 series 84

- takeover failure on FAS900 series 85

V

VI cluster adapter

- removing 111

