

# Data ONTAP® 7.3

## **Active/Active Configuration Guide**

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

Part number 210-04502\_A0

**Updated for Data ONTAP 7.3.2 on 17 August 2009**



# Contents

<b>Copyright information.....</b>	<b>9</b>
<b>Trademark information.....</b>	<b>11</b>
<b>About this guide.....</b>	<b>13</b>
Audience.....	13
Accessing Data ONTAP man pages.....	14
Terminology.....	15
Where to enter commands.....	16
Keyboard and formatting conventions.....	16
Special messages.....	17
How to send your comments.....	18
<b>Active/active configuration types and requirements.....</b>	<b>19</b>
Overview of active/active configurations.....	19
What an active/active configuration is.....	19
Benefits of HA.....	20
Characteristics of nodes in an active/active configuration.....	20
Best practices for deploying an active/active configuration.....	21
Comparison of active/active configuration types.....	22
Standard active/active configurations.....	23
How Data ONTAP works with standard active/active configurations.....	23
Standard active/active configuration diagram.....	24
Setup requirements and restrictions for standard active/active configurations.....	24
Configuration variations for standard active/active configurations.....	25
Understanding mirrored active/active configurations.....	26
Advantages of mirrored active/active configurations.....	26
Setup requirements and restrictions for mirrored active/active configurations.....	27
Configuration variations for mirrored active/active configurations.....	28
Understanding stretch MetroClusters.....	28
Continued data service after loss of one node with MetroCluster.....	29
Advantages of stretch MetroCluster configurations.....	29
Stretch MetroCluster configuration.....	29

Stretch MetroCluster configuration on 31xx systems .....	31
How Data ONTAP works with stretch MetroCluster configurations.....	31
Stretch MetroCluster and disk ownership.....	31
Setup requirements and restrictions for stretch MetroCluster configurations.....	32
Configuration variations for stretch MetroCluster configurations.....	32
MetroClusters and SnapLock volumes.....	33
Understanding fabric-attached MetroClusters.....	33
Fabric-attached MetroClusters use Brocade Fibre Channel switches.....	34
Advantages of fabric-attached MetroCluster configurations.....	34
Fabric-attached MetroCluster configuration.....	35
Fabric-attached MetroCluster configuration on 31xx systems.....	35
How Data ONTAP works with fabric-attached MetroCluster configurations.....	36
Setup requirements and restrictions for fabric-attached MetroClusters.....	36
Configuration limitations for fabric-attached MetroClusters.....	38
Configuration variations for fabric-attached MetroClusters.....	38
MetroClusters and SnapLock volumes.....	39
<b>Active/active configuration installation.....</b>	<b>41</b>
System cabinet or equipment rack installation.....	41
Active/active configurations in an equipment rack.....	41
Active/active configurations in a system cabinet.....	42
Required documentation, tools, and equipment.....	42
Required documentation.....	42
Required tools.....	43
Required equipment.....	44
Preparing your equipment.....	44
Installing the nodes in equipment racks.....	45
Installing the nodes in a system cabinet.....	45
Cabling nodes and DS14mk2 AT, DS14mk2 FC, or DS14mk4 FC disk shelves in standard or mirrored active/active configurations.....	46
Systems with two controllers in the same chassis.....	46
Which Fibre Channel ports to use for DS14mk2 or DS14mk4 FC disk shelf connections.....	46
Cabling a standard active/active configuration.....	47
Cabling a mirrored active/active configuration.....	50

Required connections for using uninterruptible power supplies with standard or mirrored active/active configurations.....	57
<b>MetroCluster installation.....</b>	<b>59</b>
Required documentation, tools, and equipment.....	59
Required documentation.....	59
Required tools.....	61
Required equipment.....	61
MetroCluster and software-based disk ownership.....	62
Converting an active/active configuration to a fabric-attached MetroCluster.....	63
Upgrading an existing MetroCluster.....	65
Cabling a stretch MetroCluster.....	67
Cabling a stretch MetroCluster between 31xx systems.....	67
Cabling a fabric-attached MetroCluster.....	68
Planning the fabric-attached MetroCluster installation.....	69
Configuration differences for fabric-attached MetroClusters on 31xx systems.....	70
Configuring the switches.....	70
Cabling Node A.....	72
Cabling Node B.....	81
Assigning disk pools (if you have software-based disk ownership).....	91
Verifying disk paths.....	92
Required connections for using uninterruptible power supplies with MetroCluster configurations.....	93
<b>Reconfiguring an active/active configuration into two stand-alone systems.....</b>	<b>95</b>
Ensure uniform disk ownership within disk shelves and loops in the system.....	95
Disabling the active/active software.....	96
Reconfiguring nodes using disk shelves for stand-alone operation.....	97
Requirements when changing an node using array LUNs to stand-alone.....	99
Reconfiguring nodes using array LUNs for stand-alone operation.....	99
<b>Configuring an active/active configuration.....</b>	<b>103</b>
Bringing up the active/active configuration.....	103
Considerations for active/active configuration setup.....	103
Configuring shared interfaces with setup.....	104
Configuring dedicated interfaces with setup.....	105

Configuring standby interfaces with setup.....	105
Enabling licenses.....	106
Setting options and parameters.....	107
Option types for active/active configurations.....	107
Setting matching node options.....	107
Parameters that must be the same on each node.....	108
Disabling the change_fsid option in MetroCluster configurations.....	108
Configuration of the hw_assist option.....	110
Configuration of network interfaces.....	112
What the networking interfaces do.....	113
IPv6 considerations in an active/active configuration.....	113
Configuring network interfaces for active/active configurations.....	114
Configuring partner addresses on different subnets (MetroClusters only).....	119
Testing takeover and giveback.....	123
<b>Management of takeover and giveback.....</b>	<b>125</b>
How takeover and giveback work.....	125
When takeovers occur.....	125
What happens during takeover.....	126
What happens after takeover.....	126
What happens during giveback.....	127
Management of an active/active configuration in normal mode.....	127
Monitoring active/active configuration status.....	127
Monitoring the hardware-assisted takeover feature.....	128
Description of active/active configuration status messages.....	130
Displaying the partner's name.....	131
Displaying disk and array LUN information on an active/active configuration.....	131
Enabling and disabling takeover.....	132
Enabling and disabling automatic takeover of a panicked partner.....	132
Halting a node without takeover.....	133
Configuration of when takeover occurs.....	133
Reasons for takeover.....	133
Commands for performing a takeover.....	135
Specifying the time period before takeover.....	136
How disk shelf comparison takeover works.....	137

Configuring VIFs or interfaces for automatic takeover.....	137
Takeover of vFiler units and the vFiler unit limit.....	137
Managing an active/active configuration in takeover mode.....	138
Determining why takeover occurred.....	138
Statistics in takeover mode.....	138
Managing emulated nodes.....	139
Management exceptions for emulated nodes.....	139
Accessing the emulated node from the takeover node.....	139
Assessing the emulated node remotely.....	141
Emulated node command exceptions.....	141
Performing dumps and restores for a failed node.....	143
Giveback operations.....	144
Performing a giveback.....	144
Configuring giveback.....	147
Enabling automatic giveback.....	148
Downloading and running the HA Configuration Checker utility.....	149
Troubleshooting takeover or giveback failures.....	149
<b>Management of DS14mk2 AT, DS14mk2 FC, or</b>	
<b>    DS14mk4 FC disk shelves in an active/active configuration.....</b>	<b>151</b>
Managing disk shelves in Multipath Storage configurations.....	151
What Multipath Storage for active/active configurations is.....	151
How the connection types are used.....	152
Advantages of Multipath Storage for active/active configurations.....	153
Requirements for Multipath Storage.....	153
Determining whether your AT-FCX modules support	
Multipath Storage.....	155
Cabling for Multipath Storage.....	156
Adding storage to a Multipath Storage loop.....	157
Adding disk shelves to non-Multipath Storage configurations.....	159
Overview of adding storage to non-multipath configurations.....	159
Adding storage to an existing non-multipath loop.....	161
Adding a new non-multipath loop.....	163
Adding storage to fabric-attached MetroClusters.....	164
Upgrading or replacing modules in an active/active configuration.....	164
About the disk shelf modules.....	165
Restrictions for changing module types.....	165

Best practices for changing module types.....	166
Testing the modules.....	166
Understanding redundant pathing in active/active configurations.....	167
Determining path status for your active/active configuration.....	167
Upgrading an LRC module to an ESH or ESH2 module.....	169
Hot-swapping a module.....	171
<b>Disaster recovery using MetroCluster.....</b>	<b>173</b>
Conditions that constitute a disaster.....	173
Ways to determine whether a disaster occurred.....	173
Failures that do not require disaster recovery.....	174
Recovering from a disaster.....	175
Restricting access to the disaster site node.....	175
Forcing a node into takeover mode.....	176
Remounting volumes of the failed node.....	177
Recovering LUNs of the failed node.....	177
Fixing failures caused by the disaster.....	178
Reestablishing the MetroCluster configuration.....	179
<b>Nondisruptive hardware changes.....</b>	<b>185</b>
Replacing a component nondisruptively.....	185
Removing the old hardware when nondisruptively changing hardware.....	186
Installing the new hardware when nondisruptively changing hardware .....	187
<b>Controller failover and single-points-of-failure.....</b>	<b>189</b>
Single-point-of-failure definition.....	189
SPOF analysis for active/active configurations.....	189
Failover event cause-and-effect table.....	192
<b>Feature update record.....</b>	<b>199</b>
<b>Abbreviations.....</b>	<b>203</b>
<b>Index.....</b>	<b>217</b>



# Copyright information

---

Copyright © 1994–2009 NetApp, Inc. All rights reserved. Printed in the U.S.A.

No part of this document 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.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. 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 NetApp.

The product described in this manual may be protected by one or more U.S.A. 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 logo, the bolt design, NetApp-the Network Appliance Company, Cryptainer, Cryptoshred, DataFabric, DataFort, Data ONTAP, Decru, FAServer, FilerView, FlexClone, FlexVol, Manage ONTAP, MultiStore, NearStore, NetCache, NOW NetApp on the Web, SANscreen, SecureShare, SnapDrive, SnapLock, SnapManager, SnapMirror, SnapMover, SnapRestore, SnapValidator, SnapVault, Spinnaker Networks, SpinCluster, SpinFS, SpinHA, SpinMove, SpinServer, StoreVault, SyncMirror, Topio, VFM, VFM Virtual File Manager, and WAFL are registered trademarks of NetApp, Inc. in the U.S.A. and/or other countries. gFiler, Network Appliance, SnapCopy, Snapshot, and The evolution of storage are trademarks of NetApp, Inc. in the U.S.A. and/or other countries and registered trademarks in some other countries. The NetApp arch logo; the StoreVault logo; ApplianceWatch; BareMetal; Camera-to-Viewer; ComplianceClock; ComplianceJournal; ContentDirector; ContentFabric; EdgeFiler; FlexShare; FPolicy; Go Further, Faster; HyperSAN; InfoFabric; Lifetime Key Management, LockVault; NOW; ONTAPI; OpenKey, RAID-DP; ReplicatorX; RoboCache; RoboFiler; SecureAdmin; Serving Data by Design; Shadow Tape; SharedStorage; Simplicore; Simulate ONTAP; Smart SAN; SnapCache; SnapDirector; SnapFilter; SnapMigrator; SnapSuite; SohoFiler; SpinMirror; SpinRestore; SpinShot; SpinStor; vFiler; VPolicy; and Web Filer are trademarks of NetApp, Inc. in the U.S.A. and other countries. NetApp Availability Assurance and NetApp ProTech Expert are service marks of NetApp, Inc. in the U.S.A.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the U.S.A. and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the U.S.A. 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 U.S.A. and/or other countries.

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

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.



# About this guide

---

You can use your product more effectively when you understand this document's intended audience and the conventions that this document uses to present information.

This guide, previously published as the *Cluster Installation and Administration Guide*, describes the following tasks and topics:

- Installing and administering a standard or mirrored active/active configuration
- Installing and administering a stretch MetroCluster
- Installing and administering a fabric-attached MetroCluster
- Managing storage in an active/active configuration
- Performing disaster recovery with a MetroCluster

See the *Data ONTAP Release Notes* for the list of storage systems that support active/active configurations.

**Note:** This guide covers administration of the Data ONTAP software for all systems in an active/active configuration, including the dual-controller FAS20xx, 31xx, or GF270c systems. However, it does not include hardware information for dual-controller systems. For more information, see the hardware documentation for the specific system.

## Next topics

[Audience](#) on page 13

[Accessing Data ONTAP man pages](#) on page 14

[Terminology](#) on page 15

[Where to enter commands](#) on page 16

[Keyboard and formatting conventions](#) on page 16

[Special messages](#) on page 17

[How to send your comments](#) on page 18

## Audience

This document is written with certain assumptions about your technical knowledge and experience.

Refer to this guide if you need to perform the following tasks:

- Cable and configure two systems into a standard or mirrored active/active configuration
- Convert stand-alone systems into a standard or mirrored active/active configuration
- Convert an active/active configuration into two stand-alone systems

- Cable and configure a fabric-attached or stretch MetroCluster
- Perform recovery in the event of a disaster at a MetroCluster node
- Manage an active/active configuration
- Manage storage on an active/active configuration

## Accessing Data ONTAP man pages

You can use the Data ONTAP manual (man) pages to access technical information.

### About this task

Data ONTAP manual pages are available for the following types of information. They are grouped into sections according to standard UNIX naming conventions.

Types of information	Man page section
Commands	1
Special files	4
File formats and conventions	5
System management and services	8

### Step

1. View man pages in the following ways:

- Enter the following command at the storage system command line:
 

```
man command_or_file_name
```
- Click the manual pages button on the main Data ONTAP navigational page in the FilerView user interface.
- Use the *Commands: Manual Page Reference*, Volumes 1 and 2 (which can be downloaded or ordered through the NOW site).

**Note:** All Data ONTAP man pages are stored on the storage system in files whose names are prefixed with the string "na\_" to distinguish them from client man pages. The prefixed names are used to distinguish storage system man pages from other man pages and sometimes appear in the NAME field of the man page, but the prefixes are not part of the command, file, or services.

## Terminology

To understand the concepts in this document, you might need to know how certain terms are used.

### Storage terms

<b>array LUN</b>	Refers to storage that third-party storage arrays provide to storage systems running Data ONTAP software. One array LUN is the equivalent of one disk on a native disk shelf.
<b>LUN (Logical Unit Number)</b>	Refers to a logical unit of storage identified by a number.
<b>native disk</b>	Refers to a disk that is sold as local storage for storage systems that run Data ONTAP software.
<b>native disk shelf</b>	Refers to a disk shelf that is sold as local storage for storage systems that run Data ONTAP software.
<b>storage controller</b>	Refers to the component of a storage system that runs the Data ONTAP operating system and controls its disk subsystem. Storage controllers are also sometimes called <i>controllers</i> , <i>storage appliances</i> , <i>appliances</i> , <i>storage engines</i> , <i>heads</i> , <i>CPU modules</i> , or <i>controller modules</i> .
<b>storage system</b>	Refers to the hardware device running Data ONTAP that receives data from and sends data to native disk shelves, third-party storage, or both. Storage systems that run Data ONTAP are sometimes referred to as <i>filers</i> , <i>appliances</i> , <i>storage appliances</i> , <i>V-Series systems</i> , or <i>systems</i> .
<b>third-party storage</b>	Refers to back-end storage arrays, such as IBM, Hitachi Data Systems, and HP, that provide storage for storage systems running Data ONTAP.

### Cluster and high-availability terms

<b>active/active configuration</b>	In the Data ONTAP 7.2 and 7.3 release families, refers to a pair of storage systems (sometimes called <i>nodes</i> ) configured to serve data for each other if one of the two systems stops functioning. Also sometimes referred to as <i>active/active pairs</i> . In the Data ONTAP 7.1 release family and earlier releases, this functionality is referred to as a <i>cluster</i> .
<b>cluster</b>	In the Data ONTAP 7.1 release family and earlier releases, refers to a pair of storage systems (sometimes called <i>nodes</i> ) configured to serve data for each other if one of the two systems stops functioning. In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an <i>active/active configuration</i> .

## Where to enter commands

You can use your product more effectively when you understand how this document uses command conventions to present information.

You can perform common administrator tasks in one or more of the following ways:

- You can enter commands either at the system console or from any client computer that can obtain access to the storage system using a Telnet or Secure Shell (SSH) session.  
In examples that illustrate command execution, the command syntax and output shown might differ from what you enter or see displayed, depending on your version of the operating system.
- You can use the FilerView graphical user interface.  
For information about accessing your system with FilerView, see the *Data ONTAP System Administration Guide*.
- You can enter Windows, ESX, HP-UX, AIX, Linux, and Solaris commands at the applicable client console.  
In examples that illustrate command execution, the command syntax and output shown might differ from what you enter or see displayed, depending on your version of the operating system.

## Keyboard and formatting conventions

You can use your product more effectively when you understand how this document uses keyboard and formatting conventions to present information.

### Keyboard conventions

Convention	What it means
The NOW site	Refers to <i>NetApp On the Web</i> at <a href="http://now.netapp.com/">http://now.netapp.com/</a> .
<i>Enter, enter</i>	<ul style="list-style-type: none"> <li>• Used to refer to the key that generates a carriage return; the key is named Return on some keyboards.</li> <li>• Used to mean pressing one or more keys on the keyboard and then pressing the Enter key, or clicking in a field in a graphical interface and then typing information into the field.</li> </ul>
hyphen (-)	Used to separate individual keys. For example, Ctrl-D means holding down the Ctrl key while pressing the D key.



Convention	What it means
type	Used to mean pressing one or more keys on the keyboard.

## Formatting conventions

Convention	What it means
<i>Italic font</i>	<ul style="list-style-type: none"> <li>Words or characters that require special attention.</li> <li>Placeholders for information that you must supply. For example, if the guide says to enter the <code>arp -d hostname</code> command, you enter the characters "arp -d" followed by the actual name of the host.</li> <li>Book titles in cross-references.</li> </ul>
Monospaced font	<ul style="list-style-type: none"> <li>Command names, option names, keywords, and daemon names.</li> <li>Information displayed on the system console or other computer monitors.</li> <li>Contents of files.</li> <li>File, path, and directory names.</li> </ul>
<b>Bold monospaced font</b>	Words or characters you type. What you type is always shown in lowercase letters, unless your program is case-sensitive and uppercase letters are necessary for it to work properly.

## Special messages

This document might contain the following types of messages to alert you to conditions that you need to be aware of.

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

**Attention:** An attention notice contains instructions that you must follow to avoid a system crash, loss of data, or damage to the equipment.

## How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by e-mail to [docomments@netapp.com](mailto:docomments@netapp.com). To help us direct your comments to the correct division, include in the subject line the name of your product and the applicable operating system. For example, *FAS6070—Data ONTAP 7.3*, or *Host Utilities—Solaris*, or *Operations Manager 3.8—Windows*.

# Active/active configuration types and requirements

---

There are four types of active/active configurations, each having different advantages and requirements.

## Next topics

[Overview of active/active configurations](#) on page 19

[Standard active/active configurations](#) on page 23

[Understanding mirrored active/active configurations](#) on page 26

[Understanding stretch MetroClusters](#) on page 28

[Understanding fabric-attached MetroClusters](#) on page 33

## Overview of active/active configurations

The different types of active/active configurations all offer access to storage through two different controllers. Each type has its own benefits and requirements.

## Next topics

[What an active/active configuration is](#) on page 19

[Benefits of HA](#) on page 20

[Characteristics of nodes in an active/active configuration](#) on page 20

[Best practices for deploying an active/active configuration](#) on page 21

[Comparison of active/active configuration types](#) on page 22

## What an active/active configuration is

An active/active configuration is two storage systems (nodes) whose controllers are connected to each other either directly or, in the case of a fabric-attached MetroCluster, through switches and FC-VI interconnect adapters.

You can configure the active/active configuration so that each node in the pair shares access to a common set of storage, subnets, and tape drives, or each node can own its own distinct set of storage and subnets.

The nodes are connected to each other through a NVRAM adapter, or, in the case of systems with two controllers in a single chassis, through an internal interconnect. This allows one node to serve data that resides on the disks of its failed partner node. Each node continually monitors its partner, mirroring the data for each other's nonvolatile memory (NVRAM or NVMEM).

## Benefits of HA

Active/active configurations provide fault tolerance and the ability to perform nondisruptive upgrades and maintenance.

Configuring storage systems in an active/active configuration provides the following benefits:

- **Fault tolerance**  
When one node fails or becomes impaired a takeover occurs, and the partner node continues to serve the failed node's data.
- **Nondisruptive software upgrades**  
When you halt one node and allow takeover, the partner node continues to serve data for the halted node while you upgrade the node you halted.
- **Nondisruptive hardware maintenance**  
When you halt one node and allow takeover, the partner node continues to serve data for the halted node while you replace or repair hardware in the node you halted.

### Related concepts

[Nondisruptive hardware changes](#) on page 185

[Management of DS14mk2 AT, DS14mk2 FC, or DS14mk4 FC disk shelves in an active/active configuration](#) on page 151

## Characteristics of nodes in an active/active configuration

To configure and manage nodes in an active/active configuration, you should be familiar with the characteristics that all types of active/active configurations have in common.

- They are connected to each other either through a cluster interconnect consisting of adapters and cable, or, in systems with two controllers in the same chassis, through an internal interconnect. The nodes use the interconnect to do the following tasks:
  - Continually check whether the other node is functioning
  - Mirror log data for each other's NVRAM
  - Synchronize each other's time
- They use two or more disk shelf loops, or third-party storage, in which the following conditions apply:
  - Each node manages its own disks or array LUNs.
  - Each node in takeover mode manages its partner's disks or array LUNs. For third-party storage, the partner node takes over read/write access to the array LUNs owned by the failed node until the failed node becomes available again.

**Note:** For systems using software-based disk ownership, disk ownership is established by Data ONTAP or the administrator, rather than by which disk shelf the disk is attached to.

For more information about disk ownership, see the *Data ONTAP Storage Management Guide*.

- They own their spare disks, spare array LUNs, or both and do not share them with the other node.
- They each have mailbox disks or array LUNs on the root volume:
  - Two if it is a FAS system (four if the root volume is mirrored using the SyncMirror feature).
  - One if it is a V-Series system (two if the root volume is mirrored using the SyncMirror feature).

The mailbox disks or LUNs are used to do the following tasks:

- Maintain consistency between the pair
  - Continually check whether the other node is running or whether it has performed a takeover
  - Store configuration information that is not specific to any particular node
- They can reside on the same Windows domain or on different domains.

## Best practices for deploying an active/active configuration

To ensure that your active/active configuration is robust and operational, you need to be familiar with configuration best practices.

- Make sure that the controllers and disk shelves are on different power supplies or grids, so that a single power outage does not affect both components.
- Use VIFs (virtual interfaces) to provide redundancy and improve availability of network communication.
- Follow the documented procedures in the *Data ONTAP Upgrade Guide* when upgrading your active/active configuration.
- Maintain consistent configuration between the two nodes. An inconsistent configuration is often the cause of failover problems.
- Test the failover capability routinely (for example, during planned maintenance) to ensure proper configuration.
- Make sure that each node has sufficient resources to adequately support the workload of both nodes during takeover mode.
- Use the HA Configuration Checker to help ensure that failovers are successful.
- If your systems support the Remote LAN Module (RLM), make sure you configure RLM properly, as described in the RLM chapter of the *Data ONTAP System Administration Guide*.
- Higher numbers of traditional and FlexVol volumes on your system can affect takeover and giveback times. When adding traditional or FlexVol volumes to an active/active configuration, consider testing the takeover and giveback times to ensure that they fall within your requirements.
- For systems using disks, check for and remove any failed disks, as described in the *Data ONTAP Storage Management Guide*.

**Related tasks**

[Downloading and running the HA Configuration Checker utility](#) on page 149

**Comparison of active/active configuration types**

Outlines the differences between the different types of active/active configurations, and when you might want to use each type.

Active/active configuration type	Data duplication?	Distance between nodes	Failover possible after loss of entire node (including storage)?	Notes
Standard active/active configuration	No	Up to 500 meters <b>Note:</b> SAS configurations are limited to 5 meters between nodes	No	Use this configuration to provide higher availability by protecting against many hardware single-points-of-failure.
Mirrored active/active configuration	Yes	Up to 500 meters <b>Note:</b> SAS configurations are limited to 5 meters between nodes	No	Use this configuration to add increased data protection to the benefits of a standard active/active configuration .
Stretch MetroCluster	Yes	Up to 500 meters (270 meters if operating at 4 Gbps)	Yes	Use this configuration to provide data and hardware duplication to protect against a local disaster (for example, a power outage to one node).

Active/active configuration type	Data duplication?	Distance between nodes	Failover possible after loss of entire node (including storage)?	Notes
Fabric-attached MetroCluster	Yes	Up to 100 kilometers, depending on switch configuration. For FAS systems, see the <i>Brocade Switch Configuration Guide for Fabric-attached MetroClusters</i> . For V-Series systems, up to 30 km.	Yes	Use this configuration to provide data and hardware duplication to protect against a larger scale disaster, such as the loss of an entire site.

## Standard active/active configurations

Standard active/active configurations provide high availability (HA) by pairing two controllers so that one can serve data for the other in case of controller failure or other unexpected events.

### Next topics

[How Data ONTAP works with standard active/active configurations](#) on page 23

[Standard active/active configuration diagram](#) on page 24

[Setup requirements and restrictions for standard active/active configurations](#) on page 24

[Configuration variations for standard active/active configurations](#) on page 25

### Related references

[SPOF analysis for active/active configurations](#) on page 189

## How Data ONTAP works with standard active/active configurations

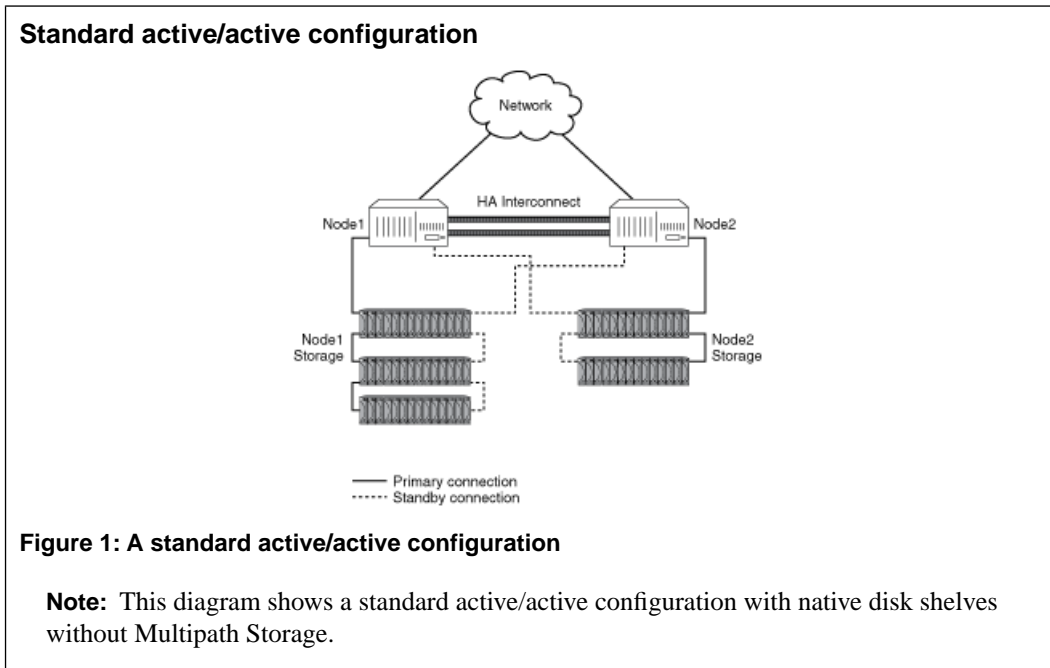
In a standard active/active configuration, Data ONTAP functions so that each node monitors the functioning of its partner through a heartbeat signal sent between the nodes. Data from the NVRAM of one node is mirrored by its partner, and each node can take over the partner's disks or array LUNs if the partner fails. Also, the nodes synchronize each other's time.

**Note:** If a node reboots (but a takeover does not occur), note that the HA interconnect link comes up prior to Data ONTAP completely loading on the rebooting partner. Commands issued on the surviving controller (that is not rebooting) that check the status of the partner or configuration may indicate that the partner could not be reached. Wait until the partner has fully rebooted and reissue the command.

In some cases (such as the `lun config_check` command) these commands are issued automatically when the interconnect comes up. The resulting error can generate an AutoSupport indicating a configuration problem when in fact the underlying problem is that Data ONTAP has not fully booted.

## Standard active/active configuration diagram

Shows an example standard active/active configuration using native disk storage without Multipath Storage.



### Related concepts

[Managing disk shelves in Multipath Storage configurations](#) on page 151

## Setup requirements and restrictions for standard active/active configurations

You must follow certain requirements and restrictions when setting up a new standard active/active configuration.

The following list specifies the requirements and restrictions to be aware of when setting up a new standard active/active configuration:

- Architecture compatibility
  - Both nodes must have the same system model and be running the same firmware version. See the *Data ONTAP Release Notes* for the list of supported systems.



**Note:** In the case of systems with two controller modules in a single chassis (except the 31xx systems), both nodes of the active/active configuration are located in the same chassis and have an internal interconnect.

- Storage capacity

The number of disks or array LUNs must not exceed the maximum configuration capacity. If your system uses both native disks and third-party storage, the combined total of disks and array LUNs cannot exceed the maximum configuration capacity. In addition, the total storage attached to each node must not exceed the capacity for a single node.

To determine the maximum capacity for a system using disks, see the *System Configuration Guide* at [http://now.netapp.com/NOW/knowledge/docs/hardware/hardware\\_index.shtml](http://now.netapp.com/NOW/knowledge/docs/hardware/hardware_index.shtml). For a system using array LUNs, disks, or both, see the *V-Series Support Matrix*.

**Note:** After a failover, the takeover node temporarily serves data from all the storage in the active/active configuration. When the single-node capacity limit is less than the total active/active configuration capacity limit, the total disk space in an active/active configuration can be greater than the single-node capacity limit. It is acceptable for the takeover node to temporarily serve more than the single-node capacity would normally allow, as long as it does not own more than the single-node capacity.

- Disks and disk shelf compatibility

- Both Fibre Channel and SATA storage are supported in standard active/active configurations, as long as the two storage types are not mixed on the same loop.
- One node can have only Fibre Channel storage and the partner node can have only SATA storage if needed.

- Cluster interconnect adapters and cables must be installed, unless the system has two controllers in the chassis and an internal interconnect.
- Nodes must be attached to the same network and the Network Interface Cards (NICs) must be configured correctly.
- The same system software, such as Common Internet File System (CIFS), Network File System (NFS), or SyncMirror, must be licensed and enabled on both nodes.

**Note:** If a takeover occurs, the takeover node can provide only the functionality for the licenses installed on it. If the takeover node does not have a license that was being used by the partner node to serve data, your active/active configuration loses functionality after a takeover.

## Configuration variations for standard active/active configurations

Active/active configurations can be configured asymmetrically, as an active/passive pair, with shared disk shelf stacks, or with Multipath Storage.

- Asymmetrical configurations

In an asymmetrical standard active/active configuration, one node has more storage than the other. This is supported, as long as neither node exceeds the maximum capacity limit for the node.

- **Active/passive configurations**  
In this configuration, the passive node has only a root volume, and the active node has all the remaining storage and services all data requests during normal operation. The passive node responds to data requests only if it has taken over the active node.
- **Shared loops or stacks**  
If your standard active/active configuration is using software-based disk ownership, you can share a loop or stack between the two nodes. This is particularly useful for active/passive configurations, as described in the preceding bullet.
- **Multipath Storage**  
Multipath Storage for active/active configurations using native disk shelves provides a redundant connection from each node to every disk. It can prevent some types of failovers.

#### Related concepts

[Managing disk shelves in Multipath Storage configurations](#) on page 151

## Understanding mirrored active/active configurations

Mirrored active/active configurations provide high availability through failover, just as standard active/active configurations do. Additionally, mirrored active/active configurations maintain two complete copies of all mirrored data. These copies are called plexes and are continually and synchronously updated every time Data ONTAP writes to a mirrored aggregate. The plexes can be physically separated to protect against the loss of one set of disks or array LUNs.

**Note:** Mirrored active/active configurations do not provide the capability to fail over to the partner node if one node is completely lost. For example, if power is lost to one entire node, including its storage, you cannot fail over to the partner node. For this capability, use a MetroCluster.

Mirrored active/active configurations use SyncMirror. For more information about SyncMirror, see the *Data ONTAP Data Protection Online Backup and Recovery Guide*.

#### Next topics

[Advantages of mirrored active/active configurations](#) on page 26

[Setup requirements and restrictions for mirrored active/active configurations](#) on page 27

[Configuration variations for mirrored active/active configurations](#) on page 28

## Advantages of mirrored active/active configurations

Data mirroring provides additional data protection in the event of disk failures and reduces the need for failover in the event of other component failures.

Mirroring your data protects it from the following problems which would cause data loss without mirroring:

- The failure or loss of two or more disks in a RAID4 aggregate
- The failure or loss of three or more disks in a RAID-DP (RAID double-parity) aggregate
- The failure of an array LUN; for example, because of a double disk failure on the storage array.
- The failure of a third-party storage array.

The failure of an FC-AL adapter, loop, or disk shelf module does not require a failover in a mirrored active/active configuration.

Similar to standard active/active configurations, if either node in a mirrored active/active configuration becomes impaired or cannot access its data, the other node can automatically serve the impaired node's data until the problem is corrected.

## Setup requirements and restrictions for mirrored active/active configurations

The restrictions and requirements for mirrored active/active configurations include those for a standard active/active configuration with these additional requirements for disk pool assignments and cabling.

- You must ensure that your pools are configured correctly:
  - Disks or array LUNs in the same plex must be from the same pool, with those in the opposite plex from the opposite pool.
  - If hardware-based ownership is used on your systems, the disk shelves must be connected to the controllers so that the disks do not have to change pools when a takeover occurs.  
For example, on a FAS3020 system, if you connect an HBA in slot 2 to Channel A (the A Input port on the disk shelf), you should connect Channel B to an HBA that is also in pool 0 on the partner node. If you connect Channel B to an HBA in, for example, slot 4, the disks would have to change from pool 0 to pool 1 when a takeover occurs.  
For more information about how Data ONTAP assigns pools ownership, see the section about hardware-based disk ownership in the *Data ONTAP Storage Management Guide*.
- There must be sufficient spares in each pool to account for a disk or array LUN failure.
  - Note:** If your systems are using hardware-based disk ownership, pool membership is determined by the physical connections between the disk shelves and the controllers. If your systems are using software-based disk ownership, pool membership is determined explicitly using the Data ONTAP command-line interface. For more information, see the section on disk ownership in the *Data ONTAP Storage Management Guide*.
- On systems using software ownership, both plexes of a mirror should not reside on the same disk shelf, as it would result in a single point of failure.

See the *Data ONTAP Data Protection Online Backup and Recovery Guide* for more information about requirements for setting up SyncMirror with third-party storage

- You must enable the following licenses on both nodes:
  - cluster
  - syncmirror\_local

- If you are using third-party storage, paths to an array LUN must be redundant.

#### Related concepts

[Setup requirements and restrictions for standard active/active configurations](#) on page 24

## Configuration variations for mirrored active/active configurations

A number of configuration variations are supported for mirrored active/active configurations .

The following list describes some configuration variations that are supported for mirrored active/active configurations :

- Asymmetrical mirroring

You can selectively mirror your storage. For example, you could mirror all the storage on one node, but none of the storage on the other node. Takeover will function normally. However, any unmirrored data is lost if the storage that contains it is damaged or destroyed.

**Note:** You must connect the unmirrored storage to both nodes, just as for mirrored storage. You cannot have storage that is connected to only one node in an active/active configuration.

- Multipath Storage

Multipath Storage for native disk shelves in active/active configurations provides a redundant connection from each node to every disk. It can help prevent some types of failovers.

#### Related concepts

[Managing disk shelves in Multipath Storage configurations](#) on page 151

## Understanding stretch MetroClusters

Stretch MetroClusters provide data mirroring and the additional ability to initiate a failover if an entire site becomes lost or unavailable.

Like mirrored active/active configurations, stretch 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 active/active configuration. These copies are called plexes and are continually and synchronously updated every time Data ONTAP writes data to the disks. Plexes are physically separated from each other across different groupings of disks.

Unlike mirrored active/active configurations, MetroClusters provide the capability to force a failover when an entire node (including the controllers and storage) is destroyed or unavailable.

**Note:** In previous versions of this document, stretch MetroClusters were called nonswitched MetroClusters.

**Note:** If you are a V-Series system customer, see the *V-Series MetroCluster Guide* for information about configuring and operating a V-Series system in a MetroCluster configuration

**Next topics**

[Continued data service after loss of one node with MetroCluster](#) on page 29

[Advantages of stretch MetroCluster configurations](#) on page 29

[Stretch MetroCluster configuration](#) on page 29

[Stretch MetroCluster configuration on 31xx systems](#) on page 31

[How Data ONTAP works with stretch MetroCluster configurations](#) on page 31

[Stretch MetroCluster and disk ownership](#) on page 31

[Setup requirements and restrictions for stretch MetroCluster configurations](#) on page 32

[Configuration variations for stretch MetroCluster configurations](#) on page 32

[MetroClusters and SnapLock volumes](#) on page 33

**Continued data service after loss of one node with MetroCluster**

The MetroCluster configuration employs SyncMirror to build a system that can continue to serve data even after complete loss of one of the nodes and the storage at that site. Data consistency is retained, even when the data is contained in more than one aggregate.

**Note:** You can have both mirrored and unmirrored volumes in a MetroCluster. However, the MetroCluster configuration can preserve data only if volumes are mirrored. Unmirrored volumes are lost if the storage where they reside is destroyed.

See the *Data ONTAP Data Protection Online Backup and Recovery Guide* for detailed information about using SyncMirror to mirror data.

**Advantages of stretch MetroCluster configurations**

MetroClusters provide the same advantages of mirroring as mirrored Active/active configurations, with the additional ability to initiate failover if an entire site becomes lost or unavailable.

- Your data is protected if there is a failure or loss of two or more disks in a RAID4 aggregate or three or more disks in a RAID-DP aggregate.
- The failure of an FC-AL adapter, loop, or ESH2 module does not require a failover.

In addition, a MetroCluster provides the `cf forcetakeover -d` command, giving you a single command to initiate a failover if an entire site becomes lost or unavailable. If a disaster occurs at one of the node locations and destroys your data there, your data not only survives on the other node, but can be served by that node while you address the issue or rebuild the configuration.

**Related concepts**

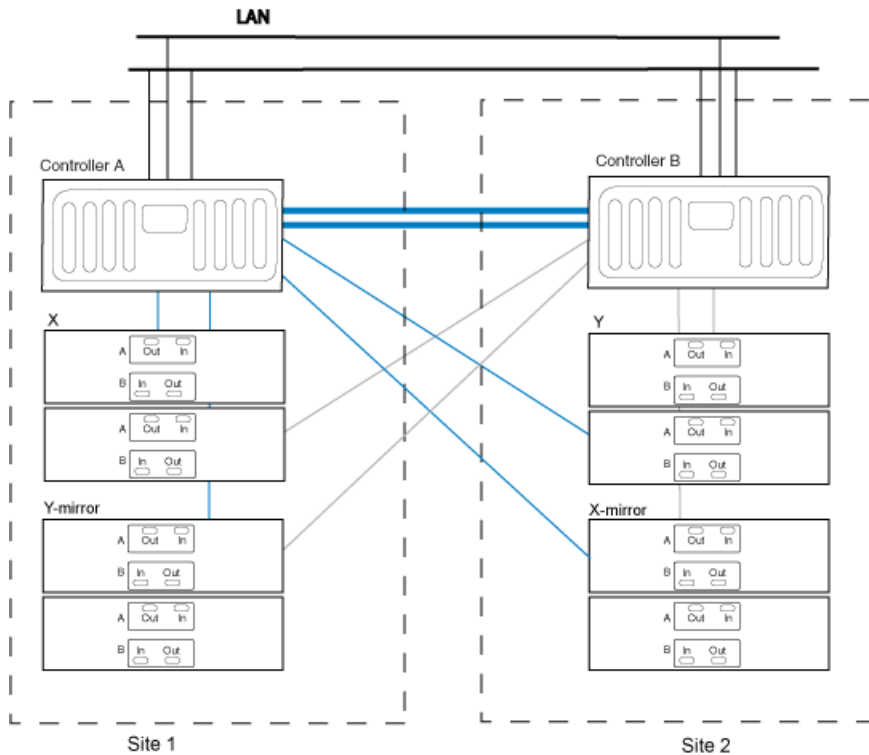
[Disaster recovery using MetroCluster](#) on page 173

**Stretch MetroCluster configuration**

You configure a stretch MetroCluster so that each controller can access its own storage and its partner's storage, with local storage mirrored at the partner site.

The following figure illustrates the stretch MetroCluster configuration. The configuration includes the following connections:

- Connections from each controller to the user network.
- The MetroCluster interconnect between the two controllers.
- Connections from each controller to its own storage:
  - Controller A to X
  - Controller B to Y
- Connections from each controller to its partner's storage:
  - Controller A to Y
  - Controller B to X
- Connections from each controller to the mirrors of its storage:
  - Controller A to X-mirror
  - Controller B to Y-mirror




**Note:** This is a simplified figure that does not show disk shelf-to-disk shelf connections.

**Figure 2: Stretch MetroCluster**

## Stretch MetroCluster configuration on 31xx systems

A stretch MetroCluster can be configured between a pair of 31xx systems in which each system has a single controller (rather than two).

To implement the stretch MetroCluster, an FC-VI adapter must be installed in each controller to provide the cluster interconnect between the systems. When the FC-VI adapter is installed in the system, the internal InfiniBand interconnect is automatically disabled. This is different from other stretch MetroClusters, which use NVRAM adapters to provide the interconnect.



**Figure 3: Stretch MetroCluster with 31xx systems**

**Note:** The dual-controller FAS200 series and FAS20xx do not support MetroClusters.

## How Data ONTAP works with stretch MetroCluster configurations

Data ONTAP divides storage across physically separated pools of disks.

During configuration, 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.

Data ONTAP can also be configured to read from both plexes, which in many cases improves read performance.

**Note:** You can determine which side of the mirrored volume (also called a plex) is read when a data request is received using the `raid.mirror_read_plex_pref` option. For more information, see the `na_options(1)` man page.

## Stretch MetroCluster and disk ownership

The type of disk ownership used by the system (hardware-based or software-based) determines how pool membership is determined for the disk shelves in the MetroCluster

For systems using hardware-based disk ownership, pool membership is determined by the physical connections between the controller and the disk shelf. For systems using software-based disk ownership, pool membership is determined either by Data ONTAP, or by the administrator, using the Data ONTAP command-line interface.

On some systems, only software-based disk ownership is available and you must use the Data ONTAP disk commands to manage pool membership.

For more information about disk ownership, see the *Data ONTAP Storage Management Guide*.

## Setup requirements and restrictions for stretch MetroCluster configurations

You must follow certain requirements and restrictions when setting up a new Stretch MetroCluster configuration.

The restrictions and requirements for stretch MetroClusters include those for a standard active/active configuration and those for a mirrored active/active configuration. In addition, the following requirements apply:

- Both SATA and Fibre Channel storage is supported on stretch MetroClusters, but both plexes of the same aggregate must use the same type of storage. For example, you cannot mirror a Fibre Channel aggregate with SATA storage.
- MetroCluster is not supported on the FAS20xx platforms.
- The following licenses must be enabled on both nodes:
  - cluster
  - syncmirror\_local
  - cluster\_remote

**Note:** See the MetroCluster Compatibility Matrix on the NOW site for more information about hardware and firmware requirements for this configuration.

### Related concepts

[Setup requirements and restrictions for standard active/active configurations](#) on page 24

[Setup requirements and restrictions for mirrored active/active configurations](#) on page 27

## Configuration variations for stretch MetroCluster configurations

Stretch MetroClusters have asymmetrical and active/passive variations.

The following list describes some common configuration variations that are supported for stretch MetroClusters:

- Asymmetrical mirroring  
You can add storage to one or both nodes that is not mirrored by the other node.



**Attention:** Any data contained in the unmirrored storage could be lost if that site experiences a disaster.

**Note:** Multiple disk failures in an unmirrored aggregate (three or more disk failures in a RAID-DP aggregate, two or more disk failures in a RAID4 aggregate) cause the node to panic, resulting in a temporary data service outage while the node reboots, a takeover occurs, or disaster recovery is performed.

You must mirror the root volumes to enable successful takeover.

**Note:** You must connect the unmirrored storage to both nodes, just as for mirrored storage. You cannot have storage that is connected to only one node in an active/active configuration.

- Active/passive MetroClusters

In this configuration, the remote (passive) node does not serve data unless it has taken over for the local (active) node. Mirroring the passive node's root volume is optional. However, both nodes must have all MetroCluster licenses installed so that remote takeover is possible.

## MetroClusters and SnapLock volumes

As with any volume on a mirrored aggregate, on a properly configured site MetroCluster enables SnapLock volumes to be mirrored from one site to the other while retaining the SnapLock characteristics. If you issue the `cf forcetakeover -d` command because of a complete disaster or other operational failure at the primary site, these mirrors are broken and the mirror site goes online. Once the failed site is restored, the mirrors can be resynchronized before performing the giveback to normal operation.

**Attention:** If for any reason the primary node has data that was not mirrored to the secondary prior to the execution of the `cf forcetakeover -d` command, data could be lost. Do not resynchronize the original disks of the primary site for a SnapLock volume until an additional backup has been made of those disks to assure availability of all data. This situation could arise, for example, if the link between the sites was down and the primary node had data written to it in the interim before the `cf forcetakeover -d` command was issued.

For more information about backing up data in SnapLock volumes using SnapMirror, see the *Data ONTAP Archive and Compliance Management Guide*.

## Understanding fabric-attached MetroClusters

Like mirrored active/active configurations, fabric-attached MetroClusters contain two complete, separate copies of the data volumes or file systems that you configured as mirrored volumes or file systems in your active/active configuration. The fabric-attached MetroCluster nodes can be physically distant from each other, beyond the 500 meter limit of a stretch MetroCluster.

**Note:** If you are a V-Series system customer, see the *V-Series MetroCluster Guide* for information about configuring and operating a V-Series system in a MetroCluster configuration.

**Next topics**

*Fabric-attached MetroClusters use Brocade Fibre Channel switches* on page 34

*Advantages of fabric-attached MetroCluster configurations* on page 34

*Fabric-attached MetroCluster configuration* on page 35

*Fabric-attached MetroCluster configuration on 31xx systems* on page 35

*How Data ONTAP works with fabric-attached MetroCluster configurations* on page 36

*Setup requirements and restrictions for fabric-attached MetroClusters* on page 36

*Configuration limitations for fabric-attached MetroClusters* on page 38

*Configuration variations for fabric-attached MetroClusters* on page 38

*MetroClusters and SnapLock volumes* on page 39

**Fabric-attached MetroClusters use Brocade Fibre Channel switches**

A MetroCluster configuration for distances greater than 500 meters connects the two nodes using four Brocade Fibre Channel switches in a dual-fabric configuration for redundancy.

Each site has two Fibre Channel switches, each of which is connected through an inter-switch link to a partner switch at the other site. The inter-switch links are fiber optic connections that provide a greater distance between nodes than other active/active configurations. See the *Brocade Switch Configuration Guide for Fabric MetroCluster*.

Each local switch combines with a partner switch to form a fabric. By using four switches instead of two, redundancy is provided to avoid single-points-of-failure in the switches and their connections.

Like a stretch MetroCluster configuration, a fabric-attached MetroCluster employs SyncMirror to build a system that can continue to serve data even after complete loss of one of the nodes and the storage at that site. Data consistency is retained, even when the data is contained in more than one aggregate.

**Related information**

*Brocade Switch Configuration Guide for Fabric MetroCluster -*

[http://now.netapp.com/NOW/download/software/sanswitch/fcp/Brocade/mc\\_ontap641\\_fabric\\_index.shtml](http://now.netapp.com/NOW/download/software/sanswitch/fcp/Brocade/mc_ontap641_fabric_index.shtml)

**Advantages of fabric-attached MetroCluster configurations**

Fabric-attached MetroClusters provide the same advantages of stretch MetroCluster configurations, while also enabling the physical nodes to be physically distant from each other.

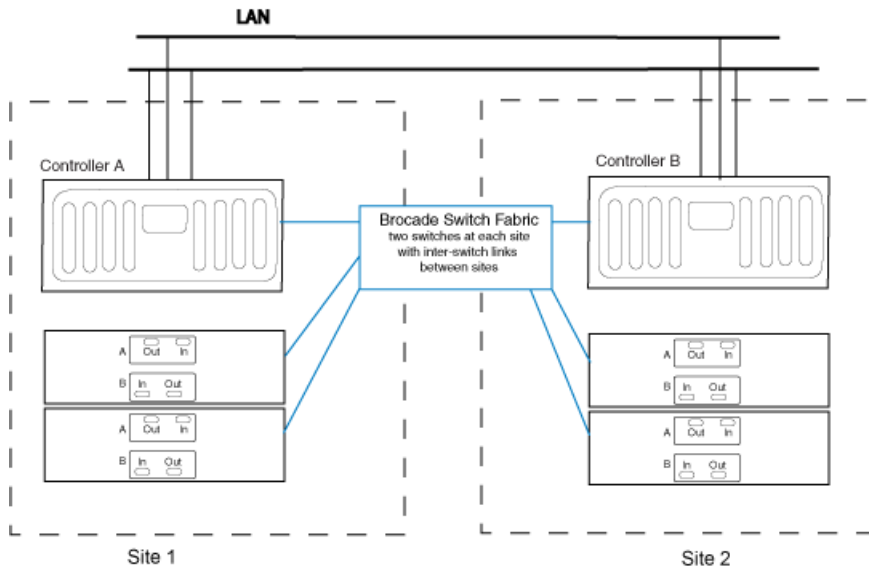
The advantages of a fabric-attached MetroCluster over a stretch MetroCluster include the following:

- The two halves of the configuration can be more than 500 meters apart, which provides increased disaster protection.
- 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.

## Fabric-attached MetroCluster configuration

A fabric-attached MetroCluster includes two Brocade Fibre Channel switch fabrics that provide long distance connectivity between the nodes. Through the Brocade switches, each controller can access its own storage and its partner's storage, with local storage mirrored at the partner site.

The following figure illustrates the fabric-attached MetroCluster configuration.



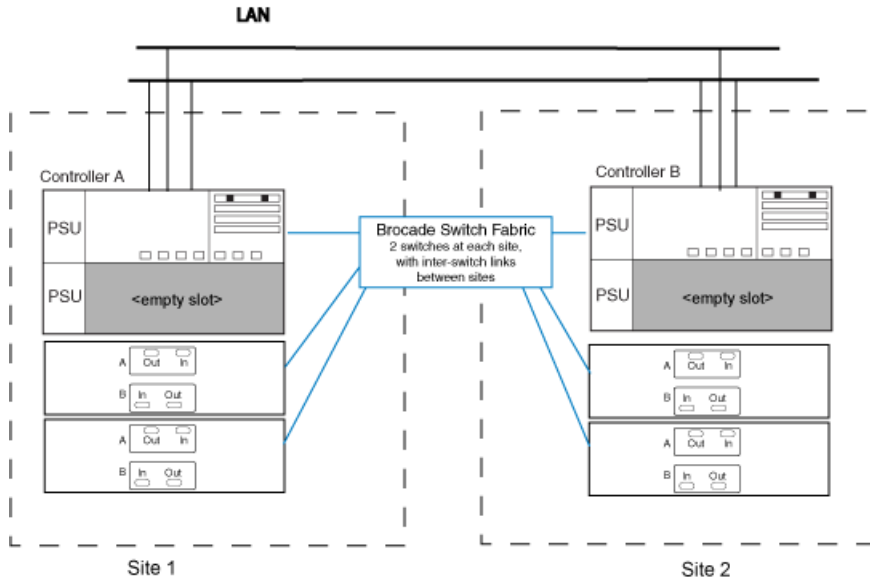
**Note:** This is a simplified figure that does not show disk shelf-to-disk shelf connections.

**Figure 4: Fabric-attached MetroCluster Configuration**

## Fabric-attached MetroCluster configuration on 31xx systems

A fabric-attached MetroCluster can be configured between a pair of 31xx systems in which each system has a single controller (rather than two).

When the system detects the presence of an FC-VI adapter, which connects the controller to the Brocade switch fabric, the internal InfiniBand connection is automatically deactivated.



**Note:** This is a simplified figure that does not show disk shelf-to-disk shelf connections.

**Figure 5: Fabric-attached MetroCluster configuration with 31xx systems**

## How Data ONTAP works with fabric-attached MetroCluster configurations

Data ONTAP functions the same way on a fabric-attached MetroCluster as on a stretch MetroCluster.

### Related concepts

[How Data ONTAP works with stretch MetroCluster configurations](#) on page 31

## Setup requirements and restrictions for fabric-attached MetroClusters

You must follow certain requirements and restrictions when setting up a new fabric-attached MetroCluster configuration.

The setup requirements for a fabric-attached MetroCluster include those for standard and mirrored Active/active configurations, with the following exceptions:

**Note:** See the MetroCluster Compatibility Matrix on the NOW site for more information about hardware and firmware requirements for this configuration.

### Node requirements

- The nodes must be one of the following system models configured for mirrored volume use; each node in the pair must be the same model.
  - FAS900 series systems

- FAS30xx systems
- 31xx systems

**Note:** The 31xx systems can have two controllers in the same chassis. When in a MetroCluster configuration, only a single controller in each system is supported (rather than two). The internal InfiniBand connections in each system are automatically deactivated.

- FAS60xx systems
- Each node requires a FC-VI (Fibre Channel/Virtual Interface) adapter; the slot position is dependent on the controller model.

The 4-Gbps FC-VI adapter is supported on the the following systems using software-based disk ownership:

- FAS60xx
- FAS3040
- FAS3070
- FAS31xx

**Note:** For information about supported cards and slot placement, see the *System Configuration Guide* on the NOW site.

The FC-VI adapter is also called a VI-MC or VI-MetroCluster adapter.

## Disk and disk shelf requirements

- Only DS14mk2 and DS14mk4 FC disk shelves are supported.
- Only Fibre Channel disks are supported; you cannot use SATA drives or AT-FCX modules for fabric-attached MetroCluster configurations.
- You can connect a maximum of two disk shelves to each loop.

## Capacity limits

The maximum capacity for a system configured in a fabric-attached MetroCluster is the smallest of the following limits:

- The maximum storage capacity for the node
  - Note:** For the maximum storage capacity, see the *System Configuration Guide* on the NOW site.
- 672 Fibre Channel disks (48 disk shelves) on the following systems:
  - FAS60xx
  - 3160 or 3170

To support 504 or more disks, the configuration must use the 24-port or 40-port Brocade Fibre Channel switch and the 4-Gbps FC-VI adapter. Zoning must be configured on the Brocade switches to separate storage and FC-VI traffic.

- 504 Fibre Channel disks (36 disk shelves) using the 16-port switch on FAS3070 systems
- 336 Fibre Channel disks (24 disk shelves) using the 24-port switch on other supported systems

### Fibre Channel switch requirements

**Note:** For the most up-to-date switch information, including supported switches and firmware downloads, see the Fabric-Attached MetroCluster Switch Description page on the NOW site. (To access this page, navigate to Download Software > Fibre Channel Switch > Brocade.)

- Each site of the MetroCluster requires two switches.
- You can use mixed switch types, but the switches must be the same type on each side of the configuration.
- Switches must be a supported Brocade model supplied by NetApp. Customer supplied switches are not supported.
- Switches must be running the correct firmware version.

### License requirements

- cluster
- syncmirror\_local
- cluster\_remote

### Related concepts

[Setup requirements and restrictions for standard active/active configurations](#) on page 24

[Setup requirements and restrictions for mirrored active/active configurations](#) on page 27

## Configuration limitations for fabric-attached MetroClusters

You must be aware of certain limitations when setting up a new fabric-attached MetroCluster configuration.

The fabric-attached MetroCluster configuration has the following limitations:

- SATA and AT-FCX storage is not supported.
- You cannot use the MetroCluster switches to connect Fibre Channel tape devices, or for Fibre Channel Protocol (FCP) traffic of any kind. You can connect only system controllers and disk shelves to the MetroCluster switches.
- You can connect a tape storage area network (SAN) to either of the nodes, but the tape SAN must not use the MetroCluster switches.

## Configuration variations for fabric-attached MetroClusters

Fabric-attached MetroClusters support asymmetrical and active/passive configurations.

The following list describes some common configuration variations that are supported for fabric-attached MetroClusters:

- Asymmetrical mirroring

You can add storage to one or both nodes that is not mirrored by the other node. However, any data contained in the unmirrored storage could be lost if that site experiences a disaster.

**Attention:** Multiple disk failures in an unmirrored aggregate (three or more disk failures in a RAID-DP aggregate, two or more disk failures in a RAID4 aggregate) will cause the node to panic, resulting in a temporary data service outage while the node reboots or disaster recovery is performed.

You must mirror the root volumes to enable successful takeover.

**Note:** You must connect the unmirrored storage to both nodes, just as for mirrored storage. You cannot have storage that is connected to only one node in an active/active configuration.

- Active/passive MetroClusters

In this configuration, the remote (passive) node does not serve data unless it has taken over for the local (active) node. Mirroring the passive node's root volume is optional. However, both nodes must have all MetroCluster licenses installed so that remote takeover is possible.

## MetroClusters and SnapLock volumes

As with any volume on a mirrored aggregate, on a properly configured site MetroCluster enables SnapLock volumes to be mirrored from one site to the other while retaining the SnapLock characteristics. If you issue the `cf forcetakeover -d` command because of a complete disaster or other operational failure at the primary site, these mirrors are broken and the mirror site goes online. Once the failed site is restored, the mirrors can be resynchronized before performing the giveback to normal operation.

**Attention:** If for any reason the primary node has data that was not mirrored to the secondary prior to the execution of the `cf forcetakeover -d` command, data could be lost. Do not resynchronize the original disks of the primary site for a SnapLock volume until an additional backup has been made of those disks to assure availability of all data. This situation could arise, for example, if the link between the sites was down and the primary node had data written to it in the interim before the `cf forcetakeover -d` command was issued.

For more information about backing up data in SnapLock volumes using SnapMirror, see the *Data ONTAP Archive and Compliance Management Guide*.





# Active/active configuration installation

---

To install and cable a new standard or mirrored active/active configuration you must have the correct tools and equipment and you must connect the controllers to the disk shelves (for FAS systems or V-Series systems using native disk shelves). You must also cable the cluster interconnect between the nodes. Active/active configurations can be installed in either NetApp system cabinets or in equipment racks.

## Next topics

[System cabinet or equipment rack installation](#) on page 41

[Required documentation, tools, and equipment](#) on page 42

[Preparing your equipment](#) on page 44

[Cabling nodes and DS14mk2 AT , DS14mk2 FC , or DS14mk4 FC disk shelves in standard or mirrored active/active configurations](#) on page 46

[Required connections for using uninterruptible power supplies with standard or mirrored active/active configurations](#) on page 57

## System cabinet or equipment rack installation

Describes the differences between an active/active configuration installed in a system cabinet or an equipment rack.

You need to install your active/active configuration in one or more NetApp system cabinets or in standard telco equipment racks.

## Next topics

[Active/active configurations in an equipment rack](#) on page 41

[Active/active configurations in a system cabinet](#) on page 42

## Active/active configurations 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.

The equipment 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 documentation that came with your equipment.

### Related concepts

[Cabling nodes and DS14mk2 AT , DS14mk2 FC , or DS14mk4 FC disk shelves in standard or mirrored active/active configurations](#) on page 46

## Active/active configurations in a system cabinet

If you ordered an active/active configuration in a system cabinet, it comes in one or more NetApp system cabinets, depending on the amount of storage.

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.

If the active/active configuration 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.

If the active/active configuration you ordered has more than six disk shelves, the active/active configuration 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 NVRAM cluster interconnects. If the active/active configuration uses switches, you must install the switches, as described in the accompanying switch documentation. The system cabinets might also need to be connected to each other. See your *System Cabinet Guide* for information about connecting your system cabinets together.

### Related concepts

[Cabling nodes and DS14mk2 AT , DS14mk2 FC , or DS14mk4 FC disk shelves in standard or mirrored active/active configurations](#) on page 46

## Required documentation, tools, and equipment

Describes the NetApp documentation and the tools required to install an active/active configuration.

### Next topics

[Required documentation](#) on page 42

[Required tools](#) on page 43

[Required equipment](#) on page 44

## Required documentation

Describes the flyers and guides required to install an active/active configuration.

NetApp hardware and service documentation is not contained within a single guide. Instead, the field-replaceable units are documented in separate flyers at the NOW site.

The following table lists and briefly describes the documentation you might need to refer to when preparing a new active/active configuration, or converting two stand-alone systems into an active/active configuration.

Manual name	Description
The appropriate system cabinet guide	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 to install NetApp equipment.
The appropriate disk shelf guide	These guides describe how to cable a disk shelf to a storage system.
The appropriate hardware documentation for your storage system model	These guides describe how to install the storage system, connect it to a network, and bring it up for the first time.
<i>Diagnostics Guide</i>	This guide describes the diagnostics tests that you can run on the storage system.
<i>Data ONTAP Upgrade Guide</i>	This guide describes how to upgrade storage system and disk firmware, and how to upgrade storage system software.
<i>Data ONTAP Data Protection Online Backup and Recovery Guide</i>	This guide describes, among other topics, SyncMirror technology, which is used for mirrored active/active configurations.
<i>Data ONTAP System Administration Guide</i>	This guide describes general storage system administration.
<i>Data ONTAP Software Setup Guide</i>	This guide describes how to configure the software of a new storage system for the first time.

**Note:** If you are installing a V-Series active/active configuration, refer also to the *V-Series Installation Requirements and Reference Guide* for information about cabling V-Series systems to storage arrays and to the V-Series Implementation Guides for information about configuring storage arrays to work with V-Series systems.

#### Related information

[Data ONTAP Information Library -](#)

[http://now.netapp.com/NOW/knowledge/docs/ontap/ontap\\_index.shtml](http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml)

## Required tools

Lists the tools you need to install the active/active configuration.

The following list specifies the tools you need to install the active/active configuration:

- #1 and #2 Phillips screwdrivers
- Hand level

- Marker

## Required equipment

When you receive your active/active configuration, you should receive the equipment listed in the following table. See the *System Configuration Guide* at the NOW site to confirm your storage system type, storage capacity, and so on.

Required equipment	Standard or mirrored active/active configuration
Storage system	Two of the same type of storage systems.
Storage	See the <i>System Configuration Guide</i> at the NOW site.
Cluster interconnect adapter	InfiniBand (IB) cluster adapter  (The NVRAM adapter functions as the cluster interconnect adapter on FAS900 series and later storage systems.)
FC-AL or FC HBA (FC HBA for Disk) adapters	Minimum of two FC-AL adapters
Fibre Channel switches	N/A
SFP (Small Form Pluggable) modules	N/A
NVRAM cluster adapter media converter	Only if using fiber cabling.
Cables (provided with shipment unless otherwise noted)	<ul style="list-style-type: none"> <li>• One optical controller-to-disk shelf cable per loop</li> <li>• Multiple disk shelf-to-disk shelf cables</li> <li>• Two 4xIB copper cables, or two 4xIB optical cables</li> </ul> <p><b>Note:</b> You must purchase longer optical cables separately for cabling distances greater than 30 meters.</p> <ul style="list-style-type: none"> <li>• Two optical cables with media converters for the IB cluster adapter</li> </ul>

## Preparing your equipment

You must install your nodes in your system cabinets or equipment racks, depending on your installation type.

Proceed to the appropriate section.

### Next topics

*Installing the nodes in equipment racks* on page 45

*Installing the nodes in a system cabinet* on page 45

## Installing the nodes in equipment racks

Before you cable your nodes together, you install the nodes and disk shelves in the equipment rack, label the disk shelves, and connect the nodes to the network.

### Steps

1. Install the nodes in the equipment rack, as described in the Fibre Channel disk shelf guide, hardware documentation, 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 setup instructions for your system.

The nodes are now in place and connected to the network and power is available.

### After you finish

Proceed to cable the active/active configuration.

## Installing the nodes in a system cabinet

Before you cable your nodes together, you must install the system cabinet and connect the nodes to the network. If you have two cabinets, the cabinets must be connected together.

### Steps

1. Install the system cabinets, as described in the *System Cabinet Guide*. If you have multiple system cabinets, remove the front and rear doors and any side panels that need to be removed, and connect the system cabinets together.
2. Connect the nodes to the network.
3. Connect the system cabinets to an appropriate power source and apply power to the cabinets.

The nodes are now in place and connected to the network and power is available.

### After you finish

Proceed to cable the active/active configuration.

## Cabling nodes and DS14mk2 AT, DS14mk2 FC, or DS14mk4 FC disk shelves in standard or mirrored active/active configurations

To cable a standard active/active configuration or a mirrored configuration in which you configure both channels of the disk shelves, you must identify the Fibre Channel port usage and cable the disk shelves to each node and cable the HA interconnect between the nodes. The specific procedure depends on whether you have a standard or mirrored active/active configuration.

**Note:** If your configuration includes DS4243 disk shelves, refer to the *DS4243 System Connectivity Guide* and the *DS4243 Hardware Service Guide* on the NOW site at <http://now.netapp.com/>.

This guide does not specify the required slot locations for the various adapters you use to cable your active/active configuration. See the *System Configuration Guide* at [http://now.netapp.com/NOW/knowledge/docs/hardware/hardware\\_index.shtml](http://now.netapp.com/NOW/knowledge/docs/hardware/hardware_index.shtml) to obtain all slot assignment information.

### Next topics

[Systems with two controllers in the same chassis](#) on page 46

[Which Fibre Channel ports to use for DS14mk2 or DS14mk4 FC disk shelf connections](#) on page 51

[Cabling a standard active/active configuration](#) on page 47

[Cabling a mirrored active/active configuration](#) on page 50

## Systems with two controllers in the same chassis

In an active/active configuration, some storage systems (such as the FAS20xx and 31xx systems) support two controller modules in the same chassis.

This simplifies cabling of these systems, because they use an internal InfiniBand connector between the two controller modules, so no interconnect adapters or cabling is required.

This is different from the examples in this section, which show systems that must be cabled together with an interconnect to enable the active/active configuration.

## Which Fibre Channel ports to use for DS14mk2 or DS14mk4 FC disk shelf connections

Before cabling your active/active configuration, you need to identify which Fibre Channel ports to use to connect your disk shelves to each storage system, and in what order to connect them.

Keep the following guidelines in mind when identifying ports to use:

- Every disk shelf loop in the active/active configuration requires a port on the node.  
A standard active/active configuration with one loop for each node uses two ports on each node.

- Onboard Fibre Channel ports should be used before using ports on expansion adapters.
- Always use the expansion slots in the order shown in the *System Configuration Guide* at the NOW site for your platform for an active/active configuration .
- If using Fibre Channel HBAs, insert the adapters in the same slots on both systems.

When complete, you should have a numbered list of Fibre Channel ports for both nodes, starting with Port 1.

### Cabling with a quad-port Fibre Channel HBA

If using ports on the quad-port, 4-Gb Fibre Channel HBAs, use the procedures in the following sections, with the following additional guidelines:

- Cable disk shelf loops using ESH4 modules to the quad-port HBA first.
- Cable disk shelf loops using AT-FCX, ESH, or ESH2 modules to dual-port HBA ports or onboard ports before using ports on the quad-port HBA.
- Connect port A of the HBA to the In port of Channel A of the first disk shelf in the loop.
- Connect the same port (port A) of the HBA in the partner node to the In port of Channel B of the first disk shelf in the loop. This ensures that disk names are the same for both nodes.
- Cable additional disk shelf loops sequentially with the HBA's ports. Use port A for the first loop, then B, and so on.
- If available, use ports C or D for multipathing after cabling all remaining disk shelf loops.
- Observe all other cabling rules described in the documentation for the HBA, and the *System Configuration Guide*.

## Cabling a standard active/active configuration

To cable a standard active/active configuration, you identify the ports you need to use on each node, then you cable the ports, and then you cable the cluster interconnect.

### About this task

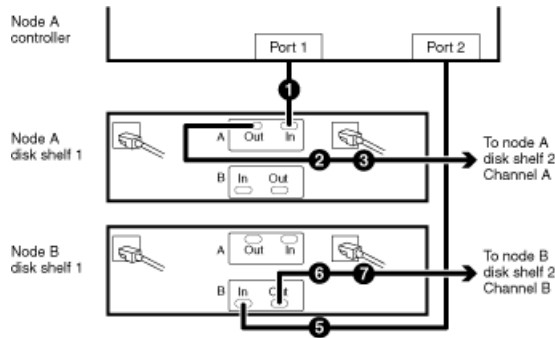
Complete the following tasks in the order shown.

**Note:** If you are using DS4243, refer to the *DS4243 Connectivity Guide* for information on cabling the disk shelves in an active/active configuration.

1. [Cabling Node A to the disk shelves](#) on page 47
2. [Cabling Node B to the disk shelves](#) on page 49
3. [Cabling the cluster interconnect for a standard active/active configuration](#) on page 50

### Cabling Node A to the disk shelves

To cable Node A, you must use the Fibre Channel ports you previously identified and cable the disk shelf loops owned by the node to these ports.



**Figure 6: Cabling Node A**

**Note:** The circled numbers in the diagram correspond to the step numbers in the procedure, and the ports (Port 1, Port 2) correspond to the Fibre Channel ports you are using.

The location of the Input and Output ports on the disk shelves vary depending on the disk shelf model. Make sure that you refer to the labeling on the disk shelf rather than to the location of the port shown in the diagram.

### Steps

1. Cable Fibre Channel port 1 of Node A to the Channel A Input port of the first disk shelf of Node A loop 1.
2. Cable the Node A disk shelf Channel A Output port to the Channel A Input port of the next disk shelf in loop 1.
3. Repeat Step 2 for any remaining disk shelves in loop 1.
4. If your disk shelf modules have terminate switches, set the terminate switches to Off on all but the last disk shelf in loop 1, and set the terminate switch on the last disk shelf to On.
5. Cable Fibre Channel port 2 of Node A to the Channel B Input port of the first disk shelf of Node B loop 1.
6. Cable the disk shelf Channel B Output port to the Channel B Input port of the next disk shelf in loop 1.
7. Repeat Step 6 for any remaining disk shelves in loop 1.
8. If your disk shelf modules have terminate switches, set the terminate switches to Off on all but the last disk shelf in loop 1, and set the terminate switch on the last disk shelf to On.
9. Repeat Step 1 to Step 8 for each pair of loops in the active/active configuration, using ports 3 and 4 for the next loop, ports 5 and 6 for the next one, and so on.

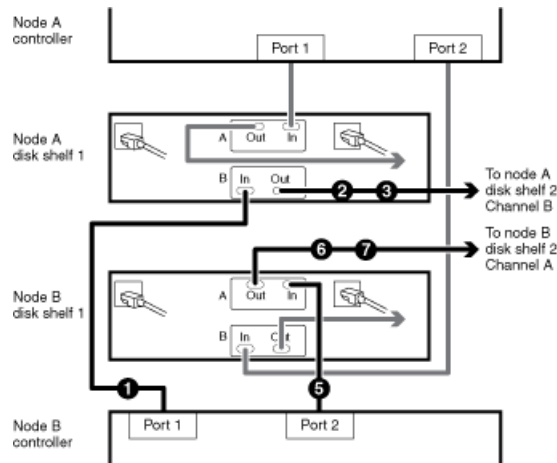
Node A is completely cabled.

Proceed to cabling Node B.



## Cabling Node B to the disk shelves

To cable Node B, you must use the Fibre Channel ports you previously identified and cable the disk shelf loops owned by the node to these ports.



**Figure 7: Cabling Node B**

### Steps

1. Cable Port 1 of Node B to the Channel B Input port of the first disk shelf of Node A loop 1.
 

Both channels of this disk shelf are connected to the same port on each node. This is not required, but it makes your active/active configuration easier to administer because the disks have the same ID on each node. This is true for Step 5 also.
2. Cable the disk shelf Channel B Output port to the Channel B Input port of the next disk shelf in loop 1.
3. Repeat Step 2 for any remaining disk shelves in loop 1.
4. If your disk shelf modules have terminate switches, set the terminate switches to Off on all but the last disk shelf in loop 1, and set the terminate switch on the last disk shelf to On.
5. Cable Fibre Channel port 2 of Node B to the Channel A Input port of the first disk shelf of Node B loop 1.
6. Cable the disk shelf Channel A Output port to the Channel A Input port of the next disk shelf in loop 1.
7. Repeat Step 6 for any remaining disk shelves in loop 1.
8. If your disk shelf modules have terminate switches, set the terminate switches to Off on all but the last disk shelf in loop 1, and set the terminate switch on the last disk shelf to On.
9. Repeat Step 1 to Step 8 for each pair of loops in the active/active configuration, using ports 3 and 4 for the next loop, ports 5 and 6 for the next one, and so on.

Node B is completely cabled.

Proceed to cable the cluster interconnect.

## Cabling the cluster interconnect for a standard active/active configuration

To cable the interconnect between the active/active configuration nodes, you must make sure that your interconnect adapter is in the correct slot and connect the adapters on each node with the optical cable.

### Steps

1. See the *System Configuration Guide* at the NOW site to ensure that your interconnect adapter is in the correct slot for your system in an active/active configuration.

For systems that use an NVRAM5 or NVRAM6 adapter, the NVRAM adapter functions as the cluster interconnect adapter.

2. Plug one end of the optical cable into one of the local node's cluster adapter ports, then plug the other end into the partner node's corresponding adapter port.

You must not cross-cable the cluster interconnect adapter. Cable the local node ports only to the identical ports on the partner node.

If the system detects a cross-cabled cluster interconnect, the following message appears:

```
Cluster interconnect port <port> of this appliance seems to be connected
to port <port> on the partner appliance.
```

3. Repeat Step 2 for the two remaining ports on the cluster adapters.

The nodes are connected to each other.

Proceed to configure the system.

## Cabling a mirrored active/active configuration

To cable a mirrored active/active configuration, you identify the ports you need to use on each node, and then you cable the ports, and then you cable the cluster interconnect.

### About this task

Complete the following tasks in the order shown:

1. [Disk pool requirements for mirrored active/active configurations](#) on page 51
2. [Which Fibre Channel ports to use for DS14mk2 or DS14mk4 FC disk shelf connections](#) on page 51
3. [Creating your port list for mirrored active/active configurations](#) on page 52
4. [Key to the mirrored active/active configuration diagrams](#) on page 52
5. [Cabling the channel A disk shelf loops](#) on page 53
6. [Cabling the channel B disk shelf loops](#) on page 54
7. [Cabling the cluster interconnect for a mirrored active/active configuration](#) on page 56

## Disk pool requirements for mirrored active/active configurations

Mirrored active/active configurations use SyncMirror to separate each aggregate into two plexes that mirror each other. One plex uses disks in pool 0 and the other plex uses disks in pool 1. To ensure proper disk pool access, your cabling depends on whether you have hardware-based or software-based disk ownership.

If your system uses hardware-based disk ownership, you must cable your mirrored active/active configuration according to the pool rules for your platform. For more information about pool rules, see the section on hardware-based disk ownership in the *Data ONTAP Storage Management Guide*.

If your system uses software-based disk ownership, follow the guidelines for software-based disk ownership in the *Data ONTAP Storage Management Guide*.

For more information about SyncMirror, see the *Data ONTAP Data Protection Online Backup and Recovery Guide*.

## Which Fibre Channel ports to use for DS14mk2 or DS14mk4 FC disk shelf connections

Before cabling your active/active configuration, you need to identify which Fibre Channel ports to use to connect your disk shelves to each storage system, and in what order to connect them.

Keep the following guidelines in mind when identifying ports to use:

- Every disk shelf loop in the active/active configuration requires a port on the node.  
A standard active/active configuration with one loop for each node uses two ports on each node.
- Onboard Fibre Channel ports should be used before using ports on expansion adapters.
- Always use the expansion slots in the order shown in the *System Configuration Guide* at the NOW site for your platform for an active/active configuration .
- If using Fibre Channel HBAs, insert the adapters in the same slots on both systems.

When complete, you should have a numbered list of Fibre Channel ports for both nodes, starting with Port 1.

## Cabling with a quad-port Fibre Channel HBA

If using ports on the quad-port, 4-Gb Fibre Channel HBAs, use the procedures in the following sections, with the following additional guidelines:

- Cable disk shelf loops using ESH4 modules to the quad-port HBA first.
- Cable disk shelf loops using AT-FCX, ESH, or ESH2 modules to dual-port HBA ports or onboard ports before using ports on the quad-port HBA.
- Connect port A of the HBA to the In port of Channel A of the first disk shelf in the loop.
- Connect the same port (port A) of the HBA in the partner node to the In port of Channel B of the first disk shelf in the loop. This ensures that disk names are the same for both nodes.

- Cable additional disk shelf loops sequentially with the HBA’s ports. Use port A for the first loop, then B, and so on.
- If available, use ports C or D for multipathing after cabling all remaining disk shelf loops.
- Observe all other cabling rules described in the documentation for the HBA, and the *System Configuration Guide*.

### Creating your port list for mirrored active/active configurations

After you determine the Fibre Channel ports to use, you create a table identifying which ports belong to which port pool.

#### Step

1. Create a table specifying the port usage; the cabling diagrams in this document use the notation “P1-3” (the third port for pool 1).

#### Example

For a 30xx active/active configuration that has two mirrored loops using hardware-based disk ownership, the port list look like the following example:

Pool 0	Pool 1
P0-1: onboard port 0a	P1-1: onboard port 0c
P0-2: onboard port 0b	P1-2: onboard port 0d
P0-3: slot 2 port A	P1-3: slot 4 port A
P0-4: slot 2 port B	P1-4: slot 4 port B

Proceed to cable the Channel A loops.

### Key to the mirrored active/active configuration diagrams

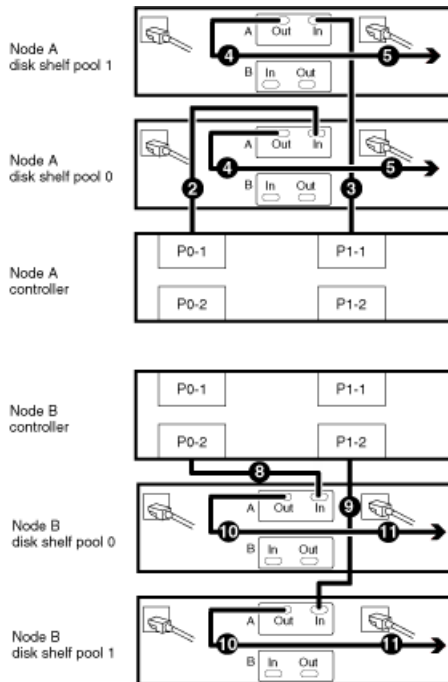
You should review these facts before using the diagrams for cabling your mirrored active/active configuration.

- The circled numbers in the diagram correspond to the step numbers in the procedure.
- The location of the Input and Output ports on the disk shelves vary depending on the disk shelf models. Make sure that you refer to the labeling on the disk shelf rather than to the location of the port shown in the diagram.
- The location of the Fibre Channel ports on the controllers is not representative of any particular storage system model; determine the locations of the ports you are using in your configuration by inspection or by using the Installation and Setup Instructions for your model.

- The port numbers refer to the list of Fibre Channel ports you created.
- The diagram only shows one loop per node and one disk shelf per loop. Your installation might have more loops, more disk shelves, or different numbers of disk shelves between nodes.

### Cabling the channel A disk shelf loops

To begin cabling of the disk shelves, you cable the appropriate pool ports on the node to the Channel A modules of the disk shelf stack for the pool.



**Figure 8: Channel A loops for a mirrored active/active configuration**

#### Steps

1. Complete your port list.
2. Cable channel A for Node A.
  - a. Cable the first port for pool 0 (P0-1) of Node A to the first Node A disk shelf Channel A Input port of disk shelf pool 0.
  - b. Cable the first port for pool 1 (P1-1) of Node A to the first Node A disk shelf Channel A Input port of disk shelf pool 1.
  - c. Cable the disk shelf Channel A Output port to the next disk shelf Channel A Input port in the loop for both disk pools.

**Note:** The illustration shows only one disk shelf per disk pool. The number of disk shelves per pool might be different for your configuration.

- d. Repeat substep 2c, connecting Channel A output to input, for any remaining disk shelves in this loop for each disk pool.
- e. If your disk shelf modules have terminate switches, set the terminate switches to Off on all but the last disk shelf in the loop, and set the terminate switch on the last disk shelf to On.

**Note:** The ESH2 and ESH4 are self-terminating and do not have a terminate switch. The AT-FCX is self-terminating if no cable is plugged into the Output port of the last disk shelf. This applies to Step 3 Substep e also.

- f. Repeat Substep a through Substep e for any additional loops for Channel A, Node A, using the odd numbered port numbers (P0-3 and P1-3, P0-5 and P1-5, and so on).

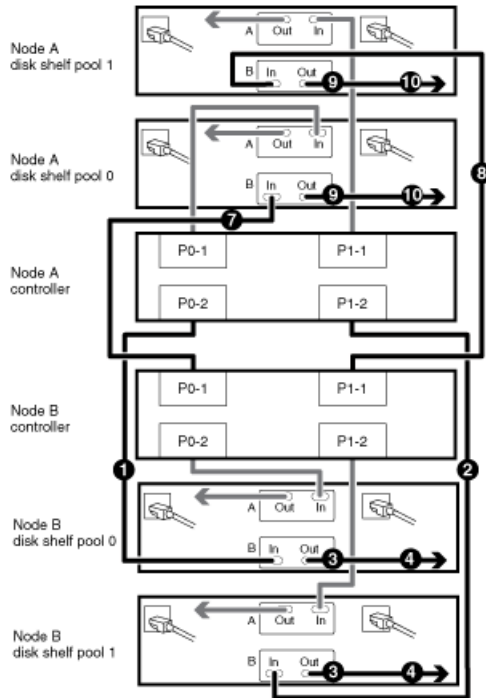
### 3. Cable Channel A for Node B

- a. Cable the second port for pool 0 (P0-2) of Node B to the first Node B disk shelf Channel A Input port of disk shelf pool 0.
- b. Cable the second port for pool 1 (P1-2) of Node B to the first Node B disk shelf Channel A Input port of disk shelf pool 1.
- c. Cable the disk shelf Channel A Output port to the next disk shelf Channel A Input port in the loop for both disk pools.
- d. Repeat substep 3.c, connecting Channel A output to input, for any remaining disk shelves in each disk pool.
- e. If your disk shelf modules have terminate switches, set the terminate switches to Off on all but the last disk shelf in the loop, and set the terminate switch on the last disk shelf to On.
- f. Repeat substep 3.a through substep 3.e for any additional loops on Channel A, Node B, using the even numbered port numbers (P0-4 and P1-4, P0-6 and P1-6, and so on).

Proceed to cable the Channel B loops.

### Cabling the channel B disk shelf loops

To provide the mirrored storage, you cable the mirrored pool ports on the node to the Channel B modules of the appropriate disk shelf stack.



**Figure 9: Channel B loops for a mirrored active/active configuration**

### Steps

#### 1. Cable Channel B for Node A

- a. Cable the second port for pool 0 (P0-2) of Node A to the first Node B disk shelf Channel B Input port of disk shelf pool 0.

**Note:** Both channels of this disk shelf are connected to the same port on each node. This is not required, but it makes your active/active configuration easier to administer because the disks have the same ID on each node.

- b. Cable the second port for pool 1 (P1-2) of Node A to the first Node B disk shelf Channel B Input port of disk shelf pool 1.
- c. Cable the disk shelf Channel B Output port to the next disk shelf Channel B Input port in the loop for both disk pools.

**Note:** The illustration shows only one disk shelf per disk pool. The number of disk shelves per pool might be different for your configuration.

- d. Repeat Substep c, connecting Channel B output to input, for any remaining disk shelves in each disk pool.
- e. If your disk shelf modules have terminate switches, set the terminate switches to Off on all but the last disk shelf in the loop, and set the terminate switch on the last disk shelf to On.

**Note:** The ESH2 and ESH4 are self-terminating and do not have a terminate switch. The AT-FCX is self-terminating if no cable is plugged into the Output port of the last shelf. This note applies to Step 2 Substep e also.

- f. Repeat Substep a through Substep e for any additional loops on Channel B, Node A, using the even numbered port numbers (P0-4 and P1-4, P0-6 and P1-6, and so on).
2. Cable Channel B for Node B
    - a. Cable the first port for pool 0 (P0-1) of Node B to the first Node A disk shelf Channel B Input port of disk shelf pool 0.
    - b. Cable the first port for pool 1 (P1-1) of Node B to the first Node A disk shelf Channel B Input port of disk shelf pool 1.
    - c. Cable the disk shelf Channel B Output port to the next disk shelf Channel B Input port in the loop for both disk pools.
    - d. Repeat Substep c, connecting Channel B output to input, for any remaining disk shelves in each disk pool.
    - e. If your disk shelf modules have terminate switches, set the terminate switches to Off on all but the last disk shelf in the loop, and set the terminate switch on the last disk shelf to On.
    - f. Repeat Substep a through Substep e for any additional loops for Channel B, Node B, using the odd numbered port numbers (P0-3 and P1-3, P0-5 and P1-5, and so on).

Proceed to cable the cluster interconnect.

## Cabling the cluster interconnect for a mirrored active/active configuration

To cable the cluster interconnect between the active/active configuration nodes, you must make sure that your interconnect adapter is in the correct slot and connect the adapters on each node with the optical cable.

### Steps

1. See the *System Configuration Guide* at the NOW site to ensure that your cluster interconnect adapter is in the correct slot for your system in an active/active configuration .

**Note:** For systems that use an NVRAM5 or NVRAM6 adapter, the NVRAM adapter functions as the cluster interconnect adapter.

2. Plug one end of the optical cable into one of the local nodes cluster adapter ports, then plug the other end into the partner nodes corresponding adapter port.

You must not cross-cable the interconnect adapter. Cable the local node ports only to the identical ports on the partner node.

If the system detects a cross-cabled interconnect, the following message appears:

```
Cluster interconnect port <port> of this appliance seems to be connected
to port <port> on the partner appliance.
```



3. Repeat Step 2 for the remaining ports on the cluster adapters.

Proceed to configure the active/active configuration .

## **Required connections for using uninterruptible power supplies with standard or mirrored active/active configurations**

You can use a UPS (uninterruptible power supply) with your active/active configuration. The UPS enables the system to fail over gracefully if power fails for one of the nodes, or to shut down gracefully if power fails for both nodes. You must ensure that the correct equipment is connected to the UPS.

To gain the full benefit of the UPS, you must ensure that all the required equipment is connected to the UPS. The equipment that needs to be connected depends on whether your configuration is a standard or a mirrored active/active configuration.

### **Equipment to connect to the UPS for standard active/active configurations**

For a standard active/active configuration, you must connect the controller, disks, and any Fibre Channel switches in use.

### **Equipment to connect to the UPS for mirrored active/active configurations**

For a mirrored active/active configuration, you must connect the controller and any Fibre Channel switches to the UPS, as for a standard active/active configuration. However, if the two sets of disk shelves have separate power sources, you do not have to connect the disks to the UPS. If power is interrupted to the local controller and disks, the controller can access the remote disks until it shuts down gracefully or the power supply is restored. In this case, if power is interrupted to both sets of disks at the same time, the active/active configuration cannot shut down gracefully.



# MetroCluster installation

---

You can install a stretch or fabric-attached MetroCluster to provide complete data mirroring and takeover capabilities if a site is lost in a disaster. Fabric-attached MetroClusters provide active/active configuration with physically separated nodes at a greater distance than that provided by stretch MetroCluster.

**Note:** If you are a V-Series system customer, see the *V-Series MetroCluster Guide* for information about configuring and operating a V-Series system in a MetroCluster configuration.

## Next topics

[Required documentation, tools, and equipment](#) on page 59

[MetroCluster and software-based disk ownership](#) on page 62

[Converting an active/active configuration to a fabric-attached MetroCluster](#) on page 63

[Upgrading an existing MetroCluster](#) on page 65

[Cabling a stretch MetroCluster](#) on page 67

[Cabling a stretch MetroCluster between 31xx systems](#) on page 67

[Cabling a fabric-attached MetroCluster](#) on page 68

[Required connections for using uninterruptible power supplies with MetroCluster configurations](#) on page 93

## Related concepts

[Disaster recovery using MetroCluster](#) on page 173

[Setup requirements and restrictions for stretch MetroCluster configurations](#) on page 32

[Setup requirements and restrictions for fabric-attached MetroClusters](#) on page 36

## Required documentation, tools, and equipment

Describes the NetApp documentation and the tools required to install a MetroCluster configuration.

### Next topics

[Required documentation](#) on page 59

[Required tools](#) on page 61

[Required equipment](#) on page 61

## Required documentation

Describes the flyers and guides required to install a new MetroCluster, or convert two stand-alone systems into a MetroCluster.

NetApp hardware and service documentation is not contained within a single guide. Instead, the field-replaceable units are documented in separate flyers at the NOW site.

The following table lists and briefly describes the documentation you might need to refer to when preparing a new MetroCluster configuration, or converting two stand-alone systems into a MetroCluster configuration.

<b>Manual name</b>	<b>Description</b>
The appropriate system cabinet guide	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 to install NetApp equipment.
The appropriate disk shelf guide	These guides describe how to cable a disk shelf to a storage system.
The appropriate hardware documentation for your storage system model	These guides describe how to install the storage system, connect it to a network, and bring it up for the first time.
<i>Diagnostics Guide</i>	This guide describes the diagnostics tests that you can run on the storage system.
<i>Upgrade Guide</i>	This guide describes how to upgrade storage system and disk firmware, and how to upgrade storage system software.
<i>Data Protection Online Backup and Recovery Guide</i>	This guide describes, among other topics, SyncMirror technology, which is used for mirrored Active/active configurations.
<i>Data ONTAP System Administration Guide</i>	This guide describes general storage system administration.
<i>Software Setup Guide</i>	This guide describes how to configure the software of a new storage system for the first time.
<i>Brocade Switch Configuration Guide for Fabric MetroCluster</i>	This document describes how to configure Brocade switches for a fabric-attached MetroCluster.  You can find this document on the Fabric-Attached MetroCluster Switch Description page on the NOW site.
The appropriate Brocade manuals	These guides describe how to configure and maintain Brocade switches.  These guides are available from the Brocade Switch Description Page at the NOW site.

#### **Related information**

*Brocade Switch Configuration Guide for Fabric MetroCluster -*

[http://now.netapp.com/NOW/download/software/sanswitch/fcp/Brocade/mc\\_ontap641\\_fabric\\_index.shtml](http://now.netapp.com/NOW/download/software/sanswitch/fcp/Brocade/mc_ontap641_fabric_index.shtml)

*Data ONTAP Information Library -*

[http://now.netapp.com/NOW/knowledge/docs/ontap/ontap\\_index.shtml](http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml)

## Required tools

Lists the tools you need to install the active/active configuration.

The following list specifies the tools you need to install the MetroCluster configuration:

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

## Required equipment

When you receive your MetroCluster, you should receive the equipment listed in the following table. See the *System Configuration Guide* at the NOW site to confirm your storage system type, storage capacity, and so on.

**Note:** For fabric-attached MetroClusters, use the information in the *System Configuration Guide* labeled for MetroClusters. For stretch MetroClusters, use the information in the *System Configuration Guide* labeled “for HA Environments.”

Required equipment	Stretch MetroCluster	Fabric-attached MetroCluster
Storage system	Two of the same type of storage systems.	
Storage	See the <i>System Configuration Guide</i> at <a href="http://now.netapp.com/NOW/knowledge/docs/hardware/hardware_index.shtml">http://now.netapp.com/NOW/knowledge/docs/hardware/hardware_index.shtml</a> .	
cluster interconnect adapter	<p>Infiniband adapter (Required only for systems that do not use an NVRAM5 or NVRAM6 adapter, which functions as the cluster interconnect adapter.)</p> <p>FC-VI adapter (Required only for the 31xx dual-controller systems.)</p> <p><b>Note:</b> When the FC-VI adapter is installed in a 31xx system, the internal InfiniBand interconnect is automatically deactivated.</p>	FC-VI adapter
FC-AL or FC HBA (FC HBA for Disk) adapters	<p>Two or four Fibre Channel HBAs. These HBAs are required for 4-Gbps MetroCluster operation. Onboard ports can be used for 2-Gbps operation.</p> <p><b>Note:</b> The ports on the Fibre Channel HBAs 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.</p>	

Required equipment	Stretch MetroCluster	Fabric-attached MetroCluster
Fibre Channel switches	N/A	Two pairs of Brocade switches  <b>Note:</b> The Fibre Channel switches must be of the same type. A mixture of switch types (such as Brocade 300 and Brocade 5100 switches) is not allowed.
SFP (Small Form Pluggable) modules	N/A	Two or four long-distance for inter-switch links, depending on whether you are using dual inter-switch links. The type of SFP needed depends on the distance between sites.  One short-distance for each switch port used.
NVRAM adapter media converter	Only if using fiber cabling.	N/A
Cables (provided with shipment unless otherwise noted)	<ul style="list-style-type: none"> <li>Four SC/LC (standard connector to low-profile connector) controller-to-disk shelf cables</li> <li>Two SC/LC IB cluster adapter cables</li> <li>Four SC/LC or LC/LC cables</li> </ul> <p><b>Note:</b> For information about required cables, see the MetroCluster Compatibility Matrix on the NOW site.</p>	<ul style="list-style-type: none"> <li>LC/LC controller-to-switch cables</li> <li>SC/LC (for DS14) or LC/LC (for DS14mk2 FC) disk shelf-to-switch cables</li> <li>Two LC/LC inter-switch link cables, not provided in the shipment</li> <li>Multiple disk shelf-to-disk shelf cables</li> </ul>

## MetroCluster and software-based disk ownership

Systems using software-based disk ownership in a MetroCluster require different configuration than systems using hardware-based disk ownership.

Some systems use software-based disk ownership to control which disks in a disk shelf loop belong to which controller and pool.

- Software commands in Data ONTAP are used to assign disks, or they are auto-assigned by the software.

This is because disk ownership is determined by the software, rather than by the physical cabling of the shelves.

- Systems that use software disk ownership require different cabling of their disk shelves when you configure your MetroCluster.

This is because different Brocade port usage rules are used with software-based disk ownership.

For details about software-based disk ownership, see the *Data ONTAP Storage Management Guide*.

### The 4-Gbps FC-VI adapter requires software disk ownership

If you want to take advantage of the performance provided by the 4-Gbps adapter, you must upgrade to a system that uses software-based disk ownership.

## Converting an active/active configuration to a fabric-attached MetroCluster

With the correct hardware, you can reconfigure an active/active configuration to a fabric-attached MetroCluster.

### Before you begin

- If you are upgrading an existing active/active configuration to a MetroCluster configuration, you must upgrade disk firmware to the latest version. After upgrading disk firmware, you must power-cycle the affected disk drives to ensure that they work correctly in a fabric-attached MetroCluster. You can download the latest disk firmware from <http://now.netapp.com/>.
- If you are upgrading from an existing active/active configuration on a system that supports both software-based and hardware-based disk ownership and is currently using software-based disk ownership, you must convert disk assignment to hardware ownership before the upgrade. However, converting an active/active configuration from software-based disk ownership to hardware-based disk ownership is a complicated process. If done incorrectly, your system might not boot. You are advised to contact technical support for assistance with this conversion.
- If you are upgrading a 31xx system, the resulting upgraded system can only have one controller in each chassis. If you have a chassis with two controllers, you must move one controller to a new chassis to form the partner node of the MetroCluster. You must also obtain and install the FC-VI interconnect card on both systems.

**Note:** For details about this conversion process, see TR-3517, *MetroCluster Upgrade Planning Guide*, on [now.netapp.com](http://now.netapp.com).

### Steps

1. Update Data ONTAP, storage system firmware, and disk firmware, as described in the *Data ONTAP Upgrade Guide*, making sure to shut down the nodes to the boot prompt.

2. Remove any ATA drives in the configuration.

ATA drives are not supported in a MetroCluster configuration.

3. Move the NVRAM adapter and FC-VI adapter to the correct slots for your model, as shown by the *System Configuration Guide* at the NOW site.
4. Determine your switch and general configuration by completing the planning worksheet.
5. Set up and configure the local switches, and verify your switch licenses, as described in the *Brocade Switch Configuration Guide for Fabric-attached MetroClusters* *Brocade Switch Configuration Guide for Fabric MetroCluster*.

You can find this document on the Brocade Switch Description Page on the NOW site.

**Note:** The configuration and firmware requirements for Brocade switches in a MetroCluster environment are different from the requirements for switches used in SAN environments. Always refer to MetroCluster documentation when installing and configuring your MetroCluster switches:

- The MetroCluster Compatibility Matrix
- The Brocade Switch Description Page
- The *Brocade Switch Configuration Guide for Fabric-attached MetroClusters*

6. Cable the local node.
7. Install the Data ONTAP licenses in the following order:
  - a. cluster
  - b. syncmirror\_local
  - c. cluster\_remote
8. Configure the local node depending on the type of active/active configuration:

If you are converting a...	Then...
<b>Standard active/active configuration</b>	Set up mirroring and configure the local node.
<b>Stretch MetroCluster</b>	Configure the local node.

9. Transport the partner node, disk shelves, and switches to the remote location.
10. Set up the remote node, disk shelves, and switches.

**After you finish**

Configure the MetroCluster.

**Related concepts**

- [Configuring an active/active configuration](#) on page 103
- [Disaster recovery using MetroCluster](#) on page 173



### Related tasks

[Cabling Node A](#) on page 72

[Cabling Node B](#) on page 81

[Disabling the `change\_fsid` option in MetroCluster configurations](#) on page 108

[Planning the fabric-attached MetroCluster installation](#) on page 69

## Upgrading an existing MetroCluster

You can upgrade an existing MetroCluster on a system using hardware-based disk ownership to a MetroCluster on a system using software-based disk ownership (FAS3040, FAS3070, or FAS60xx systems). This is useful when you are upgrading to 4-Gbps cluster interconnect support, which requires software-based disk ownership.

### About this task

When using the typical hardware upgrade procedure you upgrade your software on the old system and then use the `disk upgrade_ownership` command to apply software-based ownership to the disks. You then perform the hardware upgrade.

In the following procedure, you perform the hardware upgrade prior to using the `disk upgrade_ownership` command. This is because the old system hardware does not support the new features of the `disk upgrade_ownership` command. For the Data ONTAP 7.2.3 (or later) version of the `disk upgrade_ownership` command to run successfully, you must issue it on a system that supports software-based disk ownership.

### Steps

1. Halt the system, and then turn off the controller and disk shelves.
2. Remove the existing controller from the rack or system cabinet and install the FAS3040, FAS3070, or FAS60xx system in its place.

When replacing the controllers, use the same cabling to the Brocade switches and the disk shelves as the original controller. For the upgrade to work, you must retain the original cabling until you run the `disk upgrade_ownership` command later in this procedure.

3. Power on the disk shelves.
4. To reassign disk ownership to software-based disk ownership, complete the following substeps on both controllers:
  - a. Power on the system and boot the system into Maintenance mode.

For more information, see the *Data ONTAP System Administration Guide*.

- b. Enter the following command at the firmware prompt:

```
disk upgrade_ownership
```

This command converts the system to software-based disk ownership. Data ONTAP assigns all the disks to the same system and pool that they were assigned to for the hardware-based disk ownership.

See the *Data ONTAP Storage Management Guide* for detailed information about software-based disk ownership.

5. Verify disk ownership information by entering the following command:

```
disk show -v
```

Disk assignment is now complete.

6. Clear the mailboxes by entering the following commands:

```
mailbox destroy local
```

```
mailbox destroy partner
```

7. Enter the following command to exit Maintenance mode:

```
halt
```

8. Enter the following command for each required license:

#### Example

```
license add xxxxxx
```

xxxxx is the license code you received for the feature.

9. Enter the following command to reboot the node:

```
reboot
```

10. Configure the RLM, if applicable, as described in the *Data ONTAP Software Setup Guide*.

11. Recable the connections to the Brocade switches to conform to the virtual channel rules for the switch.

### After you finish

Configure the MetroCluster.

### Related concepts

[Configuring an active/active configuration](#) on page 103

[Disaster recovery using MetroCluster](#) on page 173

[Switch bank rules and virtual channel rules](#) on page 71

### Related tasks

[Cabling Node A](#) on page 72

[Cabling Node B](#) on page 81

[Disabling the change\\_fsid option in MetroCluster configurations](#) on page 108

## Cabling a stretch MetroCluster

The process to cable a stretch MetroCluster is the same as a mirrored active/active configuration. However, your systems must meet the requirements for a stretch MetroCluster.

### Related concepts

[Configuring an active/active configuration](#) on page 103

[Setup requirements and restrictions for stretch MetroCluster configurations](#) on page 32

[Disaster recovery using MetroCluster](#) on page 173

### Related tasks

[Cabling a mirrored active/active configuration](#) on page 50

## Cabling a stretch MetroCluster between 31xx systems

If you are configuring a stretch MetroCluster between 31xx systems, you must configure FC-VI interconnect adapter connections between the controllers.

### Steps

1. Connect port A of the FC-VI adapter on the controller of the local site to port A of the corresponding FC-VI adapter at the remote site.
2. Connect port B of the FC-VI adapter on the controller of the local site to port B of the corresponding FC-VI adapter at the remote site.
3. Cable the disk shelf loops as described in the procedure for cabling a mirrored active/active configuration.

### Related concepts

[Stretch MetroCluster configuration on 31xx systems](#) on page 31

### Related tasks

[Cabling a mirrored active/active configuration](#) on page 50

## Cabling a fabric-attached MetroCluster

You cable the fabric-attached MetroCluster so that the controller and the disk shelves at each site are connected to Brocade switches. In turn, the Brocade switches at one site are connected through inter-switch links to the Brocade switches at the other site.

### Before you begin

To cable a fabric-attached MetroCluster, you must be familiar with active/active configurations, the Brocade command-line interface, and synchronous mirroring. You must also be familiar with the characteristics of fabric-attached MetroClusters. You must also have the following information:

- Correct Brocade licenses for each switch
- Unique domain IDs for each of the switches

**Note:** You can use the switch numbers (1, 2, 3, and 4) as the switch Domain ID.


- Ethernet IP address for both the switches and nodes

**Note:** The switches ship with a default IP address (10.77.77.77), which you can use if the switches are not attached to a network.

- Ethernet subnetmask
- Gateway address

### About this task

A fabric-attached MetroCluster involves two nodes at physically separated sites. To differentiate these nodes in this documentation, the guide refers to the two nodes as Node A and Node B.



**Figure 10: Node A and Node B**

Complete the following tasks in the order shown:

1. [Planning the fabric-attached MetroCluster installation](#) on page 69
2. [Configuration differences for fabric-attached MetroClusters on 31xx systems](#) on page 70
3. [Configuring the switches](#) on page 70
4. [Cabling Node A](#) on page 72

5. *Cabling Node B* on page 81
6. *Assigning disk pools (if you have software-based disk ownership)* on page 91
7. *Verifying disk paths* on page 92

### Related concepts

- Setup requirements and restrictions for fabric-attached MetroClusters* on page 36
- Configuring an active/active configuration* on page 103
- Disaster recovery using MetroCluster* on page 173

### Related tasks

- Disabling the change\_fsid option in MetroCluster configurations* on page 108

## Planning the fabric-attached MetroCluster installation

You must fill out the fabric-attached MetroCluster worksheet 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.

### Step

1. Fill in the following tables.

Each site has two Brocade Fibre Channel switches. Use the following table to record the configured names, IP addresses, and domain IDs of these switches.

Switch number...	At site...	Is named...	IP address...	Domain ID...
1	A			
2	A			
3	B			
4	B			

In addition to on-board ports, each site has a FC-VI adapter and two Fibre Channel HBAs that connect the node to the switches. Use the following table to record which switch port these adapters are connected to.

This adapter...	At site...	Port 1 of this adapter is...		Port 2 of this adapter is...	
		Cabled to switch...	Switch port...	Cabled to switch...	Switch port...
FC-VI adapter	A	1		2	
	B	3		4	
FC HBA 1	A	1		2	
	B	3		4	

This adapter...	At site...	Port 1 of this adapter is...		Port 2 of this adapter is...	
		Cabled to switch...	Switch port...	Cabled to switch...	Switch port...
FC HBA 2	A	1		2	
	B	3		4	

Disk shelves at each site connect to the Fibre Channel switches. Use the following table to record which switch port the disk shelves are connected to.

Disk shelf...	At site...	Belonging to...	Connects to switches...	On switch port...
1	A	Node A Pool 0	1 and 2	
2				
3		Node B Pool 1		
4				
5	B	Node B Pool 0	3 and 4	
6				
7		Node A Pool 1		
8				

## Configuration differences for fabric-attached MetroClusters on 31xx systems

When configuring a fabric-attached MetroCluster between 31xx series systems, each system must only have one controller in the chassis.

### Related concepts

[Fabric-attached MetroCluster configuration on 31xx systems](#) on page 35

## Configuring the switches

To configure the switches, you refer to the *Brocade Switch Configuration Guide for Fabric-attached MetroClusters* for your Brocade switch model. The Brocade switch configuration for a MetroCluster is different than the one used for a SAN configuration.

### Step

1. To configure your Brocade switches, see the *Brocade Switch Configuration Guide for Fabric-attached MetroClusters* for your switch model. You can find this document on the MetroCluster Switch Description Page at

[http://now.netapp.com/NOW/download/software/sanswitch/fcp/Brocade/mc\\_ontap641\\_fabric\\_index.shtml](http://now.netapp.com/NOW/download/software/sanswitch/fcp/Brocade/mc_ontap641_fabric_index.shtml).  
 Scroll down the page to **Product Documentation ► Switch Configuration ► Documentation**.

**Note:** The configuration and firmware requirements for Brocade switches in a MetroCluster environment are different from the requirements for switches used in SAN environments. Always refer to MetroCluster documentation, such as the MetroCluster Compatibility Matrix or the MetroCluster Switch Description Page, when installing and configuring your MetroCluster switches.

**After you finish**

Proceed to configure Node A.


**Related concepts**

*Switch bank rules and virtual channel rules* on page 71

**Switch bank rules and virtual channel rules**

You must follow the correct switch bank rules or virtual channel rules on the Brocade switches.


If your system uses hardware-based disk ownership, you must use the switch bank rules when cabling the Brocade switch. This ensures that switch traffic is distributed across the switch quadrants to reduce potential bottlenecks.



**Figure 11: Brocade switch showing which ports belong to which switch banks and pools**

If your system does not use hardware-based disk ownership, use the switch virtual channel (VC) rules when cabling the switch. In this case, switch traffic is distributed across VCs to avoid bottlenecks. The FC-VI and inter-switch links are cabled to ports in one VC, and the disk shelf and controller connections are cabled to ports in another VC.

Virtual channel	Ports
2	0, 4, 8, 12
3	1, 5, 9, 13
4	2, 6, 10, 14
5	3, 7, 11, 15



**Figure 12: Brocade switch showing which ports belong to which virtual channels**

### Related information

*Brocade Switch Configuration Guide for Fabric MetroCluster -*

[http://now.netapp.com/NOW/download/software/sanswitch/fcp/Brocade/mc\\_ontap641\\_fabric\\_index.shtml](http://now.netapp.com/NOW/download/software/sanswitch/fcp/Brocade/mc_ontap641_fabric_index.shtml)

## Cabling Node A

To cable the local node (Node A), you need to attach the controller and the disk shelves to the switches, connect the cluster interconnect to the switches, and ensure that the disk shelves in the configuration belong to the correct pools.

### About this task

Complete the following tasks in the order shown:

1. *Cabling the controller to the switches* on page 72
2. *Cabling the disk shelves to the switches* on page 75
3. *Cabling the FC-VI adapter and inter-switch link* on page 79

## Cabling the controller to the switches

You must cable the controller to the Brocade switches.

Proceed to the appropriate section.

### Next topics

*Cabling the controller when you have hardware-based disk ownership* on page 72

*Cabling the controller when you have software-based disk ownership* on page 73


*Tips for controller-to-switch connections* on page 75

## Cabling the controller when you have hardware-based disk ownership

If you are using hardware-based disk ownership, you must cable the controller to the appropriate port bank on the switch.

Make sure you are following the switch port rules. See the *Brocade Switch Configuration Guide* for your switch. You can find this document on the MetroCluster Switch Description Page at [http://now.netapp.com/NOW/download/software/sanswitch/fcp/Brocade/mc\\_ontap641\\_fabric\\_index.shtml](http://now.netapp.com/NOW/download/software/sanswitch/fcp/Brocade/mc_ontap641_fabric_index.shtml).





**Figure 13: Controller cabling on Node A with hardware-based disk ownership**

### Steps

1. Determine which Fibre Channel ports on your system that you want to use and create a list showing the preferred order in which you want to use them.

**Attention:** The numbers in the example refer to the preferred order of usage, not the port ID. For example, Fibre Channel port 1 might be port e0a on the controller.

2. Cable the first two Fibre Channel ports of Node A to port 0 of Switch 1 and Switch 2.
3. Cable the second two Fibre Channel ports of Node A to port 1 of Switch 1 and Switch 2.

Proceed to cable disk shelves to the switches.

### Related concepts

[Switch bank rules and virtual channel rules](#) on page 71

[Which Fibre Channel ports to use for DS14mk2 or DS14mk4 FC disk shelf connections](#) on page 51

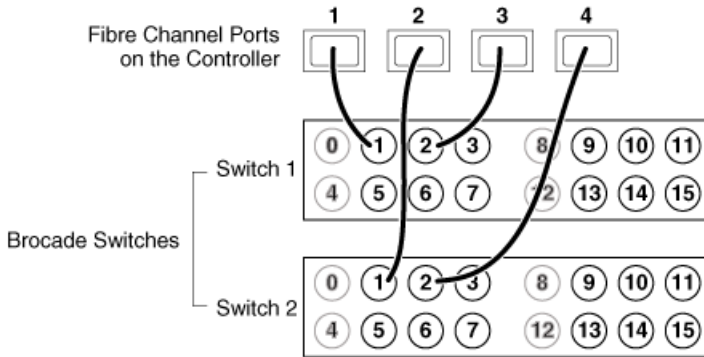
### Cabling the controller when you have software-based disk ownership

You can use this procedure to cable the Fibre Channel ports on the controller to the Brocade switches when your system uses software-based disk ownership.

- Select one virtual channel on the switch for the cluster interconnect connections. The following examples use virtual channel 2, which includes ports 0, 4, 8, and 12.
- If you are using a dual-port HBA, connecting both ports of the HBA to the same switch port number can make it easier to cable and administer your MetroCluster. (However, this is not required.)  
For example, if port 1 of the HBA is connected to port 1 of Switch 1, you should connect port 2 of that HBA to port 1 of Switch 2.

- Both Fibre Channel ports on the same dual-port HBA (or adjacent pairs of onboard ports) should never be connected to the same switch. You must connect one port to one switch and the other port to the other switch.

For example, if onboard port 0a is connected to Switch 3, you should not connect onboard port 0b to Switch 3; you must connect port 0b to Switch 4.



**Figure 14: Controller cabling on Node A with software-based disk ownership**

### Steps

- Determine which Fibre Channel ports on your system that you want to use and create a list showing the order you want to use them.

**Note:** The numbers in the example refer to the preferred order of usage, not the port ID. For example, Fibre Channel port 1 might be port e0a on the controller.

- Cable the first two Fibre Channel ports of Node A to the same numbered ports on Switch 1 and Switch 2. For example, port 1.

They must not go to ports in the virtual channel that you have reserved for the FC-VI and inter-switch link connections. In the example, we are using virtual channel 2 for the FC-VI and inter-switch link. Virtual channel 2 includes ports 0, 4, 8, and 12.

- Cable the second two Fibre Channel ports of Node A to the same numbered ports on Switch 1 and Switch 2. For example, port 2.

Again, they must not go to ports in the virtual channel that you have reserved for the FC-VI and inter-switch link connections. In the example, ports 0, 4, 8, and 12 are excluded.

**Note:** The switches in the example are 16-port switches.

Proceed to cable disk shelves to the switches.

### Related concepts

*Which Fibre Channel ports to use for DS14mk2 or DS14mk4 FC disk shelf connections* on page 51

## Tips for controller-to-switch connections

You should be aware of these important tips when you are cabling the controller.

- If you are using hardware-based disk ownership, the Fibre Channel connections from the controller to the switch must be connected to the correct switch bank:
  - For the local node (Node A), you should use switch bank 1.
  - For the remote node (Node B), you should use switch bank 2.
- If you are using software-based disk ownership, select one virtual channel on the switch for the cluster interconnect connections. The following examples use virtual channel 2, which includes ports 0, 4, 8, and 12.
- If you are using a dual-port HBA, connecting both ports of the HBA to the same switch port number can make it easier to cable and administer your MetroCluster. (However, this is not required.) For example, if port 1 of the HBA is connected to port 1 of Switch 1, you should connect port 2 of that HBA to port 1 of Switch 2.
- Both Fibre Channel ports on the same dual-port HBA (or adjacent pairs of onboard ports) should never be connected to the same switch. You must connect one port to one switch and the other port to the other switch. For example, if onboard port 0a is connected to Switch 3, you should not connect onboard port 0b to Switch 3; you must connect port 0b to Switch 4.

### Related concepts

[Switch bank rules and virtual channel rules](#) on page 71

## Cabling the disk shelves to the switches

Describes how to cable the disk shelves of Node A to the Brocade switches.

Make sure you meet disk shelf requirements and ensure that all disk shelves are operating at the same speed, either 2 Gbps or 4 Gbps.

Proceed to the appropriate section.

### Next topics

[Cabling the disk shelves when you have hardware-based disk ownership](#) on page 75


[Cabling the shelves when you have software-based disk ownership](#) on page 77

[Tips for disk shelf-to-switch connections](#) on page 78

## Cabling the disk shelves when you have hardware-based disk ownership

Describes how to cable the disk shelves on Node A to the Brocade switches.

**Note:** You can cable a maximum of two disk shelves on each loop.



**Figure 15: Cabling disk shelves on Node A with hardware-based disk ownership**

### Steps

1. Connect Node A pool 0 disks to the switches by completing the following substeps:
  - a. Connect the Input port of the A module on disk shelf 1 to port 9 on Switch 2.
  - b. Connect the Input port of the B module 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 module of disk shelf 1 to the Input ports of the corresponding module of the next disk shelf.
  - d. If your disk shelf modules have terminate switches, set them to Off on all but the last disk shelf in the disk pool, then set the terminate switches on the last disk shelf to On.

**Note:** ESH2 and ESH4 modules are self-terminating and therefore do not have a terminate switch.

2. Connect Node B pool 1 disks to the switches by completing the following substeps:

- a. Connect the Input port of the module Channel A on disk shelf 3 to port 5 on Switch 2.
  - b. Connect the Input port of the module Channel 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 module of disk shelf 3 to the Input ports of the corresponding module of the next disk shelf.
  - d. If your disk shelf modules have terminate switches, set them to Off on all but the last disk shelf in the disk pool, then set the terminate switches on the last disk shelf to On.
3. If you have more than one loop, connect the other loops in the same manner, making sure that you use the correct switch quadrant (ports 4-7 for 16-port switches, ports 2-3 for 8-port switches).

Proceed to cable the FC-VI adapter and inter-switch connections.

### Cabling the shelves when you have software-based disk ownership

You must cable the disk shelf loops on Node A directly to the Brocade switches.

- You can connect the disk shelves to any ports on that are not on the virtual channel reserved for the FC-VI adapter and the inter-switch link.
- Both disk shelf modules on the same loop must be connected to the same switch port number. For example, if the A Channel of the first loop for the local node’s disks is connected to Switch 1, port 8, then the B Channel for that loop must be connected to Switch 2, port 8.
- Both switches at a site must be the same model and have the same number of licensed ports.

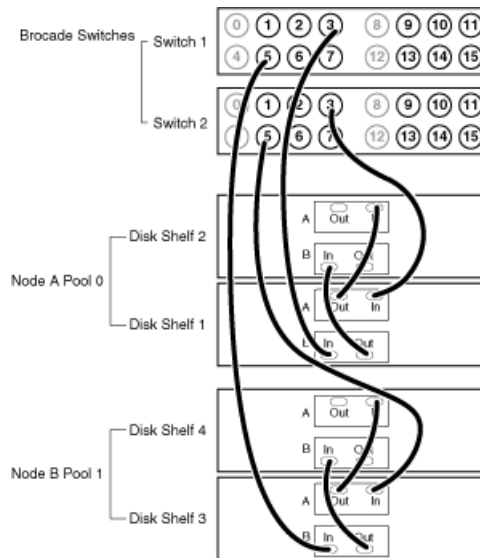


Figure 16: Cabling disk shelves on Node A

**Note:** You can cable a maximum of two disk shelves on each loop.

## Steps

1. Connect the Node A pool 0 disk shelves to the switches by completing the following substeps:
  - a. Connect the Input port of the A module on disk shelf 1 to any available port on Switch 2 other than ports 0, 4, 8, and 12. In the example, switch port 3 is used.
  - b. Connect the Input port of the B module on disk shelf 1 to the same port on Switch 1. The example uses switch port 3.
  - c. Connect disk shelf 1 to disk shelf 2 by connecting the Output ports of the module of disk shelf 1 to the Input ports of the corresponding module of the next disk shelf.
  - d. If your disk shelf modules have terminate switches, set them to Off on all but the last disk shelf in the disk pool, then set the terminate switches on the last disk shelf to On.

**Note:** ESH2 and ESH4 modules are self-terminating and therefore do not have a terminate switch.
2. Connect the Node B pool 1 disk shelves to the switches by completing the following substeps:
  - a. Connect the Input port of the module Channel A on disk shelf 3 to any available port on Switch 2 other than ports 0, 4, 8, and 12. The example uses switch port 5.
  - b. Connect the Input port of the module Channel B on disk shelf 3 to the same port on Switch 1. The example uses switch port 5.
  - c. Connect disk shelf 3 to disk shelf 4 by connecting the Output ports of the module of disk shelf 3 to the Input ports of the corresponding module of the next disk shelf.
  - d. If your disk shelf modules have terminate switches, set them to Off on all but the last disk shelf in the disk pool, then set the terminate switches on the last disk shelf to On.
3. If you have more than one loop, connect the other loops in the same manner.

Proceed to cable the FC-VI adapter and inter-switch connections.

## Tips for disk shelf-to-switch connections

You should be aware of these important tips when you are cabling the disk shelves.

- If you are using software-based disk ownership, you can connect the disk shelves to any ports on that are not on the virtual channel reserved for the FC-VI adapter and the inter-switch link.
- If you are using hardware-based disk ownership, you should connect the disk shelves to the opposite switch bank from the bank used for the controller-switch Fibre Channel connections, as follows: :
  - For the local node (Node B), you should use switch bank 2.
  - For the remote node (Node B), you should use switch bank 1.

In addition to using the correct switch bank, you should use the correct switch quadrant for the pool, as follows:

- For the local node (Node B), you should use quadrant 3 for pool 0 and quadrant 4 for pool 1.
- For the remote node (Node B), you should use quadrant 1 for pool 0 and quadrant 2 for pool 1.

**Note:** For information about how Brocade switches are divided into banks and quadrants, see the *Brocade Switch Configuration Guide* for your switch.

- Both disk shelf modules on the same loop must be connected to the same switch port number. For example, if the A Channel of the first loop for the local node's disks is connected to Switch 1, port 8, then the B Channel for that loop must be connected to Switch 2, port 8.
- Both switches at a site must be the same model and have the same number of licensed ports.

### Related concepts

[Switch bank rules and virtual channel rules](#) on page 71

### Related information

[Brocade 200E and Brocade 5000 Switch Configuration Guide -](#)

[http://now.netapp.com/NOW/download/software/sanswitch/fcp/Brocade/mc\\_ontap641\\_fabric\\_index.shtml](http://now.netapp.com/NOW/download/software/sanswitch/fcp/Brocade/mc_ontap641_fabric_index.shtml)

## Cabling the FC-VI adapter and inter-switch link

Describes how to cable the FC-VI adapter and inter-switch link on Node A.

Proceed to the appropriate section.

### Next topics

[Cabling the FC-VI adapter and inter-switch link when you have hardware-based disk ownership](#) on page 79

[Cabling the FC-VI adapter and inter-switch link when you have software-based disk ownership](#) on page 80

[Tips for cluster interconnect connections](#) on page 81

## Cabling the FC-VI adapter and inter-switch link when you have hardware-based disk ownership

You must perform the following steps in order to cable the cluster interconnect and inter-switch link on Node A.




Figure 17: Cabling the interconnects on Node A with hardware-based disk ownership

## Steps

1. Connect the FC-VI adapter to the switches, with one port connecting to Switch 1 and the second port to Switch 2.

In the example, Port 12 is used on both switches. It is in bank 2, pool 1.

**Note:** There should be one FC-VI adapter connection for each switch. The switch port connections do not have to be identical; therefore, use any available switch ports. Make sure that you have the FC-VI cluster adapter in the correct slot for your system, as shown in the *System Configuration Guide* at [http://now.netapp.com/NOW/knowledge/docs/hardware/hardware\\_index.shtml](http://now.netapp.com/NOW/knowledge/docs/hardware/hardware_index.shtml).

2. Connect an inter-switch link cable to a port on each switch, or, if using a dual inter-switch link, connect two cables in the selected quadrant.

Use the same ports on each switch. In the example, ports 13 and 14 are used on both switches.

**Note:** If using dual inter-switch links, traffic isolation must be configured on the switches.

Proceed to cable Node B.

## Related concepts


[Switch bank rules and virtual channel rules](#) on page 71

## Cabling the FC-VI adapter and inter-switch link when you have software-based disk ownership

Describes how to cable the cluster interconnect and inter-switch link on Node A.

Each port on the Interconnect (IC) cards must be connected to the same fabric.

For example, if Port A of the IC card on the local node is connected to Switch 1, and Port A of the IC card on the remote node is connected to Switch 3, then Switch 1 and Switch 3 must be connected by the inter-switch link, thereby connecting them to the same fabric.



**Figure 18: Cabling the interconnects on Node A with software-based disk ownership**



## Steps

1. Using the ports in the virtual channel you have selected for the FC-VI and inter-switch link connections, connect one port of the FC-VI adapter on switch 1 and the second port to the same port on switch 2.

In the example we are using virtual channel 2, including ports 0, 4, 8, and 12, for the FC-VI and inter-switch link connections.

**Note:** There should be one FC-VI adapter connection for each switch. Make sure that you have the FC-VI adapter in the correct slot for your system, as shown in the *System Configuration Guide* at [http://now.netapp.com/NOW/knowledge/docs/hardware/hardware\\_index.shtml](http://now.netapp.com/NOW/knowledge/docs/hardware/hardware_index.shtml).

2. Connect an inter-switch link cable to a port in the selected virtual channel on each switch, or, if using a dual inter-switch link, connect two cables in the selected virtual channel.

In the example we are using virtual channel 2, which includes ports 0, 4, 8, and 12, and are using ports 8 and 12 on switch 1 and switch 2 for the inter-switch links.

**Note:** If using dual inter-switch links, traffic isolation must be configured on the switches.

Proceed to cable Node B.

## Tips for cluster interconnect connections

Describes important tips to be aware of when you are cabling the cluster interconnect.

Each port on the Interconnect (IC) cards must be connected to the same fabric.

For example, if Port A of the IC card on the local node is connected to Switch 1, and Port A of the IC card on the remote node is connected to Switch 3, then Switch 1 and Switch 3 must be connected by the inter-switch link, thereby connecting them to the same fabric.

## Cabling Node B

To cable the remote node (Node B), you need to attach the controller and the disk shelves to the switches, connect the cluster interconnect to the switches, and ensure that the disk shelves in the configuration belong to the correct pools.

### About this task

Complete the following tasks in the order shown:

1. [Cabling the Node B controller to the switches](#) on page 82
2. [Cabling the disk shelves to the switches](#) on page 84
3. [Cabling the FC-VI adapter and inter-switch link](#) on page 88

## Cabling the Node B controller to the switches

Describes how to cable the controller to the Brocade switches.

Proceed to the appropriate section.

### Next topics


[Cabling the controller when you have hardware-based disk ownership](#) on page 82

[Cabling the controller when you have software-based disk ownership](#) on page 83

[Tips for controller-to-switch connections](#) on page 84

## Cabling the controller when you have hardware-based disk ownership

If you are using hardware-based disk ownership, you must cable the controller to the appropriate port bank on the switch.



**Figure 19: Controller cabling on Node B with hardware-based disk ownership**

### Steps

1. Determine which Fibre Channel ports on your system that you want to use and create a list showing the preferred order in which you want to use them.

**Attention:** The numbers in the example refer to the preferred order of usage, not the port ID. For example, Fibre Channel port 1 might be port e0a on the controller.

2. Cable the first two Fibre Channel ports of Node B to Switch 3 and Switch 4. The example uses port 8.
3. Cable the second two Fibre Channel ports of Node B to Switch 3 and Switch 4. The example uses port 9.

Proceed to cable disk shelves to the switches.

## Related concepts

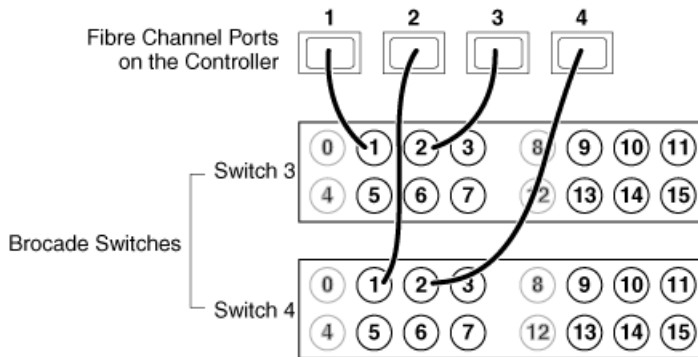
*Which Fibre Channel ports to use for DS14mk2 or DS14mk4 FC disk shelf connections* on page 51

## Cabling the controller when you have software-based disk ownership

You can use this procedure to cable the Fibre Channel ports on the controller to the Brocade switches when your system uses software-based disk ownership.

- Select one virtual channel on the switch for the cluster interconnect connections. The following examples use virtual channel 2, which includes ports 0, 4, 8, and 12.
- If you are using a dual-port HBA, connecting both ports of the HBA to the same switch port number can make it easier to cable and administer your MetroCluster. (However, this is not required.)  
For example, if port 1 of the HBA is connected to port 1 of Switch 1, you should connect port 2 of that HBA to port 1 of Switch 2.
- Both Fibre Channel ports on the same dual-port HBA (or adjacent pairs of onboard ports) should never be connected to the same switch. You must connect one port to one switch and the other port to the other switch.

For example, if onboard port 0a is connected to Switch 3, you should not connect onboard port 0b to Switch 3; you must connect port 0b to Switch 4.



**Figure 20: Controller cabling on Node B with software-based disk ownership**

## Steps

1. Determine which Fibre Channel ports on your system that you want to use and create a list showing the order you want to use them.

**Note:** The numbers in the example refer to the preferred order of usage, not the port ID. For example, Fibre Channel port 1 might be port e0a on the controller.

2. Cable the first two Fibre Channel ports of Node B to the same numbered ports Switch 3 and Switch 4. For example, port 1.

They must go to ports in the virtual channel that you have reserved for the FC-VI and inter-switch link connections. In the example, we are using virtual channel 2 for the FC-VI and inter-switch link. Virtual channel 2 includes ports 0, 4, 8, and 12.

3. Cable the second two Fibre Channel ports of Node B to the same numbered ports Switch 3 and Switch 4. For example, port 2.

Again, they must not go to ports in the virtual channel that you have reserved for the FC-VI and inter-switch link connections. In the example, ports 0, 4, 8, and 12 are excluded.

Proceed to cable disk shelves to the switches.

### Related concepts

*Which Fibre Channel ports to use for DS14mk2 or DS14mk4 FC disk shelf connections* on page 51

## Tips for controller-to-switch connections

You should be aware of these important tips when you are cabling the controller.

- If you are using hardware-based disk ownership, the Fibre Channel connections from the controller to the switch must be connected to the correct switch bank:
  - For the local node (Node B), you should use switch bank 1.
  - For the remote node (Node B), you should use switch bank 2.
- If you are using software-based disk ownership, select one virtual channel on the switch for the FC-VI connections. This virtual channel must also be used for the inter-switch link connections. The following examples use virtual channel 2, which includes ports 0, 4, 8, and 12.
- If you are using a dual-port HBA, connecting both ports of the HBA to the same switch port number can make it easier to cable and administer your MetroCluster. (However, this is not required.) For example, if port 1 of the HBA is connected to port 1 of Switch 1, you should connect port 2 of that HBA to port 1 of Switch 2.
- Both Fibre Channel ports on the same dual-port HBA (or adjacent pairs of onboard ports) should never be connected to the same switch. You must connect one port to one switch and the other port to the other switch.
 

For example, if onboard port 0a is connected to Switch 3, you should not connect onboard port 0b to Switch 3; you must connect port 0b to Switch 4.

### Related concepts

*Switch bank rules and virtual channel rules* on page 71

## Cabling the disk shelves to the switches

Describes how to cable the disk shelves of Node B to the Brocade switches.

Make sure you meet disk shelf requirements and ensure that all disk shelves are operating at the same speed, either 2 Gbps or 4 Gbps.

Proceed to the appropriate section.


**Next topics**

- Cabling the disk shelves when you have hardware-based disk ownership* on page 85
- Cabling the shelves when you have software-based disk ownership* on page 86
- Tips for disk shelf-to-switch connections* on page 88

**Cabling the disk shelves when you have hardware-based disk ownership**

Describes how to cable the disk shelves on Node B to the Brocade switches.

**Note:** You can cable a maximum of two disk shelves on each loop.



**Figure 21: Cabling disk shelves on Node B with hardware-based disk ownership**

## Steps


1. Connect Node B pool 0 disks to the switches by completing the following substeps:
  - a. Connect the Input port of the A module of disk shelf 3, to port 0 on Switch 4.
  - b. Connect the Input port of the B module 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 module of disk shelf 3 to the Input ports of the corresponding module of disk shelf 4.
  - d. If your disk shelf modules have terminate switches, set the terminate switches to Off on all but the last disk shelf in the disk pool, then set the terminate switches on the last disk shelf to On.  
**Note:** ESH2 and ESH4 modules are self-terminating and therefore do not have a terminate switch.
2. Connect Node B pool 1 disks to the switches by completing the following substeps:
  - a. Connect the Input port of the A module on disk shelf 5 to port 14 on Switch 4.
  - b. Connect the Input port of the B module on disk shelf 5 to port 14 on Switch 3.
  - c. Connect disk shelf 1 to disk shelf 2 by connecting the Output ports of the module of disk shelf 1 to the Input ports of the corresponding module of disk shelf 2.
  - d. If your disk shelf modules have terminate switches, set the terminate switches to Off on all but the last disk shelf in the disk pool, then set the terminate switches on the last disk shelf to On.
3. If you have more than one loop, connect the other loops in the same manner, making sure that you use the correct switch quadrant (ports 12 through 15 for 16-port switches or ports 6 through 7 for 8-port switches).

Proceed to cable the FC-VI adapter and inter-switch connections.

## Cabling the shelves when you have software-based disk ownership

You must cable the disk shelf loops on Node B directly to the Brocade switches.

- You can connect the disk shelves to any ports on that are not on the virtual channel reserved for the FC-VI adapter and the inter-switch link.
- Both disk shelf modules on the same loop must be connected to the same switch port number. For example, if the A Channel of the first loop for the local node's disks is connected to Switch 1, port 8, then the B Channel for that loop must be connected to Switch 2, port 8.
- Both switches at a site must be the same model and have the same number of licensed ports.



**Figure 22: Cabling disk shelves on Node B with software-based disk ownership**

**Note:** You can cable a maximum of two disk shelves on each loop.

### Steps

1. Connect the Node B pool 0 disk shelves to the switches by completing the following substeps:
  - a. Connect the Input port of the A module on disk shelf 5 to any available port on Switch 4 that is not in the virtual channel reserved for the FC-VI and inter-switch link connections. The example uses switch port 3.
  - b. Connect the Input port of the B module on disk shelf 5 to the same port on Switch 3. The example uses switch port 3.
  - c. Connect disk shelf 5 to disk shelf 6 by connecting the Output ports of the module of disk shelf 5 to the Input ports of the corresponding module of the next disk shelf.
  - d. If your disk shelf modules have terminate switches, set them to Off on all but the last disk shelf in the disk pool, then set the terminate switches on the last disk shelf to On.

**Note:** ESH2 and ESH4 modules are self-terminating and therefore do not have a terminate switch.

2. Connect the Node B pool 1 disk shelves to the switches by completing the following substeps:
  - a. Connect the Input port of the module Channel A on disk shelf 7 to any available port on Switch 4 that is not in the virtual channel reserved for the FC-VI and inter-switch link connections. The example uses switch port 5.
  - b. Connect the Input port of the module Channel B on disk shelf 7 to the same port on Switch 3. The example uses switch port 5.
  - c. Connect disk shelf 7 to disk shelf 8 by connecting the Output ports of the module of disk shelf 7 to the Input ports of the corresponding module of the next disk shelf.

- d. If your disk shelf modules have terminate switches, set them to Off on all but the last disk shelf in the disk pool, then set the terminate switches on the last disk shelf to On.
3. If you have more than one loop, connect the other loops in the same manner.

Proceed to cable the FC-VI adapter and inter-switch connections.

### Tips for disk shelf-to-switch connections

You should be aware of these important tips when you are cabling the disk shelves.

- If you are using software-based disk ownership, you can connect the disk shelves to any ports on that are not on the virtual channel reserved for the FC-VI adapter and the inter-switch link.
- If you are using hardware-based disk ownership, you should connect the disk shelves to the opposite switch bank from the bank used for the controller-switch Fibre Channel connections, as follows: :
  - For the local node (Node B), you should use switch bank 2.
  - For the remote node (Node B), you should use switch bank 1.

In addition to using the correct switch bank, you should use the correct switch quadrant for the pool, as follows:

- For the local node (Node B), you should use quadrant 3 for pool 0 and quadrant 4 for pool 1.
- For the remote node (Node B), you should use quadrant 1 for pool 0 and quadrant 2 for pool 1.

**Note:** For information about how Brocade switches are divided into banks and quadrants, see the *Brocade Switch Configuration Guide* for your switch. You can find this document on the MetroCluster Switch Description Page at:

[http://now.netapp.com/NOW/download/software/sanswitch/fcp/Brocade/mc\\_ontap641\\_fabric\\_index.shtml](http://now.netapp.com/NOW/download/software/sanswitch/fcp/Brocade/mc_ontap641_fabric_index.shtml).

- Both disk shelf modules on the same loop must be connected to the same switch port number. For example, if the A Channel of the first loop for the local node's disks is connected to Switch 1, port 8, then the B Channel for that loop must be connected to Switch 2, port 8.
- Both switches at a site must be the same model and have the same number of licensed ports.

### Related concepts

*Switch bank rules and virtual channel rules* on page 71

### Cabling the FC-VI adapter and inter-switch link

Describes how to cable the cluster interconnect and inter-switch connection on Node B.

Proceed to the appropriate section.

### Next topics

*Cabling the FC-VI adapter and inter-switch link when you have hardware-based disk ownership* on page 89




*Cabling the FC-VI adapter and inter-switch link when you have software-based disk ownership* on page 90

*Tips for cluster interconnect connections* on page 90

## Cabling the FC-VI adapter and inter-switch link when you have hardware-based disk ownership

To cable the cluster interconnect and inter-switch link on Node B you must perform the following steps in order.



**Figure 23: Cabling the interconnects on Node B with hardware-based disk ownership**

### Steps

1. Connect the FC-VI adapter to the switches, with one port connecting to Switch 3 and the second port to Switch 4.

In the example, port 12 is used on both switches. It is in bank 2, pool 1.

**Note:** There should be one FC-VI adapter connection for each switch. The switch port connections do not have to be identical; therefore, use any available switch ports. Make sure that you have the FC-VI cluster adapter in the correct slot for your system, as shown in the *System Configuration Guide* at [http://now.netapp.com/NOW/knowledge/docs/hardware/hardware\\_index.shtml](http://now.netapp.com/NOW/knowledge/docs/hardware/hardware_index.shtml).

2. Connect an inter-switch link cable to a port on each switch, or, if using a dual inter-switch link, connect two cables in the selected quadrant.

Use the same port number on each switch. In the example, ports 13 and 14 are used on both switches. It is in bank 2, pool 1.

**Note:** If using dual inter-switch links, traffic isolation must be configured on the switches.

Proceed to verify the disk paths on the system.

### Related concepts

*Switch bank rules and virtual channel rules* on page 71

### Related tasks


*Verifying disk paths* on page 92

## Cabling the FC-VI adapter and inter-switch link when you have software-based disk ownership

You must cable the cluster interconnect and inter-switch link on Node B.

Each port on the Interconnect (IC) cards must be connected to the same fabric.

For example, if port A of the IC card on the local node is connected to switch 1, and port A of the IC card on the remote node is connected to switch 3, then switch 1 and switch 3 must be connected by the inter-switch link, thereby connecting them to the same fabric.



**Figure 24: Cabling the interconnects on Node B with software-based disk ownership**

### Steps

1. Connect one port of the FC-VI adapter to a port in the virtual channel that you have reserved for the FC-VI and inter-switch link connections.

In the example, port 0 on switch 1 and port 0 on switch 2 is used.

**Note:** There should be one FC-VI adapter connection for each switch. Make sure that you have the FC-VI adapter in the correct slot for your system, as shown in the *System Configuration Guide* at [http://now.netapp.com/NOW/knowledge/docs/hardware/hardware\\_index.shtml](http://now.netapp.com/NOW/knowledge/docs/hardware/hardware_index.shtml).

2. Connect an inter-switch link cable to a port in the selected virtual channel on each switch, or if using dual inter-switch links, connect two cables in the selected virtual channel.

In the example we are using virtual channel 2, which includes ports 0, 4, 8, and 12, and are using port 8 and port 12 on switch 1 and switch 2 for the inter-switch links.

**Note:** If using dual inter-switch links, traffic isolation must be configured on the switches.

Proceed to assign disks to disk pools.

### Related tasks

*Assigning disk pools (if you have software-based disk ownership)* on page 91

## Tips for cluster interconnect connections

You should be aware of these important tips when you are cabling the cluster interconnect.

Each port on the Interconnect (IC) cards must be connected to the same fabric.

For example, if Port A of the IC card on the local node is connected to Switch 1, and Port A of the IC card on the remote node is connected to Switch 3, then Switch 1 and Switch 3 must be connected by the inter-switch link, thereby connecting them to the same fabric.

## Assigning disk pools (if you have software-based disk ownership)

If your system uses software-based disk ownership, you must assign the attached disk shelves to the appropriate pools.

### About this task

You can explicitly assign disks on the attached disk shelves to the appropriate pool with the `disk assign` command. Using wildcards in the command enables you to assign all the disks on a disk shelf with one command.

The following table shows the pool assignments for the disk shelves in the example used in this section.

Disk shelf...	At site...	Belongs to...	And is assigned to that node's	
Disk shelf 1	Site A	Node A	Pool 0	
Disk shelf 2				
Disk shelf 3		Node B		Pool 1
Disk shelf 4				
Disk shelf 5	Site B	Node B	Pool 0	
Disk shelf 6				
Disk shelf 7		Node A		Pool 1
Disk shelf 8				

**Note:** Pool 0 always contains the disks that are local to (at the same site as) the storage system that owns them.

Pool 1 always contains the disks that are remote to the storage system that owns them.

### Steps

1. Boot Node A into Maintenance mode, if you haven't already.
2. Assign the local disks to Node A pool 0 by entering the following command at the console:

```
disk assign switch2:port3.* -p0
```

This indicates that the disks attached to port 3 of switch 2 are assigned to pool 0. The asterisk (\*) indicates that all disks attached to the port are assigned.

3. Assign the remote disks to Node A pool 1 by entering the following command at the console:

```
disk assign switch4:port5.* -p1
```

This indicates that the disks attached to port 5 of switch 4 are assigned to pool 1. The asterisk (\*) indicates that all disks attached to the port are assigned.

4. Boot Node B into Maintenance mode, if you haven't already.
5. Assign the local disks to Node B pool 0 by entering the following command at the console:

```
disk assign switch4:port12.* -p0
```

This indicates that the disks attached to port 3 of switch 4 are assigned to pool 0. The asterisk (\*) indicates that all disks attached to the port are assigned.

6. Assign the remote disks to Node B pool 1 by entering the following command at the console:

```
disk assign switch2:port0.* -p1
```

This indicates that the disks attached to port 5 of switch 2 are assigned to pool 1. The asterisk (\*) indicates that all disks attached to the port are assigned.

### After you finish

Proceed to verify the disk paths on the system.

## Verifying disk paths

You should use the steps in this procedure to verify that you have correctly cabled your system.

### About this task

Proceed to the appropriate section.

### Next topics

[Verifying disk paths if you have hardware-based disk ownership](#) on page 92

[Verifying disk paths if you have software-based disk ownership](#) on page 93

## Verifying disk paths if you have hardware-based disk ownership

Use this procedure to verify your disk paths if your system uses hardware-based disk ownership.

### Steps

1. Boot Node A, if necessary.
2. Confirm that the disks are visible and have dual paths by entering the following command on the console:

```
storage show disk -p
```

**Example**

The output shows the disks connected to the switches, to what port they are connected, and to what disk and to which disk shelf they belong, as shown in the following example:

```

ha15*> storage show disk -p
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

```

3. Repeat steps 1 and 2 on Node B.

**Verifying disk paths if you have software-based disk ownership**

Use this procedure to verify your disk paths if you have software-based disk ownership.

**Steps**

1. Boot Node A into normal mode, if necessary.
2. Enter the following command to confirm that your aggregates and volumes are operational and mirrored:

```
aggr status
```

See the *Data ONTAP Storage Management Guide* for information on the `aggr status` command.

3. Repeat steps 1 and 2 on Node B.

**Required connections for using uninterruptible power supplies with MetroCluster configurations**

You can use a UPS (Uninterruptible Power Supply) with your MetroCluster. The UPS enables the system to fail over gracefully if power fails for one of the nodes, or to shut down gracefully if power fails for both nodes. You must ensure that the correct equipment is connected to the UPS.

The equipment that you need to connect to the UPS depends on how widespread a power outage you want to protect against. Always connect both controllers, any Fibre Channel switches in use, and any inter-switch link infrastructure (for example, a Dense Wavelength Division Multiplexing, or DWDM) to the UPS.

You can leave the disks on the regular power supply. In this case, if power is interrupted to one site, the controller can access the other plex until it shuts down or power is restored. If, however, power is

interrupted to both sites at the same time and the disks are not connected to the UPS, the MetroCluster cannot shut down gracefully.

# Reconfiguring an active/active configuration into two stand-alone systems

---

To remove an active/active configuration so that the nodes become stand-alone systems without redundancy, you must disable the active/active software features and then remove the hardware connections.

## About this task

Complete the following tasks in the order shown:

1. *Ensure uniform disk ownership within disk shelves and loops in the system* on page 95
2. *Disabling the active/active software* on page 96
3. *Reconfiguring nodes using disk shelves for stand-alone operation* on page 97
4. *Requirements when changing an node using array LUNs to stand-alone* on page 99
5. *Reconfiguring nodes using array LUNs for stand-alone operation* on page 99

## Ensure uniform disk ownership within disk shelves and loops in the system

In systems using software-based disk ownership, if a disk shelf or loop contains a mix of disks owned by Node A *and* Node B, you must use this procedure to move the data and make disk ownership uniform within the disk shelf or loop.

### About this task

You must ensure the following:

- Disk ownership is uniform within all disk shelves and loops in the system
- All the disks within a disk shelf or loop belong to a single node and pool

**Note:** It is a best practice to always assign all disks on the same loop to the same node and pool.

### Steps

1. Use the following command to identify any disk shelves or loops that contain both disks belonging to Node A and disks belonging to Node B:

```
disk show -v
```

2. Determine which node the disk shelf or loop with mixed ownership will be attached to when the active/active feature is unconfigured and record this information.

For example, if the majority of the disks in the loop belong to Node A, you probably want the entire loop to belong to stand-alone Node A.

### After you finish

Proceed to disable the active/active software.

## Disabling the active/active software

You need to disable the active/active configuration in the software before reconfiguring the hardware to completely unconfigure the active/active feature.

### Before you begin

Before performing this procedure you must ensure that all loops and disk shelves in the system contain disks belonging to one or the other nodes. The disk shelves and loops can't contain a mix of disks belonging to Node A and Node B. In any disk shelves or loops containing such a mix of disks, you must move data.

### Steps

1. Enter the following command on either node console:

```
cf disable
```

2. Disable the cluster license by entering the following command:

```
license delete cluster
```

3. Open the `/etc/rc` file with a text editor and remove references to the partner node in the `ifconfig` entries, as shown in the following example:

#### Example

Original entry:

```
ifconfig e0 199.9.204.254 partner 199.9.204.255
```

Edited entry:

```
ifconfig e0 199.9.204.254
```

4. Repeat Step 1 through Step 3 on the partner node.

### After you finish

Proceed to reconfigure the hardware.



## Reconfiguring nodes using disk shelves for stand-alone operation

You can use this procedure to reconfigure the hardware if you want to return to a single-controller configuration.

### Before you begin

You must disable the active/active software.

### Steps

1. Halt both nodes by entering the following command on each console:

```
halt
```

2. Using the information you recorded earlier, in the disk shelves or loops with mixed storage, physically move the disks to a disk shelf in a loop belonging to the node that owns the disk. For example, if the disk is owned by Node B, move it to a disk shelf in a loop that is owned by Node B.

**Note:** Alternatively, you can move the data on the disks using a product such as Snapshot software, rather than physically moving the disk. See the *Data ONTAP Data Protection Online Backup and Recovery Guide*.

After moving the data from the disk you can zero the disk and use the `disk remove_ownership` command to erase the ownership information from the disk. See the *Data ONTAP Storage Management Guide*.

3. If you are completely removing one node, so that all the disk shelves will belong to a single stand-alone node, complete the following substeps:
  - a. Boot the node being removed into Maintenance mode, as described in the *Data ONTAP System Administration Guide*.
  - b. Use the `disk reassign` command and reassign all disk shelves so that they all belong to the node that remains.

The `disk reassign` command has the following syntax:

```
disk reassign [-o <old_name> | -s <old_sysid>] [-n <new_name>] -d <new_sysid>
```

- c. Halt the node by entering the following command:

```
halt
```

4. Turn off the power to each node, then turn off the power to the disk shelves and unplug them from the power source.
5. Ground yourself, then remove the cluster interconnect cables from both nodes. See the hardware documentation for your system for more details.

6. Move or remove the adapter used for the cluster interconnect:

If your system uses a...	Then...
cluster interconnect adapter or an FC-VI adapter	Remove the adapter from the system.
NVRAM5 or NVRAM6 adapter	You might need to change the slot position of the adapter. See the <i>System Configuration Guide</i> for details about expansion slot usage for the adapter

7. Recable the system, depending on the type of system:

If you are converting a...	Then...
System with nonmirrored disks	<ol style="list-style-type: none"> <li>a. Disconnect all cabling from the Channel B loop on the local node.</li> <li>b. Repeat for the partner node.</li> </ol>
System with mirrored disks or a redundant Channel B loop	<ol style="list-style-type: none"> <li>a. Connect the local node to the open Channel B loop in its local disk shelves, as described in the appropriate disk shelf guide.</li> <li>b. Repeat for the partner node.</li> </ol>

8. Power on the disk shelves, then the individual nodes, monitoring the system console for error messages during the boot process.
9. Run all system diagnostics at the boot prompt by entering the following command on the system console:

```
boot diags
```

10. Unset the partner system ID by entering the following command at the prompt:

```
unsetenv partner-sysid
```

11. Boot the node by entering the following command:

```
boot
```

12. Check active/active configuration status by entering the following command:

```
cf status
```

If the active/active configuration is disabled, you see the following output:

```
Failover monitor not initialized
```

13. Repeat Step 1 through Step 10 for the partner node.

## Requirements when changing an node using array LUNs to stand-alone

After uncoupling V-Series systems in an active/active configuration, you might need to perform additional reconfiguration related to Data ONTAP ownership of array LUNs.

The following table summarizes the requirements when uncoupling a storage system using array LUNs from an active/active configuration.

If you want to...	Requirements for uncoupling systems are...	Requirements for array LUN assignments to systems are...
Make both systems in the pair stand-alone systems	Remove the active/active configuration software and interconnect cabling	No Data ONTAP reconfiguration of array LUNs is necessary. Each system can continue to own the array LUNs assigned to it.
Remove one system in the pair from service	Remove the active/active configuration software and interconnect cabling	After uncoupling the pair, you must do one of the following: <ul style="list-style-type: none"> <li>• If you want to continue to use the array LUNs for Data ONTAP, reassign the array LUNs owned by the system you are removing to another storage system.</li> <li>• Prepare the array LUNs assigned to the system you are removing for use by systems that do not run Data ONTAP.</li> </ul>

## Reconfiguring nodes using array LUNs for stand-alone operation

After uncoupling the nodes in active/active configuration, each node can continue to own its assigned array LUNs, you can reassign its array LUNs to another V-Series system, or you can release the persistent reservations on the array LUNs so the LUNs can be used by a non Data ONTAP host.

### Before you begin

You must disable the active/active configuration software.

### About this task

If you want both systems in the active/active configuration to remain in service and operate as stand-alone systems, each system can continue to own the array LUNs that were assigned to it. Each system, as a stand-alone, will continue to see the array LUNs owned by the other system because both systems are

still part of the same V-Series neighborhood. However, only the system that is the owner of the array LUNs can read from or write to the array LUN, and the systems can no longer fail over to each other.

**Steps**

1. On each node, halt the node by entering the following command at the console:

```
halt
```

2. Turn off the power to each node.
3. Ground yourself, then remove the cluster interconnect cables from both nodes. See the hardware documentation for your system for more details.
4. Move or remove the adapter used for the cluster interconnect.

If your system uses a...	Then...
<b>cluster interconnect adapter or an FC-VI adapter</b>	Remove the adapter from the system.
<b>NVRAM5 or NVRAM6 adapter</b>	You might need to change the slot position of the adapter. See the <i>System Configuration Guide</i> for details about expansion slot usage for the adapter

5. On each node, perform the following steps:
  - a. Power on the node, monitoring the system console for error messages during the boot process.
  - b. Unset the partner system ID by entering the following command at the prompt:

```
unsetenv partner-sysid
```

6. Perform the appropriate step in the following table for what you intend to do with your system and its storage.

If you want to...	Then...
<b>Keep both systems in service as stand-alone systems and continue with both systems owning the array LUNs that were already assigned to them</b>	Boot both systems by entering the following command on each system: <b>boot</b>

If you want to...	Then...
<b>Remove one of the systems from service but still use the storage that was assigned to that system for Data ONTAP</b>	<p data-bbox="637 244 1241 296"><b>a.</b> Boot the node being removed into Maintenance mode, as described in the <i>Data ONTAP System Administration Guide</i>.</p> <p data-bbox="637 305 1241 392"><b>b.</b> Use the <code>disk_reassign</code> command to reassign all the array LUNs so that they all belong to the node that remains. The <code>disk_reassign</code> command has the following syntax:</p> <pre data-bbox="673 413 1150 493"> disk_reassign [-o &lt;old_name&gt;   -s &lt;old_sysid&gt;] [-n &lt;new_name&gt;] -d &lt;new_sysid&gt; </pre> <p data-bbox="637 510 973 532"><b>c.</b> Remove the node from service.</p> <p data-bbox="637 541 1217 593"><b>d.</b> Boot the node you are keeping in service by entering the following command:</p> <pre data-bbox="673 614 731 637">boot</pre>
<b>Remove one of the systems from service and use the array LUNs that are currently assigned to it for a host that does not run Data ONTAP</b>	<p data-bbox="637 682 1241 821">Release the persistent reservations that Data ONTAP placed on those array LUNs so that the storage administrator can use those LUNs for other hosts. See the <i>Data ONTAP Storage Management Guide</i> for information about what you need to do to prepare for taking a system using array LUNs out of service.</p>



# Configuring an active/active configuration

---

Describes how to bring up a new standard or mirrored active/active configuration for the first time. Also describes how to enable licenses, set options, configure networking, and test the configuration.

Complete the following tasks in the order shown.

1. [Bringing up the active/active configuration](#) on page 103
2. [Enabling licenses](#) on page 106
3. [Setting options and parameters](#) on page 107
4. [Configuration of network interfaces](#) on page 112
5. [Testing takeover and giveback](#) on page 123

## Bringing up the active/active configuration

The first time you bring up the active/active configuration, you must ensure that the nodes are correctly connected and powered up, and then use the setup program to configure the systems.

### Next topics

[Considerations for active/active configuration setup](#) on page 103

[Configuring shared interfaces with setup](#) on page 104

[Configuring dedicated interfaces with setup](#) on page 105

[Configuring standby interfaces with setup](#) on page 105

## Considerations for active/active configuration setup

When the setup program runs on a storage system in an active/active configuration, it prompts you to answer some questions specific for active/active configurations.

The following list outlines some of the questions about your installation that you should think about before proceeding through the setup program:

- Do you want to configure virtual interfaces (VIFs) for your network interfaces?  
For information about VIFs, see the *Data ONTAP Network Management Guide*.

**Note:** You are advised to use VIFs with active/active configurations to reduce SPOFs (single-points-of-failure).

- How do you want to configure your interfaces for takeover?

**Note:** If you do not want to configure your network for use in an active/active configuration when you run setup for the first time, you can configure it later. You can do so either by running setup

again, or by using the `ifconfig` command and editing the `/etc/rc` file manually. However, you must provide at least one local IP address to exit setup.

### Related concepts

[Configuration of network interfaces](#) on page 112

### Related tasks

[Configuring shared interfaces with setup](#) on page 104

[Configuring dedicated interfaces with setup](#) on page 105

[Configuring standby interfaces with setup](#) on page 105

## Configuring shared interfaces with setup

During setup of the storage system, you can assign an IP address to a network interface and assign a partner IP address that the interface takes over if a failover occurs.

### Steps

1. Enter the IP address for the interface you are configuring.

For example:

```
Please enter the IP address for Network Interface e0 []:nnn.nn.nn.nnn
:nnn.nn.nn.nnn is the local address for the node you are configuring.
```

**Note:** The addresses for the local node and partner node can reside on different subnetworks.

2. Enter the netmask for the interface you are configuring, or press Return if the default value is correct.

For example:

```
Please enter the netmask for Network Interface e1 [255.255.0.0]:
```

3. Specify that this interface is to take over a partner IP address.

For example:

```
Should interface e1 take over a partner IP address during failover? [n]: y
```

4. Enter the IP address or interface name of the partner.

For example:

```
Please enter the IP address or interface name to be taken over by e1 []: :nnn.nn.nn.nnn
```

**Note:** If the partner is a VIF, you must use the interface name.

**Note:** The addresses for the local node and partner node can reside on different subnetworks.



## Configuring dedicated interfaces with setup

You can assign a dedicated IP address to a network interface, so that the interface does not have a partner IP address.

### About this task

This procedure is performed during setup of the storage system.

### Steps

1. Enter the IP address for the interface you are configuring.

For example:

```
Please enter the IP address for Network Interface e0 []::nnn.nn.nn.nnn
:nnn.nn.nn.nnn is the local address for the node you are configuring.
```

2. Enter the netmask for the interface you are configuring, or press Return if the default value is correct.

For example:

```
Please enter the netmask for Network Interface e1 [255.255.0.0]:
```

3. Specify that this interface does not take over a partner IP address.

For example:

```
Should interface e1 take over a partner IP address during failover? [n]: n
```

## Configuring standby interfaces with setup

You can assign a standby IP address to a network interface, so that the interface does not have a partner IP address.

### About this task

This procedure is performed during setup of the storage system.

### Steps

1. Do not enter an IP address for a standby interface; press Return.

For example:

```
Please enter the IP address for Network Interface e0 []:
```

2. Enter the netmask for the interface you are configuring, or press Return if the default value is correct.

For example:

Please enter the netmask for Network Interface e1 [255.255.0.0]:

- Specify that this interface is to take over a partner IP address.

For example:

Should interface e1 take over a partner IP address during failover? [n]: y

## Enabling licenses

You must enable the required licenses for your type of active/active configuration.

### Before you begin

The licenses you need to add depend on the type of your active/active configuration. The following table outlines the required licenses for each configuration.

**Note:** If your system is a V-Series system, you must enable the v-series license on each node in the active/active configuration.

Configuration type	Required licenses
Standard active/active configuration	cluster
Mirrored active/active configuration	<ul style="list-style-type: none"> <li>cluster</li> <li>syncmirror_local</li> </ul>
MetroCluster	<ul style="list-style-type: none"> <li>cluster</li> <li>syncmirror_local</li> <li>cluster_remote</li> </ul>

### Steps

- Enter the following command on both node consoles for each required license:

```
license add license-code
```

*license-code* is the license code you received for the feature.

- Enter the following command to reboot both nodes:

```
reboot
```

- Enter the following command on the local node console:

```
cf enable
```

4. Verify that controller failover is enabled by entering the following command on each node console:

```
cf status
```

```
Cluster enabled, filer2 is up.
```

## Setting options and parameters

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.

### Next topics

[Option types for active/active configurations](#) on page 107

[Setting matching node options](#) on page 107

[Parameters that must be the same on each node](#) on page 108

[Disabling the change\\_fsid option in MetroCluster configurations](#) on page 108

[Configuration of the hw\\_assist option](#) on page 110

## Option types for active/active configurations

Some options must be the same on both nodes in the active/active configuration, while some can be different, and some are affected by failover events.

In an active/active configuration, options are one of the following types:

- Options that must be the same on both nodes for the active/active configuration to function correctly
- Options that might be overwritten on the node that is failing over  
These options must be the same on both nodes to avoid losing system state after a failover.
- Options that should be the same on both nodes so that system behavior does not change during failover
- Options that can be different on each node

**Note:** You can find out whether an option must be the same on both nodes of an active/active configuration from the comments that accompany the option value when you enter the `option` command. If there are no comments, the option can be different on each node.

## Setting matching node options

Because some Data ONTAP options need to be the same on both the local and partner node, you need to check these options with the `options` command on each node and change them as necessary.

## Steps

1. View and note the values of the options on the local and partner nodes, using the following command on each console:

### **options**

The current option settings for the node are displayed on the console. Output similar to the following is displayed:

```
autosupport.doit DONT
autosupport.enable on
```

2. Verify that the options with comments in parentheses are set to the same value for both nodes. The comments are as follows:

```
Value might be overwritten in takeover
Same value required in local+partner
Same value in local+partner recommended
```

3. Correct any mismatched options using the following command:

```
options option_name option_value
```

**Note:** See the `na_options` man page for more information about the options.

## Parameters that must be the same on each node

Lists the parameters that must be the same so that takeover is smooth and data is transferred between the nodes correctly.

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

## Disabling the `change_fsid` option in MetroCluster configurations

In a MetroCluster configuration, you can take advantage of the `change_fsid` option in Data ONTAP to simplify site takeover when the `cf forcetakeover -d` command is used.

**About this task**

In a MetroCluster configuration, if a site takeover initiated by the `cf forcetakeover -d` command occurs, the following happens:

- Data ONTAP changes the file system IDs (FSIDs) of volumes and aggregates because ownership changes.
- Because of the FSID change, clients must remount their volumes if a takeover occurs.
- If using Logical Units (LUNs), the LUNs must also be brought back online after the takeover.

To avoid the FSID change in the case of a site takeover, you can set the `change_fsid` option to `off` (the default is `on`). Setting this option to `off` has the following results if a site takeover is initiated by the `cf forcetakeover -d` command:

- Data ONTAP refrains from changing the FSIDs of volumes and aggregates.
- Users can continue to access their volumes after site takeover without remounting.
- LUNs remain online.

**Caution:** If the option is set to `off`, any data written to the failed node that did not get written to the surviving node's NVRAM is lost. Disable the `change_fsid` option with great care.

**Step**

1. Enter the following command to disable the `change_fsid` option:

```
options cf.takeover.change_fsid off
```

By default, the `change_fsid` option is enabled (set to `on`).

**Related concepts**

[Disaster recovery using MetroCluster](#) on page 173

**Clarification of when data loss can occur when the `change_fsid` option is enabled**

Ensure that you have a good understanding of when data loss can occur before you disable the `change_fsid` option. Disabling this option can create a seamless takeover for clients in the event of a disaster, but there is potential for data loss.

If both the ISLs between the two sites in a fabric MetroCluster go down, then both the systems remain operational. However, in that scenario, client data is written only to the local plex and the plexes become unsynchronized.

If, subsequently, a disaster occurs at one site, and the `cf forcetakeover -d` command is issued, the remote plex which survived the disaster is not current. With the `change_fsid` option set to `off`, clients switch to the stale remote plex without interruption.

If the `change_fsids` option is set to `on`, the system changes the fsids when the `cf forcetakeover -d` is issued, so clients are forced to remount their volumes and can then check for the integrity of the data before proceeding.

## Configuration of the `hw_assist` option

You can use the hardware assisted takeover option to speed up takeover times. The option uses the remote management card to quickly communicate local status changes to the partner node, and has configurable parameters.

### Next topics

[Hardware-assisted takeover](#) on page 110

[Disabling and enabling the hardware-assisted takeover option](#) on page 111

[Setting the partner address for hardware-assisted takeover](#) on page 112

[Setting the partner port for hardware-assisted takeover](#) on page 112

## Hardware-assisted takeover

Hardware-assisted takeover enables systems with remote management cards to improve the speed with which takeover events are detected, thereby speeding up the takeover time.

When enabled, hardware-assisted takeover takes advantage of the remote management card capabilities to detect failures on the local machine that could require a takeover. If a failure is detected, the card sends an alert to the partner node and, depending on the type of failure, the partner performs the takeover. These alerts can speed takeover because the Data ONTAP takeover process on the partner does not have to take the time to verify that the failing system is no longer giving a heartbeat and confirm that a takeover is actually required.

The hardware-assisted takeover option (`cf.hw_assist`) is enabled by default.

### Requirements for hardware-assisted takeover

The hardware-assisted takeover feature is available only on systems that support Remote LAN Modules (RLMs) and have the RLMs installed and set up. The remote management card provides remote platform management capabilities, including remote access, monitoring, troubleshooting, logging, and alerting features.

Although a system with an RLM on both nodes provides hardware-assisted takeover on both nodes, hardware-assisted takeover is also supported on active/active configurations in which only one of the two systems has an installed RLM. The RLM does not have to be installed on both nodes in the active/active configuration. The RLM can detect failures on the system in which it is installed and provide faster takeover times if a failure occurs on the system with the RLM.

See the *Data ONTAP System Administration Guide* for information about setting up the RLM.

## System events detected by remote management

A number of events can be detected by the remote management card and generate an alert. Depending on the type of alert received, the partner node initiates takeover.

Alert	Takeover initiated upon receipt?	Description
power_loss	Yes	Power loss on the node. The remote management card has a power supply that maintains power for a short period after a power loss, allowing it to report the power loss to the partner.
l2_watchdog_reset	Yes	L2 reset detected by the system watchdog hardware.
power_off_via_rlm	Yes	The remote management card was used to power off the system.
power_cycle_via_rlm	Yes	The remote management card was used to cycle the system power off and on.
reset_via_rlm	Yes	The remote management card was used to reset the system.
abnormal_reboot	No	Abnormal reboot of the node.
loss_of_heartbeat	No	Heartbeat message from the node no longer received by the remote management card.  <b>Note:</b> This does not refer to the heartbeat messages between the nodes in the active/active configuration but the heartbeat between the node and its local remote management card.
periodic_message	No	Periodic message sent during normal hardware-assisted takeover operation.
test	No	Test message sent to verify hardware-assisted takeover operation.

## Disabling and enabling the hardware-assisted takeover option

Hardware-assisted takeover is enabled by default on systems that use an RLM. Hardware-assisted takeover speeds the takeover process by using the RLM to quickly detect potential takeover events and alerting the partner node.

### Step

1. Enter the following command to disable or enable the `cf.hw_assist` option:

```
options cf.hw_assist.enable off
```

```
options cf.hw_assist.enable on
```

## Setting the partner address for hardware-assisted takeover

The `cf.hw_assist.partner.address` option enables you to change the partner address used by the hardware-assisted takeover process on the remote management card. The default is the IP address on the e0a port of the partner. On a 31xx system, partner's e0M interface has been configured, the IP address of the e0M interface is used. If the e0M interface has not been configured, e0a is used.

### Step

1. Enter the following command to set the IP address or host name to which the hardware failure notification is sent:

```
options cf.hw_assist.partner.address address_or_hostname
```

**Note:** The hardware assisted takeover feature does not support IPv6 addresses when specifying the partner IP address in the `cf.hw_assist.partner.address address_or_hostname` option.

If a host name is specified, the host name is resolved when this command is issued.

## Setting the partner port for hardware-assisted takeover

When hardware-assisted takeover is enabled, the RLM sends hardware failure notifications to the partner. The `cf.hw_assist.partner.port` option enables you to change the partner port. The default is 4444.

### Step

1. Enter the following command to set the partner port to which the hardware failure notification is sent:

```
options cf.hw_assist.partner.port port_number
```

## Configuration of network interfaces

If you didn't configure interfaces during system setup, you need to configure them manually to ensure continued connectivity during failover.

### Next topics

[What the networking interfaces do](#) on page 113

[IPv6 considerations in an active/active configuration](#) on page 113

[Configuring network interfaces for active/active configurations](#) on page 114

[Configuring partner addresses on different subnets \(MetroClusters only\)](#) on page 119



## What the networking interfaces do

When a node in an active/active configuration 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 active/active configuration to maintain communication with the network if the partner fails.

See the *Data ONTAP Network Management Guide* for a description of available options and the function each performs.

**Note:** You should always use multiple NICs with VIFs to improve networking availability for both stand-alone storage systems and systems in an active/active configuration.

## IPv6 considerations in an active/active configuration

When enabled, IPv6 provides features such as address autoconfiguration. Using these IPv6 features requires an understanding of how these features work with the active/active configuration functionality.

For additional information about IPv6, see the *Data ONTAP Administration Guide*.

### Configuration requirements for using IPv6

To use IPv6 in an active/active configuration, IPv6 must be enabled on both nodes. If a node that does not have IPv6 enabled attempts to take over a node using IPv6, the IPv6 addresses configured on the partner's interfaces are lost because the takeover node does not recognize them.

### Using the `ifconfig` command

When using the `ifconfig` command with IPv4, the partner's interface can be mapped to a local interface or the partner's IP address. When using IPv6, you must specify the partner interface, not an IP address.

### Generation of addresses during takeover

For manually configured IPv6 addresses, during takeover, the mechanism used to configure partner's IP address remains same as in the case of IPv4.

For link-local auto-configured IPv6 addresses, during takeover, the address is auto-generated based on the partner's MAC address.

Prefix-based auto-configured addresses are also generated during takeover, based on the prefixes in router advertisements (RAs) received on the local link and on the partner's MAC address.

Duplicate Address Detection (DAD) is performed on all IPv6 partner addresses during takeover. This can potentially keep the addresses in *tentative* state for some amount of time.

### IPv6 and hardware-assisted takeover

The hardware assisted takeover feature does not support IPv6 addresses when specifying the partner IP address in the `cf.hw_assist.partner.address address_or_hostname` option.

## Configuring network interfaces for active/active configurations

Configuring network interfaces requires that you understand the available configurations for takeover and that you configure different types of interfaces (shared, dedicated, and standby) depending on your needs.

### Next topics

[Understanding interfaces in an active/active configuration](#) on page 114

[Making nondisruptive changes to the virtual interfaces](#) on page 117

[Configuring dedicated and standby interfaces](#) on page 117

## Understanding interfaces in an active/active configuration

You can configure three types of interfaces on nodes in an active/active configuration.

### Next topics

[Shared, dedicated, and standby interfaces](#) on page 114

[Interface roles in normal and takeover modes](#) on page 115

[Takeover configuration with shared interfaces](#) on page 115

[Takeover configuration with dedicated and standby interfaces](#) on page 116

[Interface types and configurations](#) on page 116

## Shared, dedicated, and standby interfaces

These different types of interfaces have different roles in normal and takeover mode.

The following table lists the three types of interface configurations that you can enable in an active/active configuration.

Interface type	Description
Shared	This type of interface supports both the local and partner nodes. It contains both the local node and partner node 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 IP address only and does not participate in network communication beyond local node support during takeover. It is paired with a standby interface.
Standby	This type of interface is on the local node, but only contains the IP address of the partner node. It is paired with a dedicated interface.

**Note:** Most active/active configuration interfaces are configured as shared interfaces because they do not require an extra NIC.

## Interface roles in normal and takeover modes

You can configure shared, dedicated, and standby interfaces in an active/active configuration. Each type has a different role in normal and takeover mode.


The following table shows the role of each interface type in normal and takeover mode.

Interface type	Normal mode	Takeover mode
Shared	Supports the identity of the local node	Supports the identity of both the local node and the failed node
Dedicated	Supports the identity of the local node	Supports the identity of the local node
Standby	Idle	Supports the identity of the failed node

## Takeover configuration with shared interfaces

You can configure two NICs on to provide two shared interfaces to each node.

In the following configuration illustration, you use two NICs to provide the two interfaces.



**Figure 25: Takeover configuration with two shared interfaces**


If Node 1 fails, interface e0 on Node 1 stops functioning, but the secondary address on e0 on Node 2 handles the Node 1 network connection with the 230 network.

If Node 2 fails, e0 on Node 2 stops functioning, but e0 on Node 1 substitutes for the failed interface and handles the Node 2 network connection with the 230 network.

### Takeover configuration with dedicated and standby interfaces

With two NICs on each node, one can provide a dedicated interface and the other can act as a standby interface.

In the following configuration illustration, you use two NICs for each interface, one on each storage system. One NIC acts as a dedicated interface and the other acts as a standby interface.



**Figure 26: Takeover configuration with dedicated and standby interfaces**

If Node 1 fails, interface e0 on Node 1 stops functioning, but e0 on Node 2 substitutes for the failed interface and handles the Node 1 network connection with the 230 network.

If Node 2 fails, e1 on Node 2 stops functioning, but e1 on Node 1 substitutes for the failed interface and handles the Node 2 network connection with the 230 network.

### Interface types and configurations

This table lists the configurations supported by each type of interface in an active/active configuration.

Interface	Shared	Dedicated	Standby	Partner parameter
Ethernet	X	X	X	IP address or interface name
Gigabit Ethernet	X	X	X	IP address or interface name
Virtual interface	X	X	X	Virtual interface name
VLAN interface	X	X	X	IP address or interface name

**Note:** Some storage systems, such as the 31xx systems, include an e0M interface that is dedicated to management traffic. This port can be partnered in an active/active configuration in the same way as a regular Ethernet interface.

## Making nondisruptive changes to the virtual interfaces

You can use the `cf takeover` and `cf giveback` commands to make changes to VIFs in the active/active configuration in a nondisruptive manner.

Changes to the `/etc/rc` file require a reboot to make the changes effective. You can use the `cf takeover` and `cf giveback` commands to take over one node in the active/active configuration, causing it to reboot while its storage is taken over by the partner.

### Steps

1. Edit the `/etc/rc` file on the desired node to modify the VIF.
2. From the partner node (the partner of the node on which you performed step 1), enter the following command:

```
cf takeover
```

3. Enter the following command:

```
cf giveback
```

The node on which the changes were made reboots and its `etc/rc` file is reread. The `rc` file is responsible for creating the VIF.

4. Repeat these steps, making any required changes to the `/etc/rc` file on the partner node.

### Related tasks

[Configuring dedicated and standby interfaces](#) on page 117

## Configuring dedicated and standby interfaces

You can configure dedicated and standby interfaces for an active/active configuration, two on each node, so that even in the event of a takeover each node still has a dedicated interface.

Both nodes in the active/active configuration must have interfaces that access the same collection of networks and subnetworks.

You must gather the following information before configuring the interfaces:

- The IP address for both the local node and partner node.
  - Note:** For MetroCluster configurations, if you use the `/etc/mcrc` file and enable the `cf.takeover.use.mcrc_file`, the addresses for the local node and partner node can reside on different subnetworks.
- The netmask for both the local node and partner node.
- The MTU size for both the local node and partner node. The MTU size must be the same on both the local and partner interface.

**Note:** You should always use multiple NICs with VIFs to improve networking availability for both stand-alone storage systems and systems in an active/active configuration.

Keep in mind that you can use interface names for specifying all interfaces.

If you configured your interfaces using setup when you first applied power to your storage systems, you do not need to configure them again.

**Note:** For information about configuring an active/active configuration to use FC, see the *Data ONTAP Block Access Management Guide for iSCSI and FC*.

## Steps

1. On nodeA, enter the following command on the command line and also enter it in the `/etc/rc` file, so that the command is permanent:

```
ifconfig interfaceA1addressA1 {other_options}
```

*interfaceA1* is the name of the dedicated local interface for nodeA.

*addressA1* is the IP address of the dedicated local interface for nodeA.

*other\_options* denotes whatever other options are needed to correctly configure the interface in your network environment.

The dedicated local interface for nodeA is configured.

2. Also on nodeA, enter the following command on the command line and in the `/etc/rc` file:

```
ifconfig interfaceA2 partner addressB1
```

*interfaceA2* is the name of the standby interface for nodeA.

*addressB1*

**Note:** When you configure virtual interfaces for takeover, you must specify the interface name and not the IP address.

The standby interface for nodeA is configured to take over the dedicated interface of nodeB on takeover.

3. On nodeB, enter the following command on the command line and in the `/etc/rc` file:

```
ifconfig interfaceB1addressB1 {other_options}
```

*interfaceB1* is the name of the dedicated local interface for nodeB.

*addressB1* is the IP address of the dedicated local interface for nodeB.

*other\_options* denotes whatever other options are needed to correctly configure the interface in your network environment.

The dedicated local interface for nodeB is configured.

4. Also on nodeB, enter the following command on the command line and in the `/etc/rc` file:

```
ifconfig interfaceB2 partner addressA1
```

*interfaceB2* is the name of the standby interface on nodeB.

*addressA1* is the IP address or interface name of the dedicated interface for nodeA.

**Note:** When you configure virtual interfaces for takeover, you must specify the interface name and not the IP address.

The standby interface on nodeB is configured to take over the dedicated interface of nodeA on takeover.

If desired, configure your interfaces for automatic takeover in case of NIC failure.

## Configuring partner addresses on different subnets (MetroClusters only)

On MetroCluster configurations, you can configure partner addresses on different subnets. To do this, you must create a separate `/etc/mcrc` file and enable the `cf.takeover.use_mcrc_file` option. When taking over its partner, the node uses the partner's `/etc/mcrc` file to configure partner addresses locally. These addresses will reside on the local subnetwork.

### Next topics

[The `/etc/mcrc` file](#) on page 119

[Creating an `/etc/mcrc` file](#) on page 121

[Setting the system to use the partner's `/etc/mcrc` file at takeover](#) on page 123

## The `/etc/mcrc` file

The `/etc/mcrc` file, in conjunction with the `cf.takeover.use_mcrc_file` option, should be used on MetroCluster configurations in which the partner nodes reside on separate subnetworks.

Normally, when a node (for example, nodeA) takes over its partner (nodeB), nodeA runs nodeB's `/etc/rc` file to configure interfaces on nodeA to handle incoming traffic for the taken-over partner, nodeB. This requires that the local and partner addresses are on the same subnetwork.

When the `cf.takeover.use_mcrc_file` option is enabled on nodeA, nodeA will use nodeB's `/etc/mcrc` file upon takeover, instead of nodeB's `/etc/rc` file. The `ifconfig` commands in the `/etc/mcrc` file can configure IP addresses on nodeA's subnetwork. With the correct `ifconfig`, virtual IP (VIP), and routing commands in the `/etc/mcrc` file, the resulting configuration allows hosts connecting to nodeB to connect to node A.

**Note:** The `/etc/mcrc` file must be maintained manually and kept in sync with the `/etc/rc` file.

### Example `/etc/rc` and `/etc/mcrc` files

NodeA's `/etc/rc` file, which configures its local addresses and a partner address (which matches the address configured in NodeB's `/etc/mcrc` file):

```
hostname nodeA
```

```

ifconfig e0a 10.1.1.1 netmask 255.255.255.0
ifconfig e0a partner 10.1.1.100
ifconfig vip add 5.5.5.5
route add default 10.1.1.50 1
routed on
options dns.domainname netapp.com
options dns.enable on
options nis.enable off
savecore

```

NodeA's `/etc/mcrc` file, which configures a partner address on NodeB's subnetwork:

```

hostname nodeA
ifconfig e0a 20.1.1.200 netmask 255.255.255.0
ifconfig vip add 5.5.5.5
route add default 20.1.1.50 1
routed on
options dns.domainname netapp.com
options dns.enable on
options nis.enable off
savecore

```

NodeB's `/etc/rc` file, which configures its local addresses and a partner address (which matches the address configured in NodeA's `/etc/mcrc` file)::

```

hostname nodeB
ifconfig e0a 20.1.1.1 netmask 255.255.255.0
ifconfig e0a partner 20.1.1.200
ifconfig vip add 7.7.7.7
route add default 20.1.1.50 1
routed on
options dns.domainname netapp.com
options dns.enable on
options nis.enable off
savecore

```

NodeB's `/etc/mcrc` file, which configures a partner address on NodeA's subnetwork:

```

hostname nodeB
ifconfig e0a 10.1.1.100 netmask 255.255.255.0
ifconfig vip add 7.7.7.7
route add default 10.1.1.50 1
routed on
options dns.domainname netapp.com
options dns.enable on

```



```
options nis.enable off
savecore
```

## Creating an `/etc/mcrc` file

You should create an `/etc/mcrc` file on each node of your MetroCluster configuration if the nodes are on separate subnetworks.

### Steps

1. Create an `/etc/mcrc` file on one node (nodeA) and place it in the `/etc` directory.

You might want to create the `/etc/mcrc` file by copying the `/etc/rc` file.

**Note:** The `/etc/mcrc` file must be configured manually. It is not updated automatically. It must include all commands necessary to implement the network configuration on the partner node in the event the node is taken over by the partner.

2. Enter the following commands in nodeA's `/etc/mcrc` file:

```
hostname nodeA
ifconfig interface MetroCluster-partner-address netmask netmask
ifconfig vip add virtual-IP-address
route add default route-for-MetroCluster-partner-address 1
routed on
other-required-options
```

*interface* is the interface on which the corresponding *MetroCluster-partner-address* will reside.

*MetroCluster-partner-address* is partner address of nodeB. It corresponds to the partner address configured by an `ifconfig` command in nodeB's `/etc/rc` file.

*virtual-IP-address* is the virtual address of the partner (nodeB).

*other-required-options* denotes whatever other options are needed to correctly configure the interface in your network environment.

### Example

Example of nodeA's `/etc/mcrc` file:

```
hostname nodeA
ifconfig e0a 20.1.1.200 netmask 255.255.255.0
ifconfig vip add 5.5.5.5
route add default 20.1.1.50 1
routed on
options dns.domainname netapp.com
options dns.enable on
```

```
options nis.enable off
savecore
```

3. Create an `/etc/mcrc` file on the other node (nodeB) and place it in the `/etc` directory.

The `/etc/mcrc` file must include an `ifconfig` command that configures the address that corresponds to the address specified in the `partner` parameter in the partner node's `/etc/rc`.

You might want to create the `/etc/mcrc` file by copying the `/etc/rc` file.

**Note:** The `/etc/mcrc` file must be configured manually. It is not updated automatically. It must include all commands necessary to configure the

Enter the result of your step here (optional).

4. Enter the following commands in nodeB's `/etc/mcrc` file:

```
hostname nodeB
ifconfig interface MetroCluster-partner-address netmask netmask
ifconfig vip add virtual-IP-address
route add default route-for-MetroCluster-partner-address 1
routed on
other-required-options
```

`interface` is the interface on which the corresponding `MetroCluster-partner-address` will reside.

`MetroCluster-partner-address` is partner address of nodeA. It corresponds to the partner address configured by an `ifconfig` command in nodeA's `/etc/rc` file.

`virtual-IP-address` is the virtual address of the partner (nodeA).

`other-required-options` denotes whatever other options are needed to correctly configure the interface in your network environment.

### Example

Example of nodeB's `/etc/mcrc` file:

```
hostname nodeB
ifconfig e0a 10.1.1.100 netmask 255.255.255.0
ifconfig vip add 7.7.7.7
route add default 10.1.1.50 1
routed on
options dns.domainname netapp.com
options dns.enable on
```

```
options nis.enable off
savecore
```

## Setting the system to use the partner's /etc/mcrc file at takeover

You must enable the `cf.takeover.use_mcrc_file` option to cause the system to use the partner's `/etc/mcrc` in the event that the local system takes over the partner. This allows the partner IP addresses to reside on separate subnetworks. This option should be set on both nodes in the MetroCluster.

### Step

1. Enter the following command on both nodes:

```
options cf.takeover.use_mcrc_file on
```

The default is off.

## Testing takeover and giveback

After you configure all aspects of your active/active configuration, verify that it operates as expected.

### Steps

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. Enter the following command from the local node console:

```
cf takeover
```

The local node takes over the partner node and gives the following output:

```
takeover completed
```

4. Test communication between the local node and partner node.

### Example

You can use the `fcstat device_map` command to ensure that one node can access the other node's disks.

5. Give back the partner node by entering the following command:

```
cf giveback
```

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
```

6. Proceed depending on whether you got the message that giveback was completed successfully.

<b>If takeover and giveback...</b>	<b>Then...</b>
<b>Is completed successfully</b>	Repeat Step 2 through Step 5 on the partner node
<b>Fails</b>	Attempt to correct the takeover or giveback failure

# Management of takeover and giveback

---

An active/active configuration allows one partner to take over the storage of the other, and return the storage using the giveback operation. Management of the nodes in the active/active configuration differs depending on whether one partner has taken over the other, and the takeover and giveback operations themselves have different options.

## Next topics

[How takeover and giveback work](#) on page 125

[Management of an active/active configuration in normal mode](#) on page 127

[Configuration of when takeover occurs](#) on page 133

[Managing an active/active configuration in takeover mode](#) on page 138

[Managing emulated nodes](#) on page 139

[Performing dumps and restores for a failed node](#) on page 143

[Giveback operations](#) on page 144

[Downloading and running the HA Configuration Checker utility](#) on page 149

[Troubleshooting takeover or giveback failures](#) on page 149

## How takeover and giveback work

Takeover is the process in which a node takes over the storage of its partner. Giveback is the process in which the storage is returned to the partner. You can initiate the processes in different ways. A number of things that affect the active/active configuration occur when takeover and giveback take place.

## Next topics

[When takeovers occur](#) on page 125

[What happens during takeover](#) on page 126

[What happens after takeover](#) on page 126

[What happens during giveback](#) on page 127

## When takeovers occur

The conditions under which takeovers occur depend on how you configure the active/active configuration.

Takeovers can be initiated when one of the following conditions occur:

- A node is in an active/active configuration that is configured for immediate takeover on panic, and it undergoes a software or system failure that leads to a panic.

- A node that is in an active/active configuration undergoes a system failure (for example, a loss of power) and cannot reboot.
  - Note:** If the storage for a node also loses power at the same time, a standard takeover is not possible. For MetroClusters, you can initiate a forced takeover in this situation.
- There is a mismatch between the disks, array LUNs, or both that one node can see and those that the other node can see.
- One or more network interfaces that are configured to support failover become unavailable.
- A node cannot send heartbeat messages to its partner. This could happen if the node experienced a hardware or software failure that did not result in a panic but still prevented it from functioning correctly.
- You halt one of the nodes without using the `-f` flag.
- You initiate a takeover manually.

## What happens during takeover

When a takeover occurs, the unimpaired partner node takes over the functions and disk drives of the failed node by creating an emulated storage system.

The emulated system performs the following tasks:

- Assumes the identity of the failed node
- Accesses the failed node's disks, array LUNs, or both and serves its data to clients

The partner node maintains its own identity and its own primary functions, but also handles the added functionality of the failed node through the emulated node.

**Note:** When a takeover occurs, existing CIFS sessions are terminated. A graceful shutdown of the CIFS sessions is not possible, and some data loss could occur for CIFS users.

## What happens after takeover

After a takeover occurs, you view the surviving partner as having two identities, its own and its partner's, that exist simultaneously on the same storage system. Each identity can access only the appropriate volumes and networks. You can send commands or log in to either storage system by using the `rsh` command, allowing remote scripts that invoke storage system commands through a Remote Shell connection to continue to operate normally.

### Access with rsh

Commands sent to the failed node through a Remote Shell connection are serviced by the partner node, as are `rsh` command login requests.

## Access with telnet

If you log in to a failed node through a Telnet session, you see a message alerting you that your storage system failed and to log in to the partner node instead. If you are logged in to the partner node, you can access the failed node or its resources from the partner node by using the partner command.

## What happens during giveback

After the partner node is repaired and is operating normally, you can use the `giveback` command to return operation to the partner.

When the failed node is functioning again, the following events can occur:

- You initiate a `giveback` command that terminates the emulated node on the partner.
- The failed node resumes normal operation, serving its own data.
- The active/active configuration resumes normal operation, with each node ready to take over for its partner if the partner fails.

## Management of an active/active configuration in normal mode

You manage an active/active configuration in normal mode by performing a number of management actions.

### Next topics

[Monitoring active/active configuration status](#) on page 127

[Monitoring the hardware-assisted takeover feature](#) on page 128

[Description of active/active configuration status messages](#) on page 130

[Displaying the partner's name](#) on page 131

[Displaying disk and array LUN information on an active/active configuration](#) on page 131

[Enabling and disabling takeover](#) on page 132

[Enabling and disabling automatic takeover of a panicked partner](#) on page 132

[Halting a node without takeover](#) on page 133

## Monitoring active/active configuration status

You can use commands on the local node to determine whether the controller failover feature is enabled and whether the other node in the active/active configuration is up.

### Step

1. Enter the following command:

```
cf status
```

**Example**

```
node1>
cf status
Cluster enabled, node2 is up.
```

**Note:** Data ONTAP can disable controller failover if a software or hardware problem exists that prevents a successful takeover. In this case, the message returned from the `cf status` command describes the reason why failover is disabled.

This verifies the link between the nodes and tells you that both `filer1` and `filer2` are functioning and available for takeover.

**Monitoring the hardware-assisted takeover feature**

You can check and test the hardware-assisted takeover configuration using the `hw_assist` command. You can also use the command to review statistics relating to hardware-assisted takeover.

**Next topics**

[Checking status](#) on page 128

[Testing the hardware-assisted takeover configuration](#) on page 129

[Checking hardware-assisted takeover statistics](#) on page 129

**Checking status**

You can check the status of the hardware-assisted takeover configuration with the `cf hw_assist status` command. It shows the current status for the local and partner nodes.

**Step**

1. Enter the following command to display the hardware-assisted takeover status:

```
cf hw_assist status
```

**Example hardware-assisted takeover status**

The following example shows output from the `cf hw_assist status` command:

```
Local Node Status - ha1
    Active: Monitoring alerts from partner(ha2)
    port 4004 IP address 172.27.1.14

Partner Node Status - ha2
```



```
Active: Monitoring alerts from partner(hal)
port 4005 IP address 172.27.1.15
```

## Testing the hardware-assisted takeover configuration

You can test the hardware-assisted takeover configuration with the `cf hw_assist test` command.

The `cf hw_assist test` command sends a test alert to the partner. If the alert is received the partner sends back an acknowledgment, and a message indicating the successful receipt of the test alert is displayed on the console.

### Step

1. Enter the following command to test the hardware-assisted takeover configuration:

```
cf hw_assist test
```

Depending on the message received from the `cf hw_assist test` command, you might need to reconfigure options so that the active/active configuration and the remote management card are operating.

## Checking hardware-assisted takeover statistics

You can display statistics about hardware-assisted takeovers with the `cf hw_assist stats` command.

### Step

1. Enter the following command to display or clear the hardware-assisted takeover statistics, respectively:

```
cf hw_assist stats
```

```
cf hw_assist stats clear
```

### Example hardware-assisted takeover statistics

The following example shows output from the `cf hw_assist stats` command on a system that has received a variety of alerts from the partner:

```
# cf hw_assist: stats
Known hw_assist alerts received from partner
  alert type          alert event          num of alerts
  -----          -
```

```

system_down      post_error      0
system_down      power_loss      0
system_down      abnormal_reboot 0
system_down      l2_watchdog_reset 0
system_down      power_off_via_rlm 0
system_down      power_cycle_via_rlm 0
system_down      reset_via_rlm 0
keep_alive       loss_of_heartbeat 0
keep_alive       periodic_message 18
test            test            6

Unknown hw_assist alerts received from partner

  Partner nvramid mismatch alerts 5

  Shared secret mismatch alerts 10

  Unknown alerts 23

Number of times hw_assist alerts throttled: 3

```

### Description of active/active configuration status messages

The `cf status` command displays information about the status of the active/active configuration.

The following table shows some of the messages that the `cf status` command can display.

Message	Meaning
cluster enabled, partner_name is up.	The active/active configuration is operating normally.
partner_name_1 has taken over partner_name_2.	One node took over the other node.
Interconnect not present.	The system does not recognize the existence of a cluster interconnect adapter.
Interconnect is down.	The cluster interconnect adapter cannot access the partner. This might be due to cabling problems or the partner might be down.
Interconnect is up.	The cluster interconnect adapter is active and can transmit data to the partner.

Message	Meaning
partner_name_1 has detected a mailbox disk error, takeover of partner_name_2 disabled.	One node cannot access multiple mailbox disks. Check access to both the local and partner root volumes and mirrors, if they exist. Also check for disk or FC-AL problems or offline storage adapters.
partner_name_2 may be down and has disabled takeover by partner_name_1.	One node might be down.
Version mismatch	The partner node has an incompatible version of Data ONTAP.
partner_name_1 is attempting takeover of partner_name_2. takeover is in module n of N modules.	A takeover is being attempted (includes information about how far the takeover has progressed).
partner_name_1 has taken over partner_name_2, giveback in progress. giveback is in module n of N modules.	A giveback is being attempted (includes information about how far the giveback has progressed).
partner_name_1 has taken over partner_name_2, partner_name_2 is ready for giveback.	The takeover node received information that the failed node is ready for giveback.
partner_name_1 has taken over partner_name_2, partner_name_2 is ready for giveback. Automatic giveback is disabled due to exceeding retry count.	The takeover node received information that the failed node is ready for giveback, but giveback cannot take place because the number of retries exceeded the limit.

## Displaying the partner's name

You can display the name of the other node with the `cf partner` command.

### Step

1. Enter the following command:

```
cf partner
```

**Note:** If the node does not yet know the name of its partner because the active/active configuration is new, this command returns “partner”.

## Displaying disk and array LUN information on an active/active configuration

To find out about the disks, array LUNs, or both on both the local and partner node, you can use the `sysconfig` and `aggr status` commands, which display information about both nodes.

### About this task

For each node, the `sysconfig` command output displays disks on both FC-AL loop A and FC-AL loop B:

- The information about disks on FC-AL loop A is the same as for storage systems not in an active/active configuration.
- The information about disks on FC-AL loop B is for hardware only; the `sysconfig` command displays information about the adapters supporting the disks. The command does not show whether a disk on FC-AL loop B is a file system disk, spare disk, or parity disk.

### Step

1. Enter one of the following commands:

```
sysconfig -r
```

or

```
aggr status -r
```

## Enabling and disabling takeover

You might want to use the `cf disable` command to disable takeover if you are doing maintenance that typically causes a takeover. You can reenable takeover with the `cf enable` command after you finish maintenance.

### Step

1. Enter the following command:

```
cf enable|disable
```

Use `cf enable` to enable takeover or `cf disable` to disable takeover.

**Note:** You can enable or disable takeover from either node.

## Enabling and disabling automatic takeover of a panicked partner

A node can be configured so it takes over immediately when its partner panics. This shortens the time between the initial failure and when service is fully restored, because the takeover can be quicker than recovery from the panic, but the subsequent giveback causes another brief outage.

### Steps

1. Ensure that you enabled controller takeover by entering the following command:

```
cf enable
```

2. Enter the following command to enable or disable takeover on panic:

```
options cf.takeover.on_panic [on|off]
```

- on** Enables immediate takeover of a failed partner or off to disable immediate takeover. This is the default value.
- off** Disables immediate takeover. If you disable this option, normal takeover procedures apply. The node still takes over if its partner panics, but might take longer to do so.

**Note:** If you enter this command on one node, the value applies to both nodes.

The setting of this option is persistent across reboots.

## Halting a node without takeover

You can halt the node and prevent its partner from taking over.

### About this task

You can halt the node and prevent its partner from taking over. For example, you might need to perform maintenance on both the storage system and its disks and want to avoid an attempt by the partner node to write to those disks.

### Step

1. Enter the following command:

```
halt -f
```

## Configuration of when takeover occurs

You can control when takeovers happen by setting the appropriate options.

### Next topics

[Reasons for takeover](#) on page 133

[Commands for performing a takeover](#) on page 135

[Specifying the time period before takeover](#) on page 136

[How disk shelf comparison takeover works](#) on page 137

[Configuring VIFs or interfaces for automatic takeover](#) on page 137

[Takeover of vFiler units and the vFiler unit limit](#) on page 137

## Reasons for takeover

Describes the different system errors that can cause a takeover and the commands to configure the related options.

Takeovers can happen for several reasons. Some system errors must cause a takeover; for example, when a system in an active/active configuration loses power, it automatically fails over to the other node.

However, for some system errors, a takeover is optional, depending on how you set up your active/active configuration. The following table outlines which system errors can cause a takeover to occur, and whether you can configure the active/active configuration for that error.

System error	Option used to configure	Default value	Notes
A node undergoes a system failure and cannot reboot.	<code>cf.takeover.on_failure</code> set to On	On	You should leave this option enabled unless instructed otherwise by technical support.
A node undergoes a software or system failure leading to a panic.	<code>cf.takeover.on_panic</code> set to On	Off, unless: <ul style="list-style-type: none"> <li>• FCP or iSCSI is licensed</li> <li>• The system is a FAS270.</li> </ul>	
There is a mismatch between the disks, array LUNs, or both that one node can see and those that the other node can see.	<code>cf.takeover.on_disk_shelf_miscompare</code> set to On	Off	
All the network interface cards (NICs) or vifs enabled for negotiated failover on a node failed.	<code>cf.takeover.on_network_interface_failure</code> set to On, <code>cf.takeover.on_network_interface_failure.policy</code> set to <code>all_nics</code>	By default, takeover on network failure is disabled.	To enable a network interface for negotiated failover, you use the <code>ifconfig if_name -nfo</code> command. For more information, see the <i>Data ONTAP MultiStore Management Guide</i> .
One or more of the NICs or vifs enabled for negotiated failover failed.  <b>Note:</b> If interfaces fail on both nodes in the active/active configuration, takeover won't occur.	<code>cf.takeover.on_network_interface_failure</code> set to On  <code>cf.takeover.on_network_interface_failure.policy</code> set to <code>any_nic</code>	By default, takeover on network failure is disabled.	To enable a network interface or vif for negotiated failover, you use the <code>ifconfig if_name -nfo</code> command. For more information, see the <i>Data ONTAP MultiStore Management Guide</i> .

System error	Option used to configure	Default value	Notes
A node fails within 60 seconds of booting up.	<code>cf.takeover.on_short_uptime</code> set to On	On	Changing the value of this option on one node automatically updates the option on the partner node.
A node cannot send heartbeat messages to its partner.	n/a		You cannot prevent this condition from causing a takeover.
You halt one of the nodes <i>without</i> using the <code>-f</code> flag.	n/a		You cannot prevent this condition from causing a takeover. If you include the <code>-f</code> flag, the takeover is prevented.
You initiate a takeover manually using the <code>cf takeover</code> command.	n/a		You cannot prevent this condition from causing a takeover.

### Related concepts

[How disk shelf comparison takeover works](#) on page 137

### Related tasks

[Enabling and disabling automatic takeover of a panicked partner](#) on page 132

## Commands for performing a takeover

Lists and describes the commands you can use when initiating a takeover. You can initiate a takeover on a node in an active/active configuration to perform maintenance on that node while still serving the data on its disks to users.

You can initiate a takeover on a node in an active/active configuration to perform maintenance on that node while still serving the data on its disks, array LUNs, or both to users. The following table lists and describes the commands you can use when initiating a takeover.

Command	Description
<code>cf takeover</code>	Initiates a takeover of the partner of the local node. Takeover is aborted if a core dump is in progress on the partner (if the <code>cf.takeover.on_panic</code> option is set to off). The takeover starts either after the partner halts successfully or after a timeout.
<code>cf takeover -f</code>	Initiates an immediate takeover of the partner of the local node regardless of whether the other node is dumping its core. The partner node is not allowed to halt gracefully.

Command	Description
<code>cf forcetakeover</code>	Tells the cluster monitor to ignore some configuration problems that would otherwise prevent a takeover, such as unsynchronized NVRAM due to a faulty cluster interconnect connection. It then initiates a takeover of the partner of the local node.
<code>cf forcetakeover -d</code>	<p>Initiates a takeover of the local partner even in the absence of a quorum of partner mailbox disks or partner mailbox LUNs.</p> <p>The <code>cf forcetakeover -d</code> command is valid only if the <code>cluster_remote</code> license is enabled.</p> <p><b>Attention:</b> Use the <code>-d</code> option only after you verify that the partner is down.</p> <p><b>Note:</b> The <code>-d</code> option is used in conjunction with RAID mirroring to recover from disasters in which one partner is not available. For more information, see the <i>Data ONTAP Data Protection Online Backup and Recovery Guide</i>.</p>
<code>cf takeover -n</code>	Initiates a takeover for a nondisruptive upgrade. For more information, see the <i>Data ONTAP Upgrade Guide</i> .

## Specifying the time period before takeover

You can specify how long (in seconds) a partner in an active/active configuration can be unresponsive before the other partner takes over.

### About this task

Both partners do not need to have the same value for this option. Thus, you can have one partner that takes over more quickly than the other.

**Note:** If your active/active configuration is failing over because one of the nodes is too busy to respond to its partner, increase the value of the `cf.takeover.detection.seconds` option on the partner.

### Step

1. Enter the following command:

```
options cf.takeover.detection.seconds number_of_seconds
```

The valid values for `number_of_seconds` are 10 through 180; the default is 15.



**Note:** If the specified time is less than 15 seconds, unnecessary takeovers can occur, and a core might not be generated for some system panics. Use caution when assigning a takeover time of less than 15 seconds.

## How disk shelf comparison takeover works

Describes the way a node uses disk shelf comparison with its partner node to determine if it is impaired.

When communication between nodes is first established through the cluster interconnect adapters, the nodes exchange a list of disk shelves that are visible on the A and B loops of each node. If, later, a system sees that the loop B disk shelf count on its partner is greater than its local loop A disk shelf count, the system concludes that it is impaired and prompts its partner to initiate a takeover.

**Note:** Disk shelf comparison does not function for Active/active configurations using software-based disk ownership, or fabric-attached MetroClusters.

## Configuring VIFs or interfaces for automatic takeover

After you configure your interfaces or VIFs to allow takeovers and givebacks to be completed successfully, you can also optionally configure them to trigger automatic takeover if any or all of them experience a persistent failure.

### Steps

1. For every VIF or interface on which you want to enable automatic takeover, enter the following command:
 

```
ifconfig interface_name nfo
```
2. Update the `/etc/rc` file with the command that you entered so that your changes persist across reboots.
3. The default policy is that takeover only occurs if all the NICs or VIFs on a node that are configured for automatic takeover fail. If you want takeover to occur if any NIC or VIF configured for automatic takeover fails, enter the following command:

```
options cf.takeover.on_network_interface_failure.policy any_nic
```

4. Enter the following command to enable takeover on interface failures:

```
options cf.takeover.on_network_interface_failure enable
```

**Note:** If interfaces fail on both nodes in the active/active configuration, takeover won't occur.

## Takeover of vFiler units and the vFiler unit limit

The vFiler limit, set with the `vfiler limit` command, determines how many vFiler units can exist on a system. In an active/active configuration, if the two systems have different vFiler limits, some vFiler units might not be taken over in the event of a takeover.

When performing a takeover, a system can take over only the number of vFiler units that were specified by that system's vFiler unit limit. For example, if the limit is set to 5, the system can only take over five vFiler units from the partner. If the partner that is being taken over had a higher vFiler limit, some vFiler units will not be successfully taken over.

For more information about setting the vFiler limit, see the *Data ONTAP MultiStore Management Guide*.

## Managing an active/active configuration in takeover mode

You manage an active/active configuration in takeover mode by performing a number of management actions.

### Next topics

[Determining why takeover occurred](#) on page 138

[Statistics in takeover mode](#) on page 138

## Determining why takeover occurred

You can use the `cf status` command to determine why a takeover occurred.

### Step

1. At the takeover prompt, enter the following command:

```
cf status
```

This command can display the following information:

- Whether controller failover is enabled or disabled
- Whether a takeover is imminent due to a negotiated failover
- Whether a takeover occurred, and the reason for the takeover

## Statistics in takeover mode

Explains differences in system statistics when in takeover mode.

In takeover mode, statistics for some commands differ from the statistics in normal mode in the following ways:

- Each display reflects the sum of operations that take place on the takeover node plus the operations on the failed node. The display does not differentiate between the operations on the takeover node and the operations on the failed node.
- The statistics displayed by each of these commands are cumulative.

- After giving back the failed partner's resources, the takeover node does not subtract the statistics it performed for the failed node in takeover mode.
- The giveback does not reset (zero out) the statistics.  
To get accurate statistics from a command after a giveback, you can reset the statistics as described in the man page for the command you are using.

**Note:** You can have different settings on each node for SNMP options, but any statistics gathered while a node was taken over do not distinguish between nodes.

## Managing emulated nodes

An emulated node is a software copy of the failed node that is hosted by the takeover node. You access the emulated node in partner mode by using the `partner` command.

### Next topics

[Management exceptions for emulated nodes](#) on page 139

[Accessing the emulated node from the takeover node](#) on page 139

[Assessing the emulated node remotely](#) on page 141

[Emulated node command exceptions](#) on page 141

## Management exceptions for emulated nodes

The management of disks and array LUNs and some other tasks are different when you are managing an emulated node.

You manage an emulated node as you do any other storage system, including managing disks or LUNs, with the following exceptions, which are described in greater detail later in this section:

- An emulated node can access only its own disks or LUNs.
- Some commands are unavailable.
- Some displays differ from normal displays.

## Accessing the emulated node from the takeover node

You access the emulated node from the takeover node in takeover mode with the `partner` command.

### About this task

You can issue the `partner` command in two forms:

- Using the `partner` command without an argument  
This toggles between *partner mode*, in which you manage the emulated node, and *takeover mode*, in which you manage the takeover node.

- Using the `partner` command with a Data ONTAP command as an argument  
This executes the command on the emulated node in partner mode and then returns to takeover mode.

### Next topics

[Accessing the remote node using the partner command without arguments](#) on page 140

[Accessing the takeover node with the partner command with arguments](#) on page 140

## Accessing the remote node using the partner command without arguments

Describes how to use the `partner` command to toggle between the partner mode, in which commands are executed on the partner node, and takeover mode.

### Step

1. From the takeover prompt, enter the following command:

```
partner
```

The prompt changes to the partner mode prompt, which has the following form:  
`emulated_node/takeover_node>`

### Example showing the change to partner mode

The following example shows the change from takeover mode to partner mode and back:

```
filer1(takeover)> partner
Login from console: filer2
Thu Aug 20 16:44:39 GMT [filer1: rc]: Login from console: filer2
filer2/filer1> partner
Logoff from console: filer2
filer1(takeover)> Thu Aug 20 16:44:54 GMT [filer1: rc]: Logoff from
console: filer2
filer1(takeover)>
```

## Accessing the takeover node with the partner command with arguments

Describes how to use the `partner` command with a Data ONTAP command as an argument.

### Step

1. From the takeover prompt, enter the following command:

```
partner command
```

`command` is the command you want to initiate on the emulated node.

**Example of issuing the partner command with an argument**

```
filer1(takeover)> partner cf status
filer2 has been taken over by filer1.
filer1(takeover)>
```

**Assessing the emulated node remotely**

You can also access the emulated node remotely using a Remote Shell (rsh) connection. You cannot access the emulated node using Secure Shell (ssh) or Telnet.

**Accessing the emulated node remotely using Remote Shell**

You can access the emulated node remotely using a Remote Shell (rsh) connection. You cannot access the emulated node using Secure Shell (ssh) or Telnet.

**Step**

1. Enter the following command:

```
rsh failed_node command
```

*failed\_node* is the name of the failed node.

*command* is the Data ONTAP command you want to run.

**Example of an rsh command**

In the following example, filer2 is the failed node.

```
rsh filer2 df
```

**Emulated node command exceptions**

Almost all the commands that are available to a takeover node are available on the emulated node. Some commands, however, are either unavailable or behave differently in emulated mode.

**Unavailable commands**

The following commands are not available on an emulated node:

- cf disable
- cf enable
- cf forcegiveback

- cf forcetakeover
- cf giveback
- cf takeover
- date
- halt
- ifconfig partner
- ifconfig -partner
- ifconfig mtusize
- license cluster
- rdate
- reboot
- timezone

### Commands with different behaviors

Command	Difference
ifconfig [interface]	<p>Displays the following:</p> <ul style="list-style-type: none"> <li>• Emulated interface mappings based on the failed node's <code>/etc/rc</code> file rather than the takeover node interface mappings</li> </ul> <p><b>Note:</b> MetroCluster nodes use the failed node's <code>/etc/mcrc</code> file if the <code>cf.takeover.use_mcrc_file</code> option is enabled.</p> <ul style="list-style-type: none"> <li>• Emulated interface names rather than the interface names of the takeover node</li> <li>• Only interfaces that have been configured, rather than all interfaces, configured or not, as displayed on the takeover node</li> </ul>
mt	Uses the tape devices on the takeover node because the failed node has no access to its tape devices.
netstat -i	Appends a plus sign (+) to shared interfaces. A shared interface is one that has two IP addresses assigned to it: an IP address for the node in which it physically resides and an IP address for its partner node in the active/active configuration.

Command	Difference
sysconfig	When it displays hardware information, the <code>sysconfig</code> command displays information only about the hardware that is attached to the takeover node. It does not display information about the hardware that is attached only to the failed node. For example, the disk adapter information that the partner <code>sysconfig -r</code> command displays is about the disk adapters on the takeover node.
uptime	Displays how long the failed node has been down and the host name of the takeover node.
aggr status	When it displays hardware information, the <code>aggr status</code> command displays information only about the hardware that is attached to the takeover node. It does not display information about the hardware that is attached only to the failed node. For example, the disk adapter information that the partner <code>aggr status -r</code> command displays is about the disk adapters on the takeover node.

## Performing dumps and restores for a failed node

You can use the emulated node and peripheral devices attached to the takeover node to perform dumps and restores for the failed node.

### Before you begin

Any `dump` commands directed to the failed node's tape drives are executed on the takeover node's tape drives. Therefore, any `dump` commands that you execute using a scheduler, such as the `cron` command, succeed only under the following conditions:

- The device names are the same on both nodes in the active/active configuration.
- The `dump` commands for the takeover node and the emulated node are not scheduled to occur during the same time period; the takeover node and the emulated node cannot access the tape drives simultaneously.

### About this task

Because the peripheral devices for a failed node are inaccessible, you perform dumps and restores for a failed node by using the emulated node (available using the `partner` command on the takeover node), making sure that you use a peripheral device attached to the takeover node.

For more information about performing dumps and restores, see the *Data ONTAP Data Protection Tape Backup and Recovery Guide*.

### Step

1. Issue the `backup` or `restore` command, either in partner mode or as an argument in the `partner` command.

#### Example

Issuing a `restore` command in partner mode:

```
filer1 (takeover)> partner
```

```
filer1/filer2> restore [options [arguments]]
```

```
filer1 (takeover)> partner
```

#### Example

Issuing a `restore` command as an argument in the `partner` command:

```
filer1 (takeover)> partner restore [options [arguments]]
```

## Giveback operations

Giveback can be implemented and configured in a number of different ways. It can also be configured to occur automatically.

### Next topics

[Performing a giveback](#) on page 144

[Configuring giveback](#) on page 147

[Enabling automatic giveback](#) on page 148

## Performing a giveback

You can perform a normal giveback, a giveback in which you terminate processes on the partner node, or a forced giveback.

**Note:** Prior to performing a giveback, you must remove failed drives in the taken-over system.

### Next topics

[Removing failed disks prior to attempting giveback](#) on page 145

[Initiating normal giveback](#) on page 145

[Troubleshooting if giveback fails](#) on page 145

[Forcing giveback](#) on page 146

[If giveback is interrupted](#) on page 147



## Removing failed disks prior to attempting giveback

For taken-over systems that use disks, you must remove the failed disk or disks prior to attempting to implement giveback.

### Step

1. Remove the failed disks, as described in the *Storage Management Guide*.

When all failed disks are removed or replaced, proceed with the giveback operation.

## Initiating normal giveback

You can return control to a taken-over partner with the `cf giveback` command.

On a fabric-attached MetroCluster, before you undertake the giveback operation, you must rejoin the aggregates on the surviving node and the partner node to reestablish the MetroCluster configuration.

### Step

1. Enter the following command on the command line of the takeover node:

```
cf giveback
```

**Note:** If the giveback fails, there might be a process running that prevents giveback. You can wait and repeat the command, or you can use the `initiate giveback` using the `-f` option to terminate the processes that are preventing giveback.

After a giveback, the takeover node's ability to take over its partner automatically is not reenabled until the partner reboots successfully. If the partner fails to reboot, you can enter the `cf takeover` command to initiate a takeover of the partner manually.

## Troubleshooting if giveback fails

If the `cf giveback` command fails, you should check for system processes that are currently running and might prevent giveback, check that the cluster interconnect is operational, and check for any failed disks on systems using disks.

### Steps

1. For systems using disks, check for and remove any failed disks, as described in the *Data ONTAP Storage Management Guide*.
2. Check for the message `cf.giveback.disk.check.fail` on the console. Both nodes should be able to detect the same disks. This message indicates that there is a disk mismatch: for some reason, one node is not seeing all the disks attached to the active/active configuration.

3. Check the cluster interconnect and verify that it is correctly connected and operating.
4. Check whether any of the following processes were taking place on the takeover node at the same time you attempted the giveback:
  - Outstanding CIFS sessions
  - RAID disk additions
  - Volume creation (traditional volume or FlexVol volume)
  - Aggregate creation
  - Disk ownership assignment
  - Disks being added to a volume (vol add)
  - Snapshot copy creation, deletion, or renaming
  - Quota initialization
  - Advanced mode repair operations, such as waffliron
  - Storage system panics
  - Backup dump and restore operations
  - SnapMirror transfers (if the partner is a SnapMirror destination)
  - SnapVault restorations
  - Disk sanitization operations

If any of these processes are taking place, either cancel the process or wait until it is complete, and then try the giveback operation again.

5. If the `cf giveback` operation still does not succeed, use the `cf giveback -f` command to force giveback.

### Related tasks

[Forcing giveback](#) on page 146

## Forcing giveback

Because the takeover node might detect an error condition on the failed node that typically prevents a complete giveback, such as data not being flushed from NVRAM to the failed node's disks, you can force a giveback, if necessary.

Use this procedure to force the takeover node to give back the resources of the failed node even if the takeover node detects an error that typically prevents a complete giveback,

### Step

1. On the takeover node, enter the following command:

```
cf forcegiveback
```

**Attention:** Use `cf forcegiveback` only when you cannot get `cf giveback` to succeed. When you use this command, you risk losing any data committed to NVRAM but not to disk.

If a `cifs terminate` command is running, allow it to finish before forcing a giveback.

## If giveback is interrupted

If the takeover node experiences a failure or a power outage during the giveback process, the giveback process stops and the takeover node returns to takeover mode when the failure is repaired or the power is restored.

## Configuring giveback

You can configure how giveback occurs, setting different Data ONTAP options to improve the speed and timing of giveback.

### Next topics

[Option for shortening giveback time](#) on page 147

[Setting giveback delay time for CIFS clients](#) on page 147

[Option for terminating long-running processes](#) on page 148

[Setting giveback to terminate long-running processes](#) on page 148

## Option for shortening giveback time

You can shorten the client service outage during giveback by using the `cf.giveback.check.partner` option. You should always set this option to `on`.

## Setting giveback delay time for CIFS clients

You can specify the number of minutes to delay an automatic giveback before terminating CIFS clients that have open files.

This option specifies the number of minutes to delay an automatic giveback before terminating CIFS clients that have open files. During the delay, the system periodically sends notices to the affected clients. If you specify 0, CIFS clients are terminated immediately.

This option is used only if automatic giveback is `On`.

### Step

1. Enter the following command:

```
options cf.giveback.auto.cifs.terminate.minutes minutes
```

Valid values for `minutes` are 0 through 999. The default is 5 minutes.

## Option for terminating long-running processes

Describes the `cf.giveback.auto.terminate.bigjobs` option, which, when on, specifies that automatic giveback should immediately terminate long-running operations.

The `cf.giveback.auto.terminate.bigjobs` option, when on, specifies that automatic giveback should immediately terminate long-running operations (dump/restore, vol verify, and so on) when initiating an automatic giveback. When this option is off, the automatic giveback is deferred until the long-running operations are complete. This option is used only if automatic giveback is On.

## Setting giveback to terminate long-running processes

You can set the automatic giveback process to terminate long-running processes that might prevent the giveback.

### Step

1. Enter the following command:

```
options cf.giveback.auto.terminate.bigjobs {on|off}
```

The `on` argument enables this option. The `off` argument disables this option. This option is On by default.

## Enabling automatic giveback

You can enable automatic giveback by using the `cf.giveback.auto.enable` option.

### About this task

Use the automatic giveback feature with care:

- Do not enable automatic giveback in MetroCluster configurations. Before the giveback operation is undertaken, you must rejoin the aggregates on the surviving node and the partner node to reestablish the MetroCluster configuration. If automatic giveback is enabled, this crucial step cannot be performed before the giveback.
- You should leave this option disabled unless your clients are unaffected by failover, or you have processes in place to handle repetitive failovers and givebacks.

### Step

1. Enable the following option to enable automatic giveback: `cf.giveback.auto.enable on`. The `on` value enables automatic giveback. The `off` value disables automatic giveback. This option is `off` by default.

## Downloading and running the HA Configuration Checker utility

You can go on the NOW site and download the Configuration Checker to check for common configuration errors.

### Before you begin

To run the HA Configuration Checker utility, you must have rsh access to both nodes.

### About this task

You can run the utility, `cf-config-check.cgi`, as a command from a UNIX shell, or you can install the Common Gateway Interface (CGI) script on a UNIX Web server and invoke it from a Web browser.

### Steps

1. To download and run the HA Configuration Checker, log in to the NOW site and go to Software Downloads > Tools & Utilities. Click “HA Configuration Checker (`cf-config-check.cgi`).”
2. Follow the directions on the web page for downloading and running the utility.

## Troubleshooting takeover or giveback failures

If takeover or giveback fails for an active/active configuration, you need to check the cluster status and proceed based on messages you receive.

### Steps

1. Check communication between the local and partner nodes by entering the following command and observing the messages:

```
cf status
```

2. Review the messages and take the appropriate action:

<b>If the error message indicates:</b>	<b>Then...</b>
<b>A cluster adapter error</b>	Check the cluster adapter cabling. Make sure that the cabling is correct and properly seated at both ends of the cable.
<b>That the NVRAM adapter is in the wrong slot number</b>	Check the NVRAM slot number. Move it to the correct slot if needed.
<b>A Channel B cabling error</b>	Check the cabling of the Channel B disk shelf loops and reseal and tighten any loose cables.

---

If the error message indicates:	Then...
<b>A networking error</b>	Check for network connectivity.  See the <i>Data ONTAP MultiStore Management Guide</i> for more information.

---

3. If you have not already done so, run the HA Configuration Checker script.
4. Correct any errors or differences displayed in the output.
5. Reboot the active/active configuration and rerun the takeover and giveback tests.
6. If you still do not have takeover enabled, contact technical support.

**Related tasks**

*[Downloading and running the HA Configuration Checker utility](#)* on page 149

## Management of DS14mk2 AT, DS14mk2 FC, or DS14mk4 FC disk shelves in an active/active configuration

---

Describes how to manage Multipath Storage, how to add disk shelves to an active/active configuration or a MetroCluster, and how to upgrade or replace disk shelf hardware in an active/active configuration.

**Note:** If your configuration includes DS4243 disk shelves, refer to the *DS4243 System Connectivity Guide* and the *DS4243 Hardware Service Guide* on the NOW site at <http://now.netapp.com/>.

### Next topics

[Managing disk shelves in Multipath Storage configurations](#) on page 151

[Adding disk shelves to non-Multipath Storage configurations](#) on page 159

[Upgrading or replacing modules in an active/active configuration](#) on page 164

## Managing disk shelves in Multipath Storage configurations

Multipath Storage for Active/active configurations provides redundancy for the path from each controller to every disk shelf in the configuration. An active/active configuration without Multipath Storage has only one path from each controller to every disk, but an active/active configuration with Multipath Storage has two paths from each controller to each disk, regardless of which node owns the disk. Multipath is the preferred configuration.

**Note:** Multipath Storage is not supported with third-party storage connected to V-Series systems but is supported with native disks connected to V-Series systems.

### Next topics

[What Multipath Storage for active/active configurations is](#) on page 151

[How the connection types are used](#) on page 152

[Advantages of Multipath Storage for active/active configurations](#) on page 153

[Requirements for Multipath Storage](#) on page 153

[Determining whether your AT-FCX modules support Multipath Storage](#) on page 155

[Cabling for Multipath Storage](#) on page 156

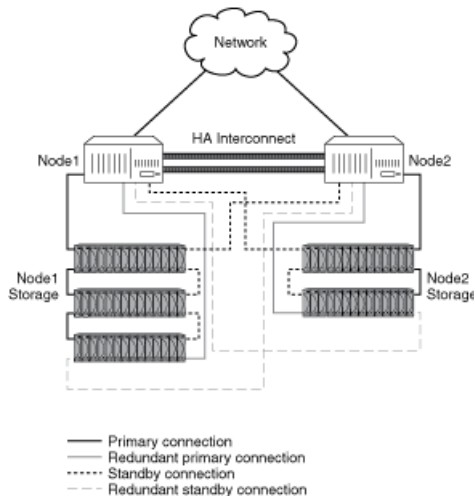
[Adding storage to a Multipath Storage loop](#) on page 157

## What Multipath Storage for active/active configurations is

Multipath Storage for active/active configurations provides redundancy for the path from each controller to every disk shelf in the configuration. It is the preferred method for cabling a storage system. An

active/active configuration without Multipath Storage has only one path from each controller to every disk, but an active/active configuration with Multipath Storage has two paths from each controller to each disk, regardless of which node owns the disk.

The following diagram shows the connections between the controllers and the disk shelves for an example active/active configuration using Multipath Storage. The redundant primary connections and the redundant standby connections are the additional connections required for Multipath Storage for active/active configurations.



**Figure 27: Multipath Storage in an active/active configuration**

## How the connection types are used

Outlines the connection types used for Multipath Storage in active/active configurations.

The following table outlines the connection types used for Multipath Storage for active/active configurations, and how the connections are used.

Connection type	How the connection is used
Primary connection	For normal operation, used to serve data (load-balanced with redundant primary connection).
Redundant primary connection	For normal operation, used to serve data (load-balanced with primary connection).
Standby connection	For normal operation, used for heartbeat information only. After a takeover, assumes role of primary connection.
Redundant standby connection	Not used for normal operation. After a takeover, assumes role of redundant primary connection. If the standby connection is unavailable at takeover time, assumes role of primary connection.



## Advantages of Multipath Storage for active/active configurations

Multipath connections reduce single-points-of-failure.

By providing two paths from each controller to every disk shelf, Multipath Storage provides the following advantages:

- The loss of a disk shelf module, connection, or host bus adapter (HBA) does not require a failover. The same storage system can continue to access the data using the redundant path.
- The loss of a single disk shelf module, connection, or HBA does not prevent a successful failover. The takeover node can access its partner's disks using the redundant path.
- You can replace modules without having to initiate a failover.

**Note:** While Multipath adds value to a stretch MetroCluster environment, it is not necessary in a fabric MetroCluster configuration since multiple paths already exist.

### Related concepts

[Understanding redundant pathing in active/active configurations](#) on page 167

## Requirements for Multipath Storage

Multipath Storage for active/active configurations has certain requirements.

### System requirements

See the *Data ONTAP® Release Notes* for a list of systems that support Multipath storage.

### Active/active configuration-type requirements

Multipath Storage is available for the following types of active/active configurations:

- Standard active/active configurations
- Mirrored active/active configurations
- Stretch MetroClusters

Your active/active configuration must be installed and fully operational. Your configuration testing should include successful failover and giveback.

**Note:** Fabric-attached MetroClusters have redundant disk paths by default; no special configuration is necessary.

### Disk shelf requirements

Multipath Storage for active/active configurations is available for only the following combinations of disk shelves and modules:

- DS14mk2 FC or DS14mk4 FC disk shelves with ESH2 or ESH4 modules

- DS14mk2 AT disk shelves with RoHS-compliant AT-FCX modules

**Note:** Only AT-FCX modules shipped in December 2005 or later support Multipath Storage for Active/active configurations. Check the version of the AT-FCX module to ensure support.

### Best practice recommendation

If any loop in your active/active configuration is cabled for Multipath Storage, every loop should be cabled for Multipath Storage. This is the recommended best practice.

**Note:** If you have a mixed configuration in which some loops are cabled for Multipath Storage and some are not, the system displays a configuration error message when you boot the system or when a disk on a loop that is cabled for Multipath becomes single-pathed.

### Software-based disk ownership requirements

Both nodes must be using software-based disk ownership.

To convert an active/active configuration to use software-based disk ownership, you must boot both nodes into Maintenance mode at the same time (during scheduled downtime).

**Note:** Plan to convert to software-based disk ownership before adding any cabling for Multipath Storage. After you add the cabling for Multipath Storage, you must manually assign all disks.

For more information about software-based disk ownership, see the chapter about disks in the *Data ONTAP Storage Management Guide*.

### Fibre Channel port requirements

Each node must have enough onboard Fibre Channel ports or HBAs to accommodate the extra cables required for Multipath Storage. Without Multipath Storage, you need one Fibre Channel port for each controller for each loop in the configuration. With Multipath Storage, you need two Fibre Channel ports for each loop.

If you are scheduling downtime to convert to software-based disk ownership, you should add the HBAs then. Otherwise, you can use the nondisruptive upgrade method to add the HBA; this method does not require downtime.

**Note:** See the *System Configuration Guide* at [http://now.netapp.com/NOW/knowledge/docs/hardware/hardware\\_index.shtml](http://now.netapp.com/NOW/knowledge/docs/hardware/hardware_index.shtml) for information about which slots to use for the HBAs and in what order to use them.

### Boot environment variable requirement for V-Series systems

To use Multipath Storage on a V-Series system, you must configure the `fc-nonarray-adapter-list` environment variable for each new loop before you connect and configure the disk shelf for Multipath Storage. See the *V-Series Implementation Guide for Native Disk Shelves*.

## Determining whether your AT-FCX modules support Multipath Storage

To use an AT-FCX module with Multipath Storage for active/active configurations, the module must be the right version.

### About this task

Modules shipped prior to December 2005 do not support Multipath Storage for active/active configurations. If you are unsure whether your module is the correct version, use the following procedure.

### Steps

1. Determine the disk address of the target SES device for the AT-FCX module by entering the following command:

```
fcadmin device_map
```

#### Example

```
node1> fcadmin device_map
Loop Map for channel 3b:
Translated Map: Port Count 17
7 46 47 48 49 50 51 52 53 54 55 56 57 58 59 60 61
Shelf mapping:
Shelf 3: 61 60 59 58 57 56 55 54 53 52 51 50 49 48
Target SES devices on this loop:
Shelf 3: 46 47
```

The device map is displayed, including the SES target device IDs.

2. Set the privilege level to advanced by entering the following command:

```
priv set advanced
```

3. Enter the following command:

```
fcadmin bridge cmd_cemi -d target_SES_device -c "cpld"
```

#### Example

```
node1*> fcadmin bridge cmd_cemi -d 3b.46 -c "cpld"
21h45m11s SelID:3(255).MA IMS601550001296 8b.03 pfuID=32
? RegisterCount: 20
Drive 00:80 Drive 01:80 Drive 02:80 Drive 03:80
Drive 04:80 Drive 05:80 Drive 06:00 Drive 07:80
Drive 08:80 Drive 09:80 Drive 10:80 Drive 11:80
Drive 12:80 Drive 13:80 Drive 14:00 Drive 15:80
HOST:20 HwId:0f SFP Reg2:52 SFP Reg1:22
Revision:24 MISC:00 0 8067:2c 1 8067:2e
Fan control:96 Controller:98 ThumbCtlrId:57 Fan status:00
PSU0:0e PSU1:0e Status:08 Control:2d
```

Several columns of information about the disk shelf are displayed.

4. Find the Revision value in the output of the `fcadmin bridge` command. If this value is 24 or higher, your AT-FCX module supports Multipath Storage for active/active configurations.
5. Return the privilege level to administrative by entering the following command:

```
priv set
```

## Cabling for Multipath Storage

To cable your active/active configuration for Multipath Storage, you need to determine your disk paths, make cabling changes, and then confirm that multiple paths are in place to each disk.

### Before you begin

Make sure that your active/active configuration meets the following prerequisites:

- Both nodes must be using software-based disk ownership.  
To convert an active/active configuration to use software-based disk ownership, you must boot both nodes into Maintenance mode at the same time (during scheduled downtime).

**Note:** Plan to convert to software-based disk ownership before adding any cabling for Multipath Storage. After you add the cabling for Multipath Storage, you must manually assign all disks.

For more information about software-based disk ownership, see the chapter about disks in the *Data ONTAP Storage Management Guide*.

- To use Multipath Storage on a V-Series system, you must configure the `fc-non-array-adapter-list` environment variable for each new loop before you connect and configure the disk shelf for Multipath Storage.
- Each node must have enough onboard Fibre Channel ports or HBAs to accommodate the extra cables required for Multipath Storage.

If you are scheduling downtime to convert to software-based disk ownership, you should add the HBAs then. Otherwise, you can use the nondisruptive upgrade method to add the HBA; this method does not require downtime.

**Note:** See the *System Configuration Guide* on the NOW site for information about which slots to use for the HBAs and in what order to use them.

- All disk shelf modules must be ESH2, ESH4, or AT-FCX.
- Your active/active configuration must be installed and fully operational. Your configuration testing should include successful failover and giveback.

**Note:** For detailed instructions about connecting disk shelves, see the hardware documentation for your disk shelf. For detailed instructions about adding HBAs to your controller, see the hardware documentation for your system model.

### Steps

1. Make sure that your active/active configuration meets the requirements.

2. (Optional) To display the current (nonredundant) disk paths, enter the following command:

```
storage show disk -p
```

3. Pick a loop on one node (Node 1) for Channel A (the cable is connected to the A port on the disk shelves), and trace the cables from the controller to the last disk shelf in the loop.

**Note:** The last disk shelf has no cable connected to the Channel A Output port.

4. Using the correct cable type for a disk shelf-to-controller connection, connect the Channel A Output port to a Fibre Channel port on the partner node (Node 2).

**Note:** When possible, do not connect the same HBA to both the primary and redundant path of the same loop. For example, if an HBA is connected to Channel B for a loop, do not use another port on that HBA for the redundant connection for Channel A of that same loop. Otherwise, the failure of the HBA could prevent the controller from accessing that loop.

Adjacent pairs of on-board ports share hardware; consider them to be the same as a single HBA. For example, do not use port 0a and port 0b for the primary and redundant paths of the same loop.

5. From the same disk shelf, using the correct cable type for a shelf-to-controller connection, connect the Channel B Output port to a Fibre Channel port on the original controller (Node 1).
6. Repeat Step 3 through Step 5 for every loop connected to Node 1.
7. Repeat Step 3 through Step 6 for Node 2.

There should be a cable in every Input and Output port of all the disk shelves.

8. Confirm that there are two paths to every disk by entering the following command:

```
storage show disk -p
```

The `sysconfig` command can also be used to confirm multiple paths to each disk.

There should be two paths listed for every disk.

### Related concepts

[Nondisruptive hardware changes](#) on page 185

[Understanding redundant pathing in active/active configurations](#) on page 167

### Related tasks

[Determining whether your AT-FCX modules support Multipath Storage](#) on page 155

[Upgrading or replacing modules in an active/active configuration](#) on page 164

[Determining path status for your active/active configuration](#) on page 167

## Adding storage to a Multipath Storage loop

To add storage to an active/active configuration configured for Multipath Storage, you need to add the new disk shelf to the end of a loop, ensuring that it is connected to the previous disk shelf and to the controller.

## Steps

1. Confirm that there are two paths to every disk by entering the following command:

```
storage show disk -p
```

**Note:** If there are not two paths listed for every disk, this procedure could result in a data service outage. Before proceeding, address any issues so that all paths are redundant. If you do not have redundant paths to every disk, you can use the nondisruptive upgrade method (failover) to add your storage.

2. Install the new disk shelf in your cabinet or equipment rack, as described in the DiskShelf14 or DiskShelf14mk2/mk4 *Hardware Service Guide*.
3. Determine whether disk shelf counting is enabled by entering the following command:

```
options cf.takeover.on_disk_shelf_miscompare
```

4. If the disk shelf counting option is set to On, turn it off by entering the following command:

```
options cf.takeover.on_disk_shelf_miscompare off
```

5. Find the last disk shelf in the loop for which you want to add the new disk shelf.

**Note:** The Channel A Output port of the last disk shelf in the loop is connected back to one of the controllers.

**Note:** In Step 6 you disconnect the cable from the disk shelf. When you do this the system displays messages about adapter resets and eventually indicates that the loop is down. These messages are normal within the context of this procedure. However, to avoid them, you can optionally disable the adapter prior to disconnecting the disk shelf.

If you choose to, disable the adapter attached to the Channel A Output port of the last disk shelf by entering the following command:

```
fcadmin config -d <adapter>
```

*<adapter>* identifies the adapter by name. For example: 0a.

6. Disconnect the SFP and cable coming from the Channel A Output port of the last disk shelf.

**Note:** Leave the other ends of the cable connected to the controller.

7. Using the correct cable for a shelf-to-shelf connection, connect the Channel A Output port of the last disk shelf to the Channel A Input port of the new disk shelf.
8. Connect the cable and SFP you removed in Step 6 to the Channel A Output port of the new disk shelf.
9. If you disabled the adapter in Step 5, reenable the adapter by entering the following command:

```
fcadmin config -e <adapter>
```

10. Repeat Step 6 through Step 9 for Channel B.

**Note:** The Channel B Output port is connected to the other controller.

11. Confirm that there are two paths to every disk by entering the following command:

```
storage show disk -p
```

There should be two paths listed for every disk.

12. If disk shelf counting was Off, reenable it by entering the following command:

```
options cf.takeover.on_disk_shelf_miscompare on
```

#### Related tasks

[Determining path status for your active/active configuration](#) on page 167

## Adding disk shelves to non-Multipath Storage configurations

Describes how to add disk shelves to Active/active configurations in systems that are not using Multipath Storage.

#### Next topics

[Overview of adding storage to non-multipath configurations](#) on page 159

[Adding storage to an existing non-multipath loop](#) on page 161

[Adding a new non-multipath loop](#) on page 163

[Adding storage to fabric-attached MetroClusters](#) on page 164

## Overview of adding storage to non-multipath configurations

You can add a disk shelf to an active/active configuration without powering down your active/active configuration, as long as your system meets the requirements. These procedures are for systems that are not using Multipath Storage.

#### Next topics

[Support for adding disk shelves without powering down](#) on page 159

[Restrictions for addition of disk shelves to an active/active configuration](#) on page 160

## Support for adding disk shelves without powering down

You can add a disk shelf to an active/active configuration without powering down your active/active configuration (sometimes referred to as *hot-adding* the disk shelf), as long as your system meets the active/active configuration requirements.

Hot-adding disk shelves enables you to add (but not swap) disk shelves without a service interruption. However, you cannot add more disk shelves to your active/active configuration than it can support, either for the entire configuration or for either node. See the *System Configuration Guide* at the NOW site for maximum storage capacity values.

**Attention:** 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.

**Related tasks**

[Adding storage to a Multipath Storage loop](#) on page 157

**Restrictions for addition of disk shelves to an active/active configuration**

The following table lists some of the restrictions for adding disk shelves to your active/active configuration.

See the *DiskShelf14 and DiskShelf14mk2 FC Hardware Guide* for a detailed discussion of restrictions.

See the *System Configuration Guide* for the maximum number of FC-AL adapters your storage system supports and for additional restrictions.

If your active/active configuration 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.
FC7, FC8, or FC9 disk shelves	Adding DS14/DS14mk2 FC disk shelves with LRCs	A DS14/DS14mk2/DS14mk4 FC must be the last disk shelf in the loop.
DS14/DS14mk2/DS14mk4 FC disk shelves with ESH, ESH2 or ESH4 modules	Hot-adding a disk shelf	<ul style="list-style-type: none"> <li>• Disks should be pre-zeroed and used as spares.</li> <li>• You can only hot-add a DS14/DS14mk2/DS14mk4 FC disk shelf to an existing DS14/DS14mk2/DS14mk4 FC disk shelf.</li> <li>• A DS14/DS14mk2/DS14mk4 FC disk shelf must be the last disk shelf in the loop.</li> <li>• DS14/DS14mk2/DS14mk4 FC disk shelves that have ESH, ESH2 or ESH4 modules can only be added to disk shelf loops that have ESH, ESH2 or ESH4 modules.</li> <li>• If the new disk shelf is a DS14/DS14mk2/DS14mk4 FC disk shelf, the loop speed switch must be set for the appropriate installation.</li> </ul>



If your active/active configuration has...	And you are...	Then...
DS14mk2 AT disk shelves with AT-FCX modules	Hot-adding a disk shelf	<ul style="list-style-type: none"> <li>• Disks should be pre-zeroed and used as spares.</li> <li>• You can only hot-add a DS14mk2 AT disk shelf to an existing a DS14mk2 AT disk shelf.</li> <li>• a DS14mk2 AT disk shelves that have AT-FCX modules can only be added to disk shelf loops that have AT-FCX modules.</li> </ul>

## Adding storage to an existing non-multipath loop

You can add storage to an existing loop in an active/active configuration without disrupting service.

### Before you begin

Before adding a new disk shelf to an existing loop, make sure that the addition will not exceed the maximum number of disk shelves for that loop. Also determine an ID for the new disk shelf that is unique for the loop you are adding the new disk shelf to.

The maximum number of disk shelves depends on the disk shelf type:

- For DS14 disk shelves, the maximum number of disk shelves in a loop is six, with disk shelf IDs of 1-6.
- For FC7, FC8 and FC9 disk shelves, the maximum number of shelves in a loop is seven, with disk shelf IDs of 0-6.
- For fabric-attached MetroClusters, the maximum number of disk shelves in a loop is 2.

**Note:** If a disk shelf is installed on a V-Series system, MetroCluster is not supported on that system.

**Note:** If you want to add more than one disk shelf, add only one new disk shelf at a time.

This procedure does not apply to adding storage to an active/active configuration using Multipath Storage.

### Steps

1. Install the new disk shelf in your cabinet or equipment rack, as described in the *DiskShelf14 or DiskShelf14mk2/mk4 Hardware Service Guide*.
2. Determine whether disk shelf counting is enabled by entering the following command:

```
options cf.takeover.on_disk_shelf_miscompare
```

3. If the disk shelf counting option is set to On, turn it off by entering the following command:

```
options cf.takeover.on_disk_shelf_miscompare off
```

4. Verify that the new disk shelf ID is unique for the loop you are adding the new disk shelf to by entering the following command:

```
fcstat device_map
```

Disk shelf IDs for all Fibre Channel devices are listed. Make sure the disk shelf ID for the new disk shelf is not already in use for this loop. If it is, change it to the next available ID.

5. Depending on the type of disk shelves you are using, take the appropriate action:

If...	Then for loops with...
<b>Disk shelves in the loop to which you are adding the new disk shelf have LRCs</b>	<ul style="list-style-type: none"> <li>• DS14 only: Go to Step 6.</li> <li>• A mix of DS14 and DS14mk2 FC: Set the loop speed to 1 Gb.</li> </ul> <p><b>Attention:</b> An incorrectly set loop speed results in an open loop condition.</p>
<b>Disk shelves in the loop to which you are adding the new disk shelf have either ESH, ESH2, or ESH4 modules exclusively, or AT-FCX modules exclusively</b>	<ul style="list-style-type: none"> <li>• DS14 only or a mix of DS14 and DS14mk2 OR DS14MK4 FC: Set the loop speed to 1 Gb.</li> <li>• DS14mk2 or DS14mk4 FC only that are running at 1 Gb: Set the loop speed to 1 Gb.</li> <li>• DS14mk2 or DS14mk4 FC only that are running at 2 Gb: Set the loop speed to 2 Gb.</li> <li>• AT-FCX modules: The loop speed is set automatically.</li> </ul> <p><b>Attention:</b> An incorrectly set loop speed causes the storage system to panic.</p>

6. Apply power to the new disk shelf and turn it on. Then wait 60 seconds for the disk shelf to fully power up and all electronics to come online.
7. Set both terminate switches (except for ESH2 or ESH4 and AT-FCX, which don't have them) on the new disk shelf to On.
8. Attach the disk shelf-to-disk shelf cable to the Channel A Input port of the new disk shelf.
9. Attach the other end of the cable to the Channel A Output port of the last existing disk shelf in the loop.
10. Set the Channel A terminate switch (except for ESH2 or ESH4 and AT-FCX) on the previous disk shelf to Off.
11. Attach the disk shelf-to-disk shelf cable to the Channel B Input port of the new disk shelf.
12. Attach the other end of the cable to the Channel B Output port of the last existing disk shelf in the loop.

13. Set the Channel B terminate switch (except for ESH2 or ESH4 and AT-FCX) on the previous disk shelf to Off.
14. If disk shelf counting was Off, reenable it by entering the following command:

```
options cf.takeover.on_disk_shelf_miscompare on
```

## Adding a new non-multipath loop

If you have an open Fibre Channel port, either an onboard port or a port on an HBA, you can create a new loop by hot-adding a disk shelf on its own loop. This procedure is for a loop that is not using Multipath Storage.

### About this task

If you have an open Fibre Channel port, either an onboard port or a port on an HBA, you can create a new loop by hot-adding a disk shelf on its own loop. To add a new loop to an active/active configuration, complete the following steps.

**Note:** If want to add more than one disk shelf, use this procedure to add one disk shelf. Then add each new disk shelf one at a time. This procedure is not for adding a loop to a system using Multipath Storage.

### Steps

1. Install the new disk shelf in your cabinet or equipment rack, as described in the *DiskShelf14 or DiskShelf14mk2/mk4 Hardware Service Guide*.
2. Determine whether disk shelf counting is enabled by entering the following command:  

```
options cf.takeover.on_disk_shelf_miscompare
```
3. If the disk shelf counting option is set to On, turn it off by entering the following command:  

```
options cf.takeover.on_disk_shelf_miscompare off
```
4. If the new disk shelf has a loop speed setting, set the loop speed to the supported speed of the lowest speed component in the loop.  
  
To support 4 Gbps, all components in the configuration must support 4 Gbps.
5. Apply power to the new disk shelf and turn it on, and then wait 60 seconds for the disk shelf to fully power up and all electronics to come online.
6. Set both terminate switches (except for ESH2 or ESH4 and AT-FCX, which don't have them) on the new disk shelf to On.
7. Connect the Channel A Input port of the new disk shelf to the Fibre Channel port on one of the controllers.

**Note:** If your active/active configuration is using hardware-based ownership, the new disk shelf will be owned by this controller.

8. Connect the Channel B Input port of the new disk shelf to the Fibre Channel port on the other controller.
9. If disk shelf counting was on, reenable it by entering the following command:

```
options cf.takeover.on_disk_shelf_miscompare on
```

## Adding storage to fabric-attached MetroClusters

Fabric-attached MetroClusters can have multiple loops attached to each node through the Brocade switch. When you need to add storage to a MetroCluster, as long as doing so does not exceed the MetroCluster disk limit, you can use one of the following methods:

- Add another disk shelf to an existing loop.

**Note:** You can add a disk shelf as a hot-add operation (you do not need to power down the MetroCluster).

You cannot attach more than two shelves to a loop in a fabric-attached MetroCluster.

You can add a new loop as a hot-add operation (you do not need to power down the MetroCluster).

- Add a new loop to the MetroCluster.

Adding a new loop requires an open port in the correct switch quadrant. 8-port switches support up to two mirrored loops on each node; 16-port switches support up to four mirrored loops on each node.

When you add a new loop to an existing fabric-attached MetroCluster, always connect the loop to a switch port in the same switch quadrant as the loops that are already connected—for both the local node and the remote node.

The newly connected switch port must be configured correctly for a MetroCluster configuration.

For information about how you configure the Brocade switch for MetroClusters, and how it is subdivided into quadrants, see the *Brocade Switch Configuration Guide* for your switch. You can find this document on the MetroCluster Switch Description Page at the NOW site.

### Related concepts

[Adding disk shelves to non-Multipath Storage configurations](#) on page 159

[Setup requirements and restrictions for fabric-attached MetroClusters](#) on page 36

## Upgrading or replacing modules in an active/active configuration

In an active/active configuration with redundant pathing, you can upgrade or replace disk shelf modules without interrupting access to storage.

## About this task

These procedures are for DS14mk2 AT, DS14mk2 FC, or DS14mk4 FC disk shelves. If your configuration includes DS4243 disk shelves, refer to the *DS4243 Hardware Service Guide* on the NOW site at <http://now.netapp.com/>.

## Next topics

[About the disk shelf modules](#) on page 165

[Restrictions for changing module types](#) on page 165

[Best practices for changing module types](#) on page 166

[Testing the modules](#) on page 166

[Understanding redundant pathing in active/active configurations](#) on page 167

[Determining path status for your active/active configuration](#) on page 167

[Upgrading an LRC module to an ESH or ESH2 module](#) on page 169

[Hot-swapping a module](#) on page 171

## About the disk shelf modules

A disk shelf module (LRC, ESH, ESH2, ESH4, or AT-FCX) in a DS14, DS14mk2, DS14mk4 FC or DS14mk2 AT includes a SCSI-3 Enclosure Services Processor that maintains the integrity of the loop when disks are swapped and provides signal retiming for enhanced loop stability. When upgrading or replacing a module, you must be sure to cable the modules correctly.

The a DS14, DS14mk2, DS14mk4 FC or DS14mk2 AT disk shelves support the LRC, ESH, ESH2, ESH4, or AT-FCX modules.

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/DS14mk4 FC shelf are the reverse of module A.

## Restrictions for changing module types

If you plan to change the type of any module in your active/active configuration, make sure that you understand the restrictions.

- You cannot mix LRC and ESH modules in the same loop. Doing so results in loop failure.
- You cannot mix LRC, ESH, ESH2 or ESH4 modules in the same loop with AT-FCX modules.
- You cannot mix ESH and ESH4 modules in the same loop.
- To replace an ESH with an ESH2, you must do the following tasks:
  - Upgrade to Data ONTAP 6.4.4 or later, if necessary.
  - Replace all ESH modules in the same disk shelf with ESH2 modules.

## Best practices for changing module types

If you plan to change the type of any module in your active/active configuration, make sure that you review the best practice guidelines.

- Whenever you remove a module from an active/active configuration, you need to know whether the path you will disrupt is redundant. If it is, you can remove the module without interfering with the storage system's ability to serve data. However, if that module provides the only path to any disk in your active/active configuration, you must take action to ensure that you do not incur system downtime.
- When you replace a module, make sure that the replacement module's termination switch is in the same position as the module it is replacing.

**Note:** ESH2 and ESH4 modules are self-terminating; this guideline does not apply to ESH2 and ESH4 modules.

- If you replace a module with a different type of module, make sure that you also change the cables, if necessary.  
For more information about supported cable types, see the hardware documentation for your disk shelf.
- Always wait 30 seconds after inserting any module before reattaching any cables in that loop.
- ESH2 and ESH4 modules should not be on the same disk shelf loop.

### Related concepts

[Understanding redundant pathing in active/active configurations](#) on page 167

## Testing the modules

You should test your disk shelf modules after replacing or upgrading them, to ensure that they are configured correctly and operating.

### Steps

1. Verify that all disk shelves are functioning properly by entering the following command:

```
environ shelf
```

2. Verify that there are no missing disks by entering the following command:

```
aggr status -r
```

Local disks displayed on the local node should be displayed as partner disks on the partner node, and vice-versa.

3. Verify that you can create and retrieve files on both nodes for each licensed protocol.

## Understanding redundant pathing in active/active configurations

Some active/active configurations have two paths from each controller to each of their disk shelves or array LUNs; this configuration is called a redundant-path or multipath configuration. Active/active configurations using Multipath Storage are redundant-path configurations. Fabric-attached MetroClusters are also redundant-path configurations.

## Determining path status for your active/active configuration

You can determine whether any module in your system provides the only path to any disk or array LUN by using the `storage show disk -p` command at your system console.

### About this task

If you want to remove a module from your active/active configuration, you need to know whether the path you will disrupt is redundant. If it is, you can remove the module without interfering with the storage system's ability to serve data. On the other hand, if that module provides the only path to any of the disks or array LUNs in your active/active configuration, you must take action to ensure that you do not incur system downtime.

### Step

1. Use the `storage show disk -p` command at your system console.

This command displays the following information for every disk or array LUN in the active/active configuration:

- Primary port
- Secondary port
- Disk shelf
- Bay

### Examples for configurations with and without redundant paths

The following example shows what the `storage show disk -p` command output might look like for a redundant-path active/active configuration consisting of FAS systems:

PRIMARY	PORT	SECONDARY	PORT	SHELF	BAY
0c.112	A	0b.112	B	7	0
0b.113	B	0c.113	A	7	1
0b.114	B	0c.114	A	7	2
0c.115	A	0b.115	B	7	3
0c.116	A	0b.116	B	7	4
0c.117	A	0b.117	B	7	5
0b.118	B	0c.118	A	7	6
0b.119	B	0c.119	A	7	7

0b.120	B	0c.120	A	7	8
0c.121	A	0b.121	B	7	9
0c.122	A	0b.122	B	7	10
0b.123	B	0c.123	A	7	11

Notice that every disk (for example, 0c.112/0b.112) has two ports active: one for A and one for B. The presence of the redundant path means that you do not need to fail over one system before removing modules from the system.

**Attention:** Make sure that every disk or array LUN has two paths. Even in an active/active configuration configured for redundant paths, a hardware or configuration problem can cause one or more disks to have only one path. If any disk or array LUN in your active/active configuration has only one path, you must treat that loop as if it were in a single-path active/active configuration when removing modules.

The following example shows what the `storage show disk -p` command output might look like for an active/active configuration consisting of FAS systems that does not use redundant paths:

```
filer1> storage show disk -p
PRIMARY PORT SECONDARY PORT SHELF BAY
-----
5b.16      B                1      0
5b.17      B                1      1
5b.18      B                1      2
5b.19      B                1      3
5b.20      B                1      4
5b.21      B                1      5
5b.22      B                1      6
5b.23      B                1      7
5b.24      B                1      8
5b.25      B                1      9
5b.26      B                1     10
5b.27      B                1     11
5b.28      B                1     12
5b.29      B                1     13
5b.32      B                2      0
5b.33      B                2      1
5b.34      B                2      2
5b.35      B                2      3
5b.36      B                2      4
5b.37      B                2      5
5b.38      B                2      6
5b.39      B                2      7
5b.40      B                2      8
5b.41      B                2      9
5b.42      B                2     10
5b.43      B                2     11
5b.44      B                2     12
5b.45      B                2     13
```



For this active/active configuration, there is only one path to each disk. This means that you cannot remove a module from the configuration, thereby disabling that path, without first performing a takeover.

## Upgrading an LRC module to an ESH or ESH2 module

You can upgrade an LRC module to a newer ESH-series module without disrupting data availability.

### Before you begin

The following table describes which types of hot-upgrade and hot-downgrade each type of disk shelf supports.

Disk shelf loop speed	Hot-upgrade	Hot-downgrade
DS14 (only available at 1-Gb LRC to ESH or ESH2ESH or ESH2 hot-upgrade is supported. loop speed)	LRC to ESH or ESH2ESH 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 ESH2ESH, ESH2 or ESH4 hot-upgrade is supported.	ESH or ESH2 to LRC hot-downgrade is supported.
DS14mk2 FC set at 2-Gb loop speed	<ul style="list-style-type: none"> <li>• Not applicable</li> <li>• ESH to ESH2 is supported.</li> </ul>	<ul style="list-style-type: none"> <li>• ESH or ESH2 to LRC hot-downgrade is not supported.</li> <li>• ESH2 to ESH hot-downgrade is supported.</li> </ul>
DS14mk4 FC set at 4-Gb loop speed	Not applicable	Not applicable

### About this task

If your active/active configuration has disk shelves that have LRC modules and you add a DS14/DS14mk2/DS14mk4 FC disk shelf that has ESH, ESH2 or ESH4 modules, you need to replace the LRC modules with ESH, ESH2 or ESH4 modules in the existing DS14 disk shelves.

**Note:** Do not mix LRC and ESH modules in the same loop; loop failure results.

Replacing 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 with ESH2 modules.

For more information about upgrading disk shelves, see the DiskShelf14, DiskShelf14mk2 FC, DiskShelf14mk4 FC or DiskShelf14mk2 AT Hardware Guide.

## Steps

1. Determine which loop you will upgrade first, and determine whether any disks in the active/active configuration are single-pathed through that loop.
2. If any disks use this loop for their only path to a controller, complete the following steps:
  - a. On one node, called Node A, enter the following command:

```
cf takeover
```

- b. Wait for takeover to be complete and make sure that the partner node, or Node B, reboots and is waiting for giveback.

You can now replace all modules on all loops attached to Node B.

3. Note the terminate switch position of the LRCs you are replacing in this loop. If the new modules have terminate switches, they should be set to the same setting as the LRCs they are replacing.

**Note:** The ESH2 is self-terminating and does not have a terminate switch. The ESH2 and ESH4 are self-terminating and do not have a terminate switch.

4. While grounded, unplug the cabling from the dormant loop and note the cable locations.
5. Remove the LRC modules from the disk shelves that do not have a cable attached.
6. Insert all of the new modules.

**Attention:** After inserting all modules, wait 30 seconds before proceeding to the next step.

7. Recable the disk shelves to their original locations.
8. Check the operation of the new modules by entering the following command from the console of the node that is still running:

```
environ shelf
```

The node reports the status of the modified shelves as good.

9. If there is more than one loop attached to the node that has been taken over, repeat Step 3 through Step 8 for all other loops on that node.
10. If you performed a takeover previously, complete the following steps:
  - a. At the console of the takeover node, return control of Node B's disk shelves by entering the following command:

```
cf giveback
```

- b. Wait for the giveback to complete before proceeding to the next step.

11. Check the operation of the new modules by entering the following command on the console of both nodes:

```
environ shelf
```

Each node reports the status of the modified shelves as good.

12. Check the next loop you will upgrade to determine whether any disks are single-pathed through that loop.
13. If any disks use this loop for their only path to a controller, complete the following steps:
  - a. On Node B, enter the following command:  

```
cf takeover
```
  - b. Wait for takeover to be complete and make sure that the partner node, or Node A, reboots and is waiting for giveback.

You can now replace all modules on all loops attached to Node A.

14. Repeat Step 3 through Step 11, with Node B as the takeover node.
15. Test the upgraded modules.
16. Test the configuration.

### Related tasks

[Hot-swapping a module](#) on page 171

[Determining path status for your active/active configuration](#) on page 167

## Hot-swapping a module

You can hot-swap a faulty disk shelf module, removing the faulty module and replacing it without disrupting data availability.

### About this task

When you hot-swap a disk shelf module, you must ensure that you never disable the only path to a disk, which results in a system outage.

**Attention:** If there is newer firmware in the `/etc/shelf_fw` directory than that on the replacement module, the system automatically runs a firmware update. On non-Multipath HA AT-FCX installations, Multipath HA configurations running versions of Data ONTAP prior to 7.3.1, and non-RoHS modules, this firmware update causes a service interruption.

### Steps

1. Verify that your storage system meets the minimum software requirements to support the disk shelf modules that you are hot-swapping. See the *DiskShelf14*, *DiskShelf14mk2 FC*, or *DiskShelf14mk2 AT Hardware Guide* for more information.
2. Determine which loop contains the module you are removing, and determine whether any disks are single-pathed through that loop.
3. If any disks use this loop for their only path to a controller, complete the following steps:
  - a. Follow the cables from the module you want to replace back to one of the nodes, called Node A.
  - b. At the Node B console, enter the following command:

**cf takeover**

- c. Wait for takeover to be complete and make sure that the partner node, or Node A, reboots and is waiting for giveback.

Any module in the loop that is attached to Node A can now be replaced.

4. Note whether the module you are replacing has a terminate switch. If it does, set the terminate switch of the new module to the same setting.

**Note:** The ESH2 and ESH4 are self-terminating and do not have a terminate switch.

5. Put on the antistatic wrist strap and grounding leash.
6. Disconnect the module that you are removing from the Fibre Channel cabling.
7. 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.
8. Slide the replacement module into the slot at the rear of the disk shelf and push the levers of the CAM mechanism into place.

**Attention:** Do not use excessive force when sliding the module into the disk shelf; you might damage the connector.

Wait 30 seconds after inserting the module before proceeding to the next step.

9. Recable the disk shelf to its original location.
10. Check the operation of the new module by entering the following command from the console of the node that is still running:

```
environ shelf
```

The node reports the status of the modified disk shelves.

11. If you performed a takeover previously, complete the following steps:
  - a. At the console of the takeover node, return control of Node B's disk shelves by entering the following command:

```
cf giveback
```

- b. Wait for the giveback to be completed before proceeding to the next step.
12. Test the replacement module.
13. Test the configuration.

**Related concepts**

*Best practices for changing module types* on page 166

**Related tasks**

*Determining path status for your active/active configuration* on page 167

# Disaster recovery using MetroCluster

---

In situations such as prolonged power outages or natural disasters, you can use the optional MetroCluster feature of Data ONTAP to provide a quick failover to another site that contains a nearly real time copy of the data at the disaster site.

**Note:** If you are a V-Series system customer, see the *V-Series MetroCluster Guide* for information about configuring and operating a V-Series system in a MetroCluster configuration.

## Next topics

[Conditions that constitute a disaster](#) on page 173

[Recovering from a disaster](#) on page 175

## Conditions that constitute a disaster

The disaster recovery procedure is an extreme measure that you should use only if the failure disrupts all communication from one MetroCluster site to the other for a prolonged period of time.

The following are examples of disasters that could cause such a failure:

- Fire
- Earthquake
- Prolonged power outages at a site
- Prolonged loss of connectivity from clients to the storage systems at a site

## Next topics

[Ways to determine whether a disaster occurred](#) on page 173

[Failures that do not require disaster recovery](#) on page 174

## Ways to determine whether a disaster occurred

You should declare a disaster only after using predefined procedures to verify that service cannot be restored.

It is critical that you follow a predefined procedure to confirm that a disaster occurred. The procedure should include determining the status of the disaster site by:

- Using external interfaces to the storage system, such as the following:
  - `ping` command to verify network connectivity
  - Remote shell

- FilerView administration tool
- Using network management tools to verify connectivity to the disaster site
- Physically inspecting the disaster site, if possible

You should declare a disaster only after verifying that service cannot be restored.

## Failures that do not require disaster recovery

If you can reestablish the MetroCluster connection after fixing the problem, you should not perform the disaster recovery procedure.

Do not perform the disaster recovery procedure for the following failures:

- A failure of the cluster interconnect between the two sites. This can be caused by the following:
  - Failure of the interconnect cable
  - Failure of one of the FC-VI adapters
  - If using switches, a failure of the SFP connecting a node to the switch

With this type of failure, both nodes remain running. Automatic takeover is disabled because Data ONTAP cannot synchronize the NVRAM logs. After you fix the problem and reestablish the connection, the nodes resynchronize their NVRAM logs and the MetroCluster returns to normal operation.

- The storage from one site (site A) is not accessible to the node at the other site (site B). This can be caused by the following:
  - Failure of any of the cables connecting the storage at one site to the node at the other site or switch
  - If using switches, failure of any of the SFPs connecting the storage to the switch or the node to the switch
  - Failure of the Fibre Channel adapter on the node
  - Failure of a storage disk shelf (disk shelf module; power; access to disk shelves; and so on)

With this type of failure, you see a “mailbox disk invalid” message on the console of the storage system that cannot see the storage. After you fix the problem and reestablish the connection, the MetroCluster returns to normal operation.

- If you are using switches, the inter-switch link between each pair of switches fails.
 

With this type of failure, both nodes remain running. You see a “mailbox disk invalid” message because a storage system at one site cannot see the storage at the other site. You also see a message because the two nodes cannot communicate with each other. After you fix the problem and reestablish the connection, the nodes resynchronize their NVRAM logs and the MetroCluster returns to normal operation.

## Recovering from a disaster

After determining that there is a disaster, you should take steps to recover access to data, fix problems at the disaster site, and re-create the MetroCluster configuration.

### About this task

Complete the following tasks in the order shown.

**Attention:** If for any reason the primary node has data that was not mirrored to the secondary prior to the execution of the `cf forcetakeover -d` command, data could be lost. Do not resynchronize the original disks of the primary site for a SnapLock volume until an additional backup has been made of those disks to assure availability of all data. This situation could arise, for example, if the link between the sites was down and the primary node had data written to it in the interim before the `cf forcetakeover -d` command was issued.

For more information about backing up data in SnapLock volumes using SnapMirror, see the *Data ONTAP Archive and Compliance Management Guide*.

1. [Restricting access to the disaster site node](#) on page 175
2. [Forcing a node into takeover mode](#) on page 176
3. [Remounting volumes of the failed node](#) on page 177
4. [Recovering LUNs of the failed node](#) on page 177
5. [Fixing failures caused by the disaster](#) on page 178
6. [Reestablishing the MetroCluster configuration](#) on page 179

## Restricting access to the disaster site node

You must restrict access to the disaster site node to prevent the node from resuming service. If you do not restrict access, you risk the possibility of data corruption.

### About this task

You can restrict access to the disaster site node in the following ways:

- Turning off power to the disaster site node.
- Manually fencing off the node.

1. [Restricting access to the node by turning off power](#) on page 175
2. [Restricting access to the node by fencing off](#) on page 176

## Restricting access to the node by turning off power

This is the preferred method for restricting access to the disaster site node. You can perform this task at the disaster site or remotely, if you have that capability.

**Step**

1. Switch off the power at the back of the storage system.

**Restricting access to the node by fencing off**

You can use manual fencing as an alternative to turning off power to the disaster site node. The manual fencing method restricts access using software and physical means.

**Steps**

1. Disconnect the cluster interconnect and Fibre Channel adapter cables of the node at the surviving site.
2. Use the appropriate fencing method depending on the type of failover you are using:

<b>If you are using...</b>	<b>Then fencing is achieved by...</b>
<b>Application failover</b>	Using any application-specified method that either prevents the application from restarting at the disaster site or prevents the application clients from accessing the application servers at the disaster site. Methods can include turning off the application server, removing an application server from the network, or any other method that prevents the application server from running applications.
<b>IP failover</b>	Using network management procedures to ensure that the storage systems at the disaster site are isolated from the external public network.

**Forcing a node into takeover mode**

If a disaster has occurred, you can force the surviving node into takeover mode, so that the surviving node serves the data of the failed node.

**Step**

1. Enter the following command on the surviving node:

```
cf forcetakeover -d
```

Data ONTAP causes the following to occur:

- The surviving node takes over the functions of the failed node.
- The mirrored relationships between the two plexes of mirrored aggregates are broken, thereby creating two unmirrored aggregates. This is called splitting the mirrored aggregates.

The overall result of using the `cf forcetakeover -d` command is that a node at the surviving site is running in takeover mode with all the data in unmirrored aggregates.



## Remounting volumes of the failed node

If the `cf.takeover.change_fsids` option is set to `on`, you must remount the volumes of the failed node because the volumes are accessed through the surviving node.

### About this task

For more information about mounting volumes, see the *Data ONTAP File Access and Protocols Management Guide*.

**Note:** You can disable the `change_fsids` option to avoid the necessity of remounting the volumes.

### Steps

1. On an NFS client at the surviving site, create a directory to act as a mount point.

#### Example

```
mkdir /n/toaster/home
```

2. Mount the volume.

#### Example

```
mount toaster:/vol/vol0/home /n/toaster/home
```

### Related tasks

[Disabling the `change\_fsids` option in MetroCluster configurations](#) on page 108

## Recovering LUNs of the failed node

You must actively track whether LUNs are online or offline in a MetroCluster configuration. If the `cf.takeover.change_fsids` option is set to `on`, and there is a disaster, all LUNs in the aggregates that were mirrored at the surviving site are offline. You can't determine if they were online prior to the disaster unless you track their state.

### About this task

If you have a MetroCluster configuration, you must actively track the state of LUNs (track whether they are online or offline) on the node at each site. If there is a failure to a MetroCluster configuration that qualifies as a disaster and the node at one site is inaccessible, all LUNs in the aggregates that were mirrored at the surviving site are offline. There is no way to distinguish the LUNs that were offline before the disaster from the LUNs that were online before the disaster unless you have been tracking their status.

When you recover access to the failed node's LUNs, it is important to bring back online only the LUNs that were online before the disaster. To avoid igroup mapping conflicts, do not bring a LUN online if it was offline before the disaster. For example, suppose you have two LUNs with IDs of 5 mapped to the same igroup, but one of these LUNs was offline before the disaster. If you bring the previously

offline LUN online first, you cannot bring the second LUN online because you cannot have two LUNs with the same ID mapped to the same host.

**Note:** You can disable the `change_fsid` option to avoid the necessity of remounting the volumes.

### Steps

1. Identify the LUNs that were online before the disaster occurred.
2. Make sure that the LUNs are mapped to an igroup that contains the hosts attached to the surviving node.

For more information about mapping LUNs to igroups, see your *Data ONTAP Block Access Management Guide for iSCSI and FC*.

3. On the surviving node, enter the following command:

```
lun online lun-path ...
```

*lun-path* is the path to the LUN you want to bring online. You can specify more than one path to bring multiple LUNs online.

#### Example

```
lun online /vol/vol1/lun5
```

#### Example

```
lun online /vol/vol1/lun3 /vol/vol1/lun4
```

**Note:** After you bring LUNs back online, you might have to perform some application or host-side recovery procedures. For example, the File System Identifiers (FSIDs) are rewritten, which can cause the LUN disk signatures to change. For more information, see the documentation for your application and for your host operating system.

## Fixing failures caused by the disaster

You need to fix the failures caused by the disaster, if possible. For example, if a prolonged power outage to one of the MetroCluster sites caused the failure, restoring the power fixes the failure.

### Before you begin

You cannot fix failures if the disaster causes a site to be destroyed. For example, a fire or an earthquake could destroy one of the MetroCluster sites. In this case, you fix the failure by creating a new MetroCluster-configured partner at a different site.

### Step

1. Fix the failures at the disaster site.

**After you finish**

After the node at the surviving site can see the disk shelves at the disaster site, Data ONTAP renames the mirrored aggregates that were split, and the volumes they contain, by adding a number in parenthesis to the name. For example, if a volume name was vol1 before the disaster and the split, the renamed volume name could be vol1(1).

**Reestablishing the MetroCluster configuration**

Describes how to reestablish a MetroCluster after a disaster, depending on the state of the mirrored aggregate at the time of the takeover.

**About this task**

Depending on the state of a mirrored aggregate before you forced the surviving node to take over its partner, you use one of two procedures to reestablish the MetroCluster configuration:

- If the mirrored aggregate was in a normal state before the forced takeover, you can rejoin the two aggregates to reestablish the MetroCluster configuration. This is the most typical case.
- If the mirrored aggregate was in an initial resynchronization state (level-0) before the forced takeover, you cannot rejoin the two aggregates. You must re-create the synchronous mirror to reestablish the MetroCluster configuration.

**Next topics**

[Rejoining the mirrored aggregates to reestablish a MetroCluster](#) on page 179

[Re-creating mirrored aggregates to return a MetroCluster to normal operation](#) on page 181

**Rejoining the mirrored aggregates to reestablish a MetroCluster**

Describes how to rejoin the mirrored aggregates if the mirrored aggregate was in a normal state before the forced takeover.

**Attention:** If you attempt a giveback operation prior to rejoining the aggregates, you might cause the node to boot with a previously failed plex, resulting in a data service outage.

**Steps**

1. Validate that you can access the remote storage by entering the following command:

```
aggr status -r
```

2. Turn on power to the node at the disaster site.

After the node at the disaster site boots, it displays the following message:

```
Waiting for Giveback...
```

3. Determine which aggregates are at the surviving site and which aggregates are at the disaster site by entering the following command:

**aggr status**

Aggregates at the disaster site show plexes that are in a failed state with an out-of-date status. Aggregates at the surviving site show plexes as online.

4. If aggregates at the disaster site are online, take them offline by entering the following command for each online aggregate:

```
aggr offline disaster_aggr
```

*disaster\_aggr* is the name of the aggregate at the disaster site.

**Note:** An error message appears if the aggregate is already offline.

5. Re-create the mirrored aggregates by entering the following command for each aggregate that was split:

```
aggr mirror aggr_name -v disaster_aggr
```

*aggr\_name* is the aggregate on the surviving site's node.

*disaster\_aggr* is the aggregate on the disaster site's node.

The *aggr\_name* aggregate rejoins the *disaster\_aggr* aggregate to reestablish the MetroCluster configuration.

6. Verify that the mirrored aggregates have been re-created by entering the following command:

```
aggr status -r mir
```

The giveback operation only succeeds if the aggregates have been rejoined.

7. Enter the following command at the partner node:

```
cf giveback
```

The node at the disaster site reboots.

### Example of rejoining aggregates

The following example shows the commands and status output when you rejoin aggregates to reestablish the MetroCluster configuration.

First, the aggregate status of the disaster site's storage after reestablishing access to the partner node at the surviving site is shown.

```
filer1> aggr status -r
Aggregate mir (online, normal) (zoned checksums)
  Plex /mir/plex5 (online, normal, active)
    RAID group /filer1/plex5/rg0 (normal)

RAID Disk Device HA  SHELF BAY CHAN  Used (MB/blks) Phys (MB/blks)
-----
parity  8a.2  8a   0    2    FC:B  34500/70656000  35003/71687368
```

```

data      8a.8   8a    1     0    FC:B  34500/70656000 35003/71687368

Aggregate mir(1) (failed, out-of-date) (zoned checksums)
  Plex /mir(1)/plex1 (offline, normal, out-of-date)
  RAID group /mir(1)/plex1/rg0 (normal)

RAID Disk Device HA  SHELF BAY CHAN  Used (MB/blks) Phys (MB/blks)
-----
parity    6a.0   6a    0     0    FC:B  34500/70656000 35003/71687368
data     6a.1   6a    0     1    FC:B  34500/70656000 35003/71687368

  Plex /mir(1)/plex5 (offline, failed, out-of-date)

```

Next, the mirror is reestablished using the `aggr mirror -v` command.

**Note:** The node at the surviving site is called `filer1`; the node at the disaster site is called `filer2`.

```

filer1> aggr mirror mir -v mir(1)
This will destroy the contents of mir(1).  Are you sure? y
Mon Nov 18 15:36:59 GMT [filer1: raid.mirror.resync.snapcertok:info]:
mir: created mirror resynchronization snapshot
mirror_resync.1118153658(filer2)
Mon Nov 18 15:36:59 GMT [filer1: raid.rg.resync.start:notice]:
/mir/plex6/rg0: start resynchronization (level 1)
Mon Nov 18 15:36:59 GMT [filer1: raid.mirror.resync.start:notice]: /mir:
start resynchronize to target /mir/plex6

```

After the aggregates rejoin, the synchronous mirrors of the MetroCluster configuration are reestablished.

```

filer1> aggr status -r mir
Aggregate mir (online, mirrored) (zoned checksums)
  Plex /mir/plex5 (online, normal, active)
  RAID group /mir/plex5/rg0 (normal)

RAID Disk Device HA  SHELF BAY CHAN  Used (MB/blks) Phys (MB/blks)
-----
parity    8a.2   8a    0     2    FC:B  34500/70656000 35003/71687368
data     8a.8   8a    1     0    FC:B  34500/70656000 35003/71687368

  Plex /mir/plex6 (online, normal, active)
  RAID group /mir/plex6/rg0 (normal)

RAID Disk Device HA  SHELF BAY CHAN  Used (MB/blks) Phys (MB/blks)
-----
parity    6a.0   6a    0     0    FC:B  34500/70656000 35003/71687368
data     6a.1   6a    0     1    FC:B  34500/70656000 35003/71687368

```

## Re-creating mirrored aggregates to return a MetroCluster to normal operation

Describes how to return the MetroCluster to normal operation by re-creating the MetroCluster mirror.

## Steps

1. Validate that you can access the remote storage by entering the following command:

```
aggr status -r
```

**Note:** A (level-0 resync in progress) message indicates that a plex cannot be rejoined.

2. Turn on the power to the node at the disaster site.

After the node at the disaster site boots up, it displays the following:

```
Waiting for Giveback...
```

3. If the aggregates at the disaster site are online, take them offline by entering the following command for each aggregate that was split:

```
aggr offline disaster_aggr
```

*disaster\_aggr* is the name of the aggregate at the disaster site.

**Note:** An error message appears if the aggregate is already offline.

4. Destroy every target plex that is in a level-0 resync state by entering the following command:

```
aggr destroy plex_name
```

For more information about the SyncMirror feature, see the *Data ONTAP Data Protection Online Backup and Recovery Guide*.

5. Re-create the mirrored aggregates by entering the following command for each aggregate that was split:

```
aggr mirror aggr_name
```

6. Enter the following command at the partner node:

```
cf giveback
```

The node at the disaster site reboots.

### Example of re-creating a mirrored aggregate

The following example shows the commands and status output when re-creating aggregates to reestablish the MetroCluster configuration.

First, the aggregate status of the disaster site's storage after reestablishing access to the partner at the surviving site is shown.

```
filer1>aggr status -r
Aggregate mir1 (online, normal) (zoned checksums)
  Plex /mir1/plex0 (online, normal, active)
  RAID group /mir1/plex0/rg0 (normal)
```

RAID	Disk	Device	HA	SHELF	BAY	CHAN	Used (MB/blks)	Phys (MB/blks)
parity	8a.3	8a	0	3	FC:B	34500/70656000	35003/71687368	
data	8a.4	8a	0	4	FC:B	34500/70656000	35003/71687368	
data	8a.6	8a	0	6	FC:B	34500/70656000	35003/71687368	
data	8a.5	8a	0	5	FC:B	34500/70656000	35003/71687368	

```
Aggregate mir1(1) (failed, partial) (zoned checksums)
  Plex /mir1(1)/plex0 (offline, failed, inactive)

  Plex /mir1(1)/plex6 (online, normal, resyncing)
    RAID group /mir1(1)/plex6/rg0 (level-0 resync in progress)
```

RAID	Disk	Device	HA	SHELF	BAY	CHAN	Used (MB/blks)	Phys (MB/blks)
parity	6a.6	6a	0	6	FC:B	34500/70656000	35003/71687368	
data	6a.2	6a	0	2	FC:B	34500/70656000	35003/71687368	
data	6a.3	6a	0	3	FC:B	34500/70656000	35003/71687368	
data	6a.5	6a	0	5	FC:B	34500/70656000	35003/71687368	

The mir1(1)/plex6 plex shows that a level-0 resynchronization was in progress; therefore, an attempt to rejoin the plexes fails, as shown in the following output:

```
filer1> aggr mirror mir1 -v mir1(1)
aggr mirror: Illegal mirror state for aggregate 'mir1(1)'
```

Because the mir1(1)/plex6 plex had a level-0 resynchronization in progress, the mir1(1) aggregate must be destroyed and the mir aggregate remirrored to reestablish a synchronous mirror, as shown in the following output:

```
filer1> aggr mirror mir1 -v mir1(1)
aggr mirror: Illegal mirror state for aggregate 'mir1(1)'
filer1> aggr destroy mir1(1)
Are you sure you want to destroy this aggregate? y
Aggregate 'mir1(1)' destroyed.
filer1> aggr mirror mir1
Creation of a mirror plex with 4 disks has been initiated. The disks
need to be zeroed before addition to the aggregate. The process has been
initiated and you will be notified via the system log as disks are
added.
```





# Nondisruptive hardware changes

---

By taking advantage of an active/active configuration's takeover and giveback operations, you can change hardware in your configuration without disrupting access to system storage.

For more information about nondisruptive Data ONTAP upgrades of active/active configurations, see the *Data ONTAP Upgrade Guide*.

**Note:** See the hardware documentation for your storage systems before upgrading hardware components.

## Replacing a component nondisruptively

You can use the nondisruptive upgrade method to perform a variety of hardware upgrade procedures.

### Before you begin

- If you are upgrading within a release family, you can upgrade the hardware and Data ONTAP together.
- If hardware and software are upgraded together, the new software must be downloaded first.
- If you are upgrading between release families, you need to upgrade Data ONTAP first, which might cause a small disruption, but you can use the nondisruptive upgrade procedure to cover the longer period required for hardware upgrades.

### About this task

You can perform hardware maintenance and upgrades using the nondisruptive upgrade method. You can perform nondisruptive hardware upgrades independently or in conjunction with a Data ONTAP upgrade of an active/active configuration.

In a nondisruptive upgrade, each node is successively “failed” and its partner is put in takeover state at one point in the procedure. While the “failed” node is in the waiting for giveback state, it can be interrupted and brought to the boot prompt or powered off.

You can use the nondisruptive upgrade method to perform the following hardware upgrade procedures:

- Replacing the controller (with a controller of the same type, and with the same adapters)
- Replacing the motherboard
- Replacing or adding an adapter

You can replace NVRAM, disk, or NIC components, either with the same component or with an improved component (for example, you can upgrade from 2-port to 4-port Gigabit Ethernet (GbE), or 1-port to 2-port Fibre Channel).

- Replacing the cluster interconnect adapter
- Installing onboard firmware on various platforms

You can also run diagnostics on various components of the “failed” machine (for example, motherboard or NVRAM).

**Note:** Running diagnostics on parts of the machine that are still visible to or in use by the partner—notably disks and the cluster interconnect adapter—can present problems. See the Diagnostics Guide.

Replacing a component nondisruptively involves the following tasks:

- Removing the component
- Installing the new component

The nondisruptive component replacement procedures use the following terms:

- The target node is the node that contains the failed component.
- The partner node is the node that is functioning normally, and that serves data for the active/active configuration while you replace the component.

Complete the following tasks in the order shown:

1. [Removing the old hardware when nondisruptively changing hardware](#) on page 186
2. [Installing the new hardware when nondisruptively changing hardware](#) on page 187

## Removing the old hardware when nondisruptively changing hardware

Describes how to remove a hardware component from a node when doing a nondisruptive hardware upgrade.

### Steps

1. Put on an antistatic wrist strap.
2. Take over the node that has the component you want to replace by entering the following command from the partner node’s console:

```
cf takeover
```

The partner node takes over the target node. You see a message similar to the following on the partner node’s console:

```
takeover completed
```

3. If your active/active configuration includes AT-FCX disk shelf modules, wait 5 minutes before proceeding to the next step.

**Attention:** If you attempt to proceed without the 5-minute wait, your systems could experience disk reservation conflicts, resulting in a system panic and data service outage.

4. Turn off the power to the target node, then unplug it from the power source.
5. Remove any applicable cabling to the component, then remove the component, as described in the *Hardware and Service Guide* or the hardware documentation for that component.

### After you finish

Proceed to install the new replacement hardware.

## Installing the new hardware when nondisruptively changing hardware

Describes how to install new hardware when doing a nondisruptive hardware upgrade.

### Steps

1. While grounded, install the new component in the target node, as described in the *Hardware and Service Guide* or the hardware documentation for that component.
2. Close the chassis and reconnect any applicable cables to the component.
3. Plug the target node back into the power source, and then apply power.

**Note:** Before applying power, read the next step to determine whether you need to interrupt the boot process.

4. If your active/active configuration is using software-based disk ownership and you replaced the NVRAM adapter, the system ID has changed and you must reassign the disks. To reassign the disks, complete the following substeps:

- a. Interrupt the target node boot process and boot into Maintenance mode.
- b. Determine the new target node system ID by entering the following command on the target node:

```
disk show -v
```

- c. On the partner node, reassign the disks by entering the following commands:

```
priv set advanced
```

```
disk reassign -d new_sysID
```

**Note:** You cannot reassign more than 500 disks from one controller to another by using the `disk reassign` command. If you try to do so, the system reports an error. If you want to reassign more than 500 disks, contact technical support.

**Note:** The `disk reassign` command does not modify SyncMirror disk pool settings in cases where software-based disk ownership is enabled.

- d. Return to the boot process by entering the following commands on the target node:

```
halt
```

```
boot
```

5. After the target node boots and is at the ‘Waiting for giveback’ prompt, run giveback by entering the following command on the partner node’s console:

```
cf giveback
```

The target node reboots and functions normally. A successful giveback ends with the following message on the partner node:

```
giveback completed
```

# Controller failover and single-points-of-failure

---

Lists the single-points-of-failure (SPOFs) and the failover causes. A storage system has a variety of SPOFs that you can reduce through active/active configuration. In an active/active configuration, there are a number of failures that can cause a controller to fail over.

## Benefits of controller failover

You can use controller failover, a high-availability data service solution, to further increase the uptime of storage systems. It protects against controller failure by transferring the data service from the failed node to its partner node. Controller failover can also protect against other hardware failures, such as problems with network interface cards, Fibre Channel-Arbitrated Loops (FC-AL loops), and disk shelf modules. Controller failover is also an effective tool for reducing planned downtime of one of the nodes.

**Note:** You might also see the term cluster failover; this is equivalent to the term controller failover used in this document.

## Next topics

[Single-point-of-failure definition](#) on page 189

[SPOF analysis for active/active configurations](#) on page 189

[Failover event cause-and-effect table](#) on page 192

## Single-point-of-failure definition

Explains what a single-point-of-failure is in the context of your storage system.

A single-point-of failure (SPOF) represents the failure of a single hardware component that can lead to loss of data access or potential loss of data. SPOF does not include multiple/rolling hardware errors, such as triple disk failure, dual disk shelf module failure, and so on.

All hardware components included with your storage system have demonstrated very good reliability with low failure rates. If a hardware component fails, such as a controller or adapter, you can use controller failover to provide continuous data availability and preserve data integrity for client applications and users.

## SPOF analysis for active/active configurations

Enables you to see which hardware components are SPOFs, and how controller failover can eliminate these SPOFs to improve data availability.

Hardware components	SPOF		How controller failover eliminates SPOF
	Stand-alone	Active/active configuration	
Controller	Yes	No	If a controller fails, the node automatically fails over to its partner node. The partner (takeover) node serves data for both of the nodes.
NVRAM	Yes	No	If an NVRAM adapter fails, the node automatically fails over to its partner node. The partner (takeover) node serves data for both of the nodes.
CPU fan	Yes	No	If the CPU fan fails, the node gracefully shuts down and automatically fails over to its partner node. The partner (takeover) node serves data for both of the nodes.
Multiple NICs with VIFs (virtual interfaces)	No	No	If one of the networking links within a VIF fails, the networking traffic is automatically sent over the remaining networking links on the same node. No failover is needed in this situation.  If all the NICs in a VIF fail, the node automatically fails over to its partner node, if failover is enabled for the VIF.
Single NIC	Yes	No	If a NIC fails, the node automatically fails over to its partner node, if failover is enabled for the NIC.
FC-AL adapter	Yes	No	If an FC-AL adapter for the primary loop fails for a configuration without redundant paths, the partner node attempts a failover at the time of failure. With redundant paths, no failover is required.  If the FC-AL adapter for the secondary loop fails for a configuration without redundant paths, the failover capability is disabled, but both nodes continue to serve data to their respective applications and users, with no impact or delay. With redundant paths, failover capability is not affected.
FC-AL cable (controller-to-shelf, shelf-to-shelf)	Yes	No	If an FC-AL loop breaks in a configuration that does not have redundant paths, the break could lead to a failover, depending on the shelf type. The partnered nodes invoke the negotiated failover feature to determine which node is best for serving data, based on the disk shelf count. When redundant paths are present, no failover is required.

Hardware components	SPOF		How controller failover eliminates SPOF
	Stand-alone	Active/active configuration	
Disk shelf module	Yes	No	If a disk shelf module fails in a configuration that does not have redundant paths, the failure could lead to a failover. The partnered nodes invoke the negotiated failover feature to determine which node is best for serving data, based on the disk shelf count. When redundant paths are present, there is no impact.
Disk drive	No	No	If a disk fails, the node can reconstruct data from the RAID4 parity disk. No failover is needed in this situation.
Disk shelf (including backplane)	No	No	A disk shelf is a passive backplane with dual power supplies, dual fans, dual modules, and dual FC-AL loops. It is the single most reliable component in a storage system.
Power supply	No	No	Both the controller and disk shelf have dual power supplies. If one power supply fails, the second power supply automatically kicks in. No failover is needed in this situation. If both power supplies fail, the node automatically fails over to its partner node, which serves data for both nodes.
Fan (controller or disk shelf)	No	No	Both the controller and disk shelf have multiple fans. If one fan fails, the second fan automatically provides cooling. No failover is needed in this situation. If both fans fail, the node automatically fails over to its partner node, which serves data for both nodes.
Cluster adapter	N/A	No	If a cluster adapter fails, the failover capability is disabled but both nodes continue to serve data to their respective applications and users.
Cluster interconnect cable	N/A	No	The cluster adapter supports dual cluster interconnect cables. If one cable fails, the heartbeat and NVRAM data are automatically sent over the second cable with no delay or interruption.  If both cables fail, the failover capability is disabled but both nodes continue to serve data to their respective applications and users.

## Failover event cause-and-effect table

Helps you understand what happens when a failover event occurs on FAS systems or V-Series systems using native disk shelves, and how the various active/active configurations handle these events.

Event	Does the event trigger failover?		Does the event prevent a future failover from occurring, or a failover from occurring successfully?		Is data still available on the affected volume after the event?		
	Standard or Mirrored	MetroCluster	Standard or Mirrored	MetroCluster	Single Storage System	Standard or Mirrored	Fabric Attached MetroCluster
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 SyncMirror or RAID-DP, then no.	No	Maybe. If root volume has double disk failure or if the mailbox disks are affected, no failover is possible.	No	No, unless you are using RAID-DP or SyncMirror, then yes.	No, unless you are using RAID-DP or SyncMirror, then yes.	Yes, with no failover necessary.
Triple disk failure (3 disks fail in same RAID group)	Maybe. If SyncMirror is being used, there won't be a takeover; otherwise, yes.	No	Maybe. If root volume has triple disk failure, no failover is possible.	No	No	No	Yes, with no failover necessary.
Single HBA (initiator) failure, Loop A	Maybe. If SyncMirror or redundant paths are in use, then no; otherwise, yes.	No	Maybe. If root volume has double disk failure, no failover is possible.	No	Yes, if redundant paths or SyncMirror is being used.	Yes, if redundant paths or SyncMirror is being used, or if failover succeeds.	Yes, with no failover necessary.



Event	Does the event trigger failover?		Does the event prevent a future failover from occurring, or a failover from occurring successfully?		Is data still available on the affected volume after the event?		
	Standard or Mirrored	MetroCluster	Standard or Mirrored	MetroCluster	Single Storage System	Standard or Mirrored	Fabric Attached MetroCluster
Single HBA (initiator) failure, Loop B	No	No	Yes, unless you are using SyncMirror or redundant paths and the mailbox disks aren't affected, then no.	No	Yes, if redundant paths or SyncMirror is being used.	Yes, if redundant paths or SyncMirror is being used, or if failover succeeds.	Yes, with no failover necessary.
Single HBA initiator failure, (both loops at the same time)	Yes, unless the data is mirrored on a different (up) loop or redundant paths are in use, then no takeover needed.	No	Maybe. If the data is mirrored or redundant paths are being used and the mailbox disks are not affected, then no; otherwise, yes.	No	No, unless the data is mirrored or redundant paths are in use, then yes.	No failover needed if data is mirrored or redundant paths are in use.	Yes, with no failover necessary.
LRC failure (Loop A)	Only if multidisk volume failure or open loop condition occurs, and redundant paths are not in use.	No	Maybe. If redundant paths are not in use and the root volume has a double disk failure, no failover is possible because this impacts the mailbox disks.	No	Yes, if redundant paths or SyncMirror is in use.	Yes, if failover succeeds or if redundant paths or SyncMirror is in use.	Yes, with no failover necessary.

Event	Does the event trigger failover?		Does the event prevent a future failover from occurring, or a failover from occurring successfully?		Is data still available on the affected volume after the event?		
	Standard or Mirrored	MetroCluster	Standard or Mirrored	MetroCluster	Single Storage System	Standard or Mirrored	Fabric Attached MetroCluster
LRC failure (Loop B)	No	No	Maybe. If the data is mirrored or redundant paths are in use, and the mailbox disks aren't affected, automatic failover happens.	No	Yes, if redundant paths or SyncMirror is in use.	No	No
ESH2 or AT-FCX failure (Loop A)	Only if multidisk volume failure or open loop condition occurs, and neither SyncMirror nor redundant paths are in use.	No	Maybe. If root volume has double disk failure, no failover is possible.	No	No	Yes, if failover succeeds.	Yes, with no failover necessary.
ESH2 or AT-FCX failure (Loop B)	No	No	Maybe. If SyncMirror or redundant paths are in use, then no; otherwise, yes.	No	Yes, if redundant paths or SyncMirror is in use.	Yes	Yes

Event	Does the event trigger failover?		Does the event prevent a future failover from occurring, or a failover from occurring successfully?		Is data still available on the affected volume after the event?		
	Standard or Mirrored	MetroCluster	Standard or Mirrored	MetroCluster	Single Storage System	Standard or Mirrored	Fabric Attached MetroCluster
Shelf (backplane) failure	Only if multidisk volume failure or open loop condition occurs, and data isn't mirrored.	No	Maybe. If root volume has double disk failure or if the mailboxes are affected, no failover is possible.	No	Maybe. If data is mirrored, then yes; otherwise, no.	Maybe. If data is mirrored, then yes; otherwise, no.	Yes, with no failover necessary.
Shelf, single power failure	No	No	No	No	Yes	Yes	Yes
Shelf, dual power failure	Only if multidisk volume failure or open loop condition occurs and data isn't mirrored.	No	Maybe. If root volume has double disk failure or if the mailbox disks are affected, no failover is possible.	No	Maybe. If data is mirrored, then yes; otherwise, no.	Maybe. If data is mirrored, then yes; otherwise, no.	Yes, with no failover necessary.
Controller, single power failure	No	No	No	No	Yes	Yes	Yes
Controller, dual power failure	Yes	Yes	Yes, until power is restored.	Yes, until power is restored.	No	Yes, if failover succeeds.	Yes, if failover succeeds.
Cluster interconnect failure (1 port)	No	No	No	No	n/a	Yes	Yes
Cluster interconnect failure (both ports)	No	No	Yes	No. Failover is possible.	n/a	Yes	Yes

Event	Does the event trigger failover?		Does the event prevent a future failover from occurring, or a failover from occurring successfully?		Is data still available on the affected volume after the event?		
	Standard or Mirrored	MetroCluster	Standard or Mirrored	MetroCluster	Single Storage System	Standard or Mirrored	Fabric Attached MetroCluster
Ethernet interface failure (primary, no VIF)	Yes, if set up to do so	Yes, if set up to do so	No	No	Yes	Yes	Yes
Ethernet interface failure (primary, VIF)	Yes, if set up to do so	Yes, if set up to do so	No	No	Yes	Yes	Yes
Ethernet interface failure (secondary, VIF)	Yes, if set up to do so	Yes, if set up to do so	No	No	Yes	Yes	Yes
Ethernet interface failure (VIF, all ports)	Yes, if set up to do so	Yes, if set up to do so	No	No	Yes	Yes	Yes
Tape interface failure	No	No	No	No	Yes	Yes	Yes
Heat exceeds permissible amount	Yes	Yes	No	No	No	No	No
Fan failures (disk shelves or controller)	No	No	No	No	Yes	Yes	Yes
Reboot	No	No	No	No	Maybe. Depends on cause of reboot.	Maybe. Depends on cause of reboot.	Maybe. Depends on cause of reboot.

Event	Does the event trigger failover?		Does the event prevent a future failover from occurring, or a failvoer from occurring successfully?		Is data still available on the affected volume after the event?		
	Standard or Mirrored	MetroCluster	Standard or Mirrored	MetroCluster	Single Storage System	Standard or Mirrored	Fabric Attached MetroCluster
Panic	No	No	No	No	Maybe. Depends on cause of panic.	Maybe. Depends on cause of panic.	Maybe. Depends on cause of panic.



## Feature update record

---

Provides a record of the history of changes made to this guide. 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
<ul style="list-style-type: none"> <li>• Updates for FAS920</li> <li>• Update for NVRAM5</li> <li>• Illustration updates</li> </ul>	Data ONTAP 6.5.1	May 2004
<ul style="list-style-type: none"> <li>• Updates for NVRAM5 support in FAS900 series active/active configurations, except for MetroCluster</li> <li>• Failover event cause-and-effect table</li> <li>• Declaration of Conformity update</li> <li>• Addition of controller failover and single-point-of-failure analysis</li> </ul>	Data ONTAP 7.0	November 2004
<ul style="list-style-type: none"> <li>• FAS30xx information</li> <li>• Corrections were made to the <i>Upgrading an LRC to ESH/ESH2/AT-FCX</i> procedure</li> </ul>	Data ONTAP 7.0.1	April 2005
<ul style="list-style-type: none"> <li>• Incorporation of the Cluster Administration chapter from the <i>Data ONTAP System Administration Guide</i> and the <i>Disaster Protection Using MetroCluster</i> appendix from the <i>Data ONTAP Data Protection Online Backup and Recovery Guide</i>.</li> </ul>	Data ONTAP 7.1	June 2005
<ul style="list-style-type: none"> <li>• Updated MetroCluster information for FAS30xx</li> </ul>	Data ONTAP 7.1	October 2005

Feature updates	Feature first implemented in	Feature release date
<ul style="list-style-type: none"> <li>• Updated module replacement information</li> <li>• Fixed problem in Brocade switch configuration information</li> </ul>	Data ONTAP 7.1	December 2005
<ul style="list-style-type: none"> <li>• Updated and extended active/active configuration information</li> <li>• Moved Brocade switch configuration to Brocade Switch Description Page.</li> <li>• Moved from <i>cluster</i> to <i>active/active configuration</i></li> <li>• Added information about Multipath Storage for Active/active configurations</li> </ul>	Data ONTAP 7.1.1	June 2006
<ul style="list-style-type: none"> <li>• Generalized standard and mirrored active/active configuration cabling instructions</li> <li>• Updated standard and mirrored active/active configuration cabling instructions to include FAS60xx</li> </ul>	Data ONTAP 7.2 RC1	February 2006
<ul style="list-style-type: none"> <li>• Changed name of document from <i>Cluster Installation and Administration Guide</i> to <i>Active/Active Configuration Guide</i>.</li> <li>• Added FAS60xx information</li> <li>• Updated and extended active/active configurations configuration information</li> <li>• Moved Brocade switch configuration to Brocade Switch Description Page.</li> <li>• Moved from <i>cluster</i> to <i>active/active configuration</i></li> </ul>	Data ONTAP 7.2 RC3	May 2006



Feature updates	Feature first implemented in	Feature release date
<ul style="list-style-type: none"> <li>Added information about Multipath Storage for Active/active configurations.</li> </ul>	Data ONTAP 7.2.1	November 2006
<ul style="list-style-type: none"> <li>Added quad-port, 4-Gb Fibre Channel HBA, ESH4 module, DS14mk4 FC disk shelf information</li> <li>Added information to explain that automatic giveback should not be used in MetroClusters</li> <li>Updated Multipath Storage information</li> <li>Updated MetroCluster disaster recovery information</li> <li>Corrected failover and single-point-of-failure table</li> </ul>	Data ONTAP 7.2.2	March 2007
<ul style="list-style-type: none"> <li>Added procedures for configuring fabric-attached MetroClusters on systems using software-based disk management</li> <li>Added procedure for unconfiguring an active/active pair and returning to stand-alone operation</li> </ul>	Data ONTAP 7.2.3	June 2007
<ul style="list-style-type: none"> <li>Added support for 504 disks in MetroClusters</li> <li>Added support for the FAS6040 and FAS6080 systems</li> <li>Added support for the <code>change_fsid</code> option</li> <li>Added procedure for removing an active/active configuration</li> </ul>	Data ONTAP 7.2.4	November 2007

Feature updates	Feature first implemented in	Feature release date
<ul style="list-style-type: none"><li>• Added V-Series content</li><li>• Added references to the DS4243 disk shelf documentation</li><li>• Added support for 672 disks in MetroClusters</li><li>• Added MetroCluster support for the Brocade 300 and 5100 switches</li><li>• Add support for MetroCluster nodes on separate subnetworks</li></ul>	Data ONTAP 7.3.2	September 2009

# Abbreviations

---

A list of abbreviations and their spelled-out forms are included here for your reference.

## A

*ABE (Access-Based Enumeration)*

*ACE (Access Control Entry)*

*ACL (access control list)*

*AD (Active Directory)*

*ALPA (arbitrated loop physical address)*

*ALUA (Asymmetric Logical Unit Access)*

*AMS (Account Migrator Service)*

*API (Application Program Interface)*

*ARP (Address Resolution Protocol)*

*ASCII (American Standard Code for Information Interchange)*

*ASP (Active Server Page)*

*ATA (Advanced Technology Attachment)*

## B

*BCO (Business Continuance Option)*

*BIOS (Basic Input Output System)*

*BURT (Bug Reporting Tool)*

*BCS (block checksum )*

*BLI (block-level incremental)*

**C**

*CD-ROM (compact disc read-only memory)*

*CDDI (Copper Distributed Data Interface)*

*CDN (content delivery network)*

*CFE (Common Firmware Environment)*

*CFO (cluster failover)*

*CGI (Common Gateway Interface)*

*CHA (channel adapter)*

*CHAP (Challenge Handshake Authentication Protocol)*

*CHIP (Client-Host Interface Processor)*

*CIDR (Classless Inter-Domain Routing)*

*CIFS (Common Internet File System)*

*CIM (Common Information Model)*

*CLI (command-line interface)*

*CP (consistency point)*

*CPU (central processing unit)*

*CRC (cyclic redundancy check)*

*CSP (communication service provider)*

**D**

*DAFS (Direct Access File System)*

*DBBC (database consistency checker)*

*DCE (Distributed Computing Environment)*

*DDS (Decru Data Decryption Software)*

*dedupe (deduplication)*

*DES (Data Encryption Standard)*

*DFS (Distributed File System)*

*DHA (Decru Host Authentication)*

*DHCP (Dynamic Host Configuration Protocol)*

*DIMM (dual-inline memory module)*

*DITA (Darwin Information Typing Architecture)*

*DLL (Dynamic Link Library)*

*DMA (direct memory access)*

*DMTD (Distributed Management Task Force)*

*DNS (Domain Name System)*

*DOS (Disk Operating System)*

*DPG (Data Protection Guide)*

*DTE (Data Terminal Equipment)*

**E**

*ECC (Elliptic Curve Cryptography) or (EMC Control Center)*  
*ECDN (enterprise content delivery network)*  
*ECN (Engineering Change Notification)*  
*EEPROM (electrically erasable programmable read-only memory)*  
*EFB (environmental fault bus)*  
*EFS (Encrypted File System)*  
*EGA (Enterprise Grid Alliance)*  
*EISA (Extended Infrastructure Support Architecture)*  
*ELAN (Emulated LAN)*  
*EMU environmental monitoring unit)*  
*ESH (embedded switching hub)*

**F**

*FAQs (frequently asked questions)*  
*FAS (fabric-attached storage)*  
*FC (Fibre Channel)*  
*FC-AL (Fibre Channel-Arbitrated Loop)*  
*FC SAN (Fibre Channel storage area network)*  
*FC Tape SAN (Fibre Channel Tape storage area network)*  
*FC-VI (virtual interface over Fibre Channel)*  
*FCP (Fibre Channel Protocol)*  
*FDDI (Fiber Distributed Data Interface)*  
*FQDN (fully qualified domain name)*  
*FRS (File Replication Service)*  
*FSID (file system ID)*  
*FSRM (File Storage Resource Manager)*  
*FTP (File Transfer Protocol)*

**G**

*GbE (Gigabit Ethernet)*

*GID (group identification number)*

*GMT (Greenwich Mean Time)*

*GUI (graphical user interface)*

*GUID (globally unique identifier)*

**H**

*HA (high availability)*

*HBA (host bus adapter)*

*HDM (Hitachi Device Manager Server)*

*HP (Hewlett-Packard Company)*

*HTML (hypertext markup language)*

*HTTP (Hypertext Transfer Protocol)*

**I**

*IB (InfiniBand)*

*IBM (International Business Machines Corporation)*

*ICAP (Internet Content Adaptation Protocol)*

*ICP (Internet Cache Protocol)*

*ID (identification)*

*IDL (Interface Definition Language)*

*ILM (information lifecycle management)*

*IMS (If-Modified-Since)*

*I/O (input/output)*

*IP (Internet Protocol)*

*IP SAN (Internet Protocol storage area network)*

*IQN (iSCSI Qualified Name)*

*iSCSI (Internet Small Computer System Interface)*

*ISL (Inter-Switch Link)*

*iSNS (Internet Storage Name Service)*

*ISP (Internet storage provider)*

**J**

*JBOD (just a bunch of disks)*

*JPEG (Joint Photographic Experts Group)*

**K**

*KB (Knowledge Base)*

*Kbps (kilobits per second)*

*KDC (Kerberos Distribution Center)*



**L**

*LAN (local area network)*

*LBA (Logical Block Access)*

*LCD (liquid crystal display)*

*LDAP (Lightweight Directory Access Protocol)*

*LDEV (logical device)*

*LED (light emitting diode)*

*LFS (log-structured file system)*

*LKM (Lifetime Key Management)*

*LPAR (system logical partition)*

*LRC (Loop Resiliency Circuit)*

*LREP (logical replication tool utility)*

*LUN (logical unit number)*

*LUSE (Logical Unit Size Expansion)*

*LVM (Logical Volume Manager)*

**M**

*MAC (Media Access Control)*

*Mbps (megabits per second)*

*MCS (multiple connections per session)*

*MD5 (Message Digest 5)*

*MDG (managed disk group)*

*MDisk (managed disk)*

*MIB (Management Information Base)*

*MIME (Multipurpose Internet Mail Extension)*

*MMC (Microsoft Management Console)*

*MMS (Microsoft Media Streaming)*

*MPEG (Moving Picture Experts Group)*

*MPIO (multipath network input/output)*

*MRTG (Multi-Router Traffic Grapher)*

*MSCS (Microsoft Cluster Service)*

*MSDE (Microsoft SQL Server Desktop Engine)*

*MTU (Maximum Transmission Unit)*

**N**

*NAS (network-attached storage)*

*NDMP (Network Data Management Protocol)*

*NFS (Network File System)*

*NIC (network interface card)*

*NMC (Network Management Console)*

*NMS (network management station)*

*NNTP (Network News Transport Protocol)*

*NTFS (New Technology File System)*

*NTLM (NetLanMan)*

*NTP (Network Time Protocol)*

*NVMEM (nonvolatile memory management)*

*NVRAM (nonvolatile random-access memory)*

**O**

*OFM (Open File Manager)*

*OFW (Open Firmware)*

*OLAP (Online Analytical Processing)*

*OS/2 (Operating System 2)*

*OSMS (Open Systems Management Software)*

*OSSV (Open Systems SnapVault)*

**P**

*PC (personal computer)*

*PCB (printed circuit board)*

*PCI (Peripheral Component Interconnect)*

*pcnfsd (storage daemon)*

*(PC)NFS (Personal Computer Network File System)*

*PDU (protocol data unit)*

*PKI (Public Key Infrastructure)*

*POP (Post Office Protocol)*

*POST (power-on self-test)*

*PPN (physical path name)*

*PROM (programmable read-only memory)*

*PSU power supply unit)*

*PVC (permanent virtual circuit)*

**Q**

*QoS (Quality of Service)*

*QSM (Qtree SnapMirror)*

**R**

*RAD (report archive directory)*

*RADIUS (Remote Authentication Dial-In Service)*

*RAID (redundant array of independent disks)*

*RAID-DP (redundant array of independent disks, double-parity)*

*RAM (random access memory)*

*RARP (Reverse Address Resolution Protocol)*

*RBAC (role-based access control)*

*RDB (replicated database)*

*RDMA (Remote Direct Memory Access)*

*RIP (Routing Information Protocol)*

*RISC (Reduced Instruction Set Computer)*

*RLM (Remote LAN Module)*

*RMC (remote management controller)*

*ROM (read-only memory)*

*RPM (revolutions per minute)*

*rsh (Remote Shell)*

*RTCP (Real-time Transport Control Protocol)*

*RTP (Real-time Transport Protocol)*

*RTSP (Real Time Streaming Protocol)*

**S**

*SACL (system access control list)*

*SAN (storage area network)*

*SAS (storage area network attached storage) or (serial-attached SCSI)*

*SATA (serial advanced technology attachment)*

*SCSI (Small Computer System Interface)*

*SFO (storage failover)*

*SFSR (Single File SnapRestore operation)*

*SID (Secure ID)*

*SIMM (single inline memory module)*

*SLB (Server Load Balancer)*

*SLP (Service Location Protocol)*

*SNMP (Simple Network Management Protocol)*

*SNTP (Simple Network Time Protocol)*

*SP (Storage Processor)*

*SPN (service principal name)*

*SPOF (single point of failure)*

*SQL (Structured Query Language)*

*SRM (Storage Resource Management)*

*SSH (Secure Shell)*

*SSL (Secure Sockets Layer)*

*STP (shielded twisted pair)*

*SVC (switched virtual circuit)*

**T**

*TapeSAN (tape storage area network)*

*TCO (total cost of ownership)*

*TCP (Transmission Control Protocol)*

*TCP/IP (Transmission Control Protocol/Internet Protocol)*

*TOE (TCP offload engine)*

*TP (twisted pair)*

*TSM (Tivoli Storage Manager)*

*TTL (Time To Live)*

**U**

*UDP (User Datagram Protocol)*

*UI (user interface)*

*UID (user identification number)*

*Ultra ATA (Ultra Advanced Technology Attachment)*

*UNC (Uniform Naming Convention)*

*UPS (uninterruptible power supply)*

*URI (universal resource identifier)*

*URL (uniform resource locator)*

*USP (Universal Storage Platform)*

*UTC (Universal Coordinated Time)*

*UTP (unshielded twisted pair)*

*UUID (universal unique identifier)*

*UWN (unique world wide number)*

**V**

*VCI (virtual channel identifier)*

*VCMDB (Volume Configuration Management Database)*

*VDI (Virtual Device Interface)*

*VDisk (virtual disk)*

*VDS (Virtual Disk Service)*

*VFM (Virtual File Manager)*

*VFS (virtual file system)*

*VI (virtual interface)*

*vif (virtual interface)*

*VIRD (Virtual Router ID)*

*VLAN (virtual local area network)*

*VLD (virtual local disk)*

*VOD (video on demand)*

*VOIP (voice over IP)*

*VRML (Virtual Reality Modeling Language)*

*VTL (Virtual Tape Library)*

**W**

*WAFL (Write Anywhere File Layout)*

*WAN (wide area network)*

*WBEM (Web-Based Enterprise Management)*

*WHQL (Windows Hardware Quality Lab)*

*WINS (Windows Internet Name Service)*

*WORM (write once, read many)*

*WWN (worldwide name)*

*WWNN (worldwide node name)*

*WWPN (worldwide port name)*

*www (worldwide web)*

**X**

**Y**

**Z**

*ZCS (zoned checksum)*



# Index

- ## A
- active/active configurations
    - benefits of [20](#)
    - changing nodes to stand-alone [95, 96, 97, 99](#)
    - characteristics of [20](#)
    - converting to MetroCluster [63](#)
    - definition of [19](#)
    - restrictions [24](#)
    - setup requirements for [24](#)
    - types of
      - compared [22](#)
      - fabric-attached MetroClusters [33](#)
      - installed in equipment racks [41](#)
      - installed in system cabinets [42](#)
      - mirrored [26](#)
      - standard [23](#)
      - stretch MetroClusters [28](#)
  - adapters [44, 46, 51, 61, 186, 187](#)
    - installing [187](#)
    - quad-port FibreChannel HBA [46, 51](#)
    - removing [186](#)
  - aggregates
    - recreating mirrored after disaster [181](#)
    - rejoining after disaster [179](#)
  - AT-FCX modules and Multipath Storage [155](#)
  - automatic giveback [148](#)
- ## B
- bring up [107, 112](#)
    - configuring interfaces for [112](#)
    - manually setting options for [107](#)
  - Brocade switch configuration [70, 71](#)
    - switch bank rules [71](#)
    - virtual channel rules [71](#)
- ## C
- cable [44, 61](#)
  - cabling
    - Channel A
      - for mirrored active/active configuration [53](#)
      - for standard active/active configuration [48](#)
    - Channel B
      - for mirrored active/active configuration [54](#)
      - for standard active/active configuration [49](#)
    - cluster interconnect for fabric-attached MetroClusters
      - with hardware-based disk ownership [79, 89](#)
      - with software-based disk ownership [90](#)
    - cluster interconnect for mirrored active/active configuration [56](#)
    - cluster interconnect for standard active/active configuration [50](#)
    - cross-cabled cluster interconnect [50](#)
    - error message, cross-cabled cluster interconnect [50](#)
    - fabric-attached MetroClusters [68](#)
    - FC-VI adapter for fabric-attached MetroClusters
      - with hardware-based disk ownership [79, 89](#)
      - with software-based disk ownership [80, 90](#)
    - local controller in fabric-attached MetroCluster
      - with hardware-based disk ownership [72](#)
      - with software-based disk ownership [73](#)
    - local disk shelves in fabric-attached MetroCluster
      - with hardware-based disk ownership [75](#)
      - with software-based disk ownership [77](#)
    - Multipath Storage [156](#)
    - preparing equipment racks for [45](#)
    - preparing system cabinets for [45](#)
    - remote controller in fabric-attached MetroCluster
      - with hardware-based disk ownership [82](#)
      - with software-based disk ownership [83](#)
    - remote disk shelves in fabric-attached MetroCluster
      - with hardware-based disk ownership [85](#)
      - with software-based disk ownership [86](#)
    - requirements [44, 61](#)
    - stretch MetroClusters [67](#)
  - cf forcegiveback command [146](#)
  - cf giveback command [145](#)
  - cf-config-check.cgi utility [149](#)
  - cf.giveback.auto.cifs.terminate.minutes options [147](#)
  - cf.giveback.auto.enable option [148](#)
  - cf.giveback.auto.terminate.bigjobs option [148](#)
  - cf.giveback.check.partner option [147](#)
  - cf.takeover.on\_network\_interface\_failure option [137](#)
  - cf.takeover.on\_network\_interface\_failure.policy option [137](#)
  - cf.takeover.use\_mcrf\_file [123](#)

change\_fsid option [108](#)

Channel A

cabling [48, 53](#)

defined [27](#)

Channel B

cabling [49, 54](#)

checking configuration through a utility [149](#)

CIFS clients and giveback delay [147](#)

CIFS sessions terminated on takeover [126](#)

cluster interconnect connections, tips [81, 90](#)

cluster interconnect, cabling [50, 56, 79, 80, 89, 90](#)

command exceptions for emulated nodes [141](#)

commands

cf (enables and disables takeover) [132](#)

cf forcesgiveback (forces giveback) [146](#)

cf forcetakeover -d (forces takeover) [176](#)

cf forcetakeover (forces takeover) [135](#)

cf giveback (enables giveback) [123](#)

cf giveback (initiates giveback) [145](#)

cf partner (displays partner's name) [131](#)

cf status (displays status) [127, 138](#)

cf takeover (initiates takeover) [135](#)

cf takeover (initiates takeover) [123](#)

halt (halts system without takeover) [133](#)

license add (license cluster) [106](#)

partner (accesses emulated node) [139](#)

storage show disk -p (displays paths) [167](#)

sysconfig [131](#)

takeover (description of all takeover commands) [135](#)

comparison of types of active/active configurations [22](#)

configuration variations

fabric-attached MetroCluster configurations [38](#)

mirrored active/active configurations [28](#)

standard active/active configurations [25](#)

stretch MetroClusters [32](#)

configurations

reestablishing MetroCluster configuration [179](#)

testing [123](#)

configuring

dedicated and standby interfaces [117](#)

shared interfaces [117](#)

controller failover

benefits [189](#)

controller-to-switch cabling, fabric-attached MetroClusters [72, 73, 82, 83](#)

controller-to-switch connections, tips [75, 84](#)

## D

Data ONTAP

in a standard active/active configurations [23](#)

in fabric-attached MetroCluster configurations [36](#)

in stretch MetroCluster configurations [31](#)

dedicated interfaces

configuring using ifconfig [117](#)

configuring using setup [105](#)

described [114](#)

diagram [116](#)

delay, specifying before takeover [136](#)

disabling takeover (cf) [132](#)

disasters

determining whether one occurred [173](#)

recognizing [173](#)

recovery from

forcing takeover [176](#)

manually fencing off the disaster site node [176](#)

reestablishing MetroCluster configuration [179](#)

restricting access to the failed node [175](#)

steps [175](#)

using MetroCluster [173](#)

when not to perform [174](#)

disk information, displaying [131](#)

disk ownership [31](#)

disk paths, verifying in a fabric-attached MetroCluster

with hardware-based disk ownership [92](#)

with software-based disk ownership [93](#)

disk shelf pool assignments, fabric-attached MetroClusters [91](#)

disk shelf-to-switch connections, tips [78, 88](#)

disk shelves

about modules for [165](#)

adding to an active/active configuration [161](#)

adding to an active/active configuration with Multipath Storage [158](#)

comparison [137](#)

hot adding [159](#)

hot swapping modules in [171](#)

restrictions on adding [160](#)

upgrades supported [169](#)

upgrading modules [169](#)

disk-shelf-to-switch cabling, fabric-attached MetroClusters [75, 77, 85, 86](#)

documentation, required [42, 59](#)

## E

eOM management interface [116](#)

- eliminating single-point-of-failure (SPOF) 190
- emulated LANs
  - considerations for 117
- emulated node
  - accessing from the takeover node 139
  - accessing remotely 141
  - backing up 143
  - commands that are unavailable in 141
  - description of 139
  - dumps and restores 143
  - managing 139
- enabling takeover (cf) 132
- equipment racks
  - installation in 41
  - preparation of 45
- events triggering failover 192

## F

- fabric-attached MetroCluster configuration
  - adding disk shelves and loops 164
  - assigning disk pools 91
  - behavior of Data ONTAP with 36
  - local node
    - cabling controller to switch
      - with hardware-based disk ownership 72
      - with software-based disk ownership 73
    - cabling disk shelves to switch
      - with hardware-based disk ownership 75
      - with software-based disk ownership 77
  - remote node
    - cabling controller to switch
      - with hardware-based disk ownership 82
      - with software-based disk ownership 83
    - cabling disk shelves to switch
      - with hardware-based disk ownership 85
      - with software-based disk ownership 86
  - verifying disk paths
    - with hardware-based disk ownership 92
    - with software-based disk ownership 93
- fabric-attached MetroClusters
  - about 33
  - advantages of 34
  - Brocade switch configuration 70
  - cabling 68, 72, 73, 75, 77, 82, 83, 85, 86
  - illustration of 68
  - limitations 38
  - planning worksheet 69
  - restrictions 36
  - setup requirements for 36

- fabric-attached MetroClusters (*continued*)
  - upgrading from hardware-based to software-based disk ownership 65
  - variations 38
- fabric-attached MetroClusters configuration
  - cabling cluster interconnect for
    - cabling FC-VI adapter for
      - with hardware-based disk ownership 79, 89
      - with software-based disk ownership 80, 90
    - with hardware-based disk ownership 79, 89
    - with software-based disk ownership 80, 90
- failover 138, 189, 192
  - cause-and-effect table 192
  - determining status (cf status) 138
- failures that trigger failover 192
- FC-VI adapter, cabling 79, 80, 89, 90
- fencing, manual 176
- Fibre Channel ports
  - identifying for active/active configuration 46, 51
  - mirrored active/active configurations and 51
- Fibre Channel switches 61
- forcing
  - giveback 146
  - takeover 135

## G

- giveback
  - automatic 148
  - automatically terminating long-running processes 148
  - delay time for CIFS clients 147
  - description of 127
  - managing 144
  - normal 145
  - performing a 144
  - shortening 147
  - testing 123
  - troubleshooting 149

## H

- HA configuration checker 149
- halting system without takeover 133
- hardware
  - active/active components described 24
  - components described 24
  - installing a component 187
  - removing a component 186
  - single-point-of-failure 189
  - upgrading nondisruptively 185

hardware assisted takeover [110](#)  
 hardware-based disk ownership [51](#), [65](#), [72](#), [75](#), [82](#), [85](#)

## I

immediate takeover, enabling or disabling [132](#)  
 installation  
   equipment rack [41](#)  
   system cabinet [42](#)  
 installing hardware components [187](#)  
 interface configurations  
   dedicated [114](#)  
   shared [114](#)  
   standby [114](#)  
 interfaces  
   configuration for takeover [116](#)  
   configuring dedicated [105](#)  
   configuring for automatic takeover [137](#)  
   configuring shared [104](#)  
   configuring standby [105](#)  
   dedicated, diagram [116](#)  
   shared, diagram [115](#)  
   standby, diagram [116](#)  
   types and configurations [113](#), [116](#)  
 IPv6 considerations [113](#)

## L

licenses  
   enabling cluster [106](#)  
   required [106](#)  
 loop  
   adding to a fabric-attached MetroCluster [164](#)  
   adding to an active/active configuration [163](#)  
 lun commands, lun online [177](#)  
 LUNs, bringing online [177](#)

## M

mailbox disks [20](#)  
 managing in normal mode [127](#)  
 manual fencing [176](#)  
 MetroClusters  
   converting to from a standard or mirrored active/active configuration [63](#)  
   disaster recovery using [173](#)  
   LUNs and [177](#)  
   reestablishing configuration after disaster [179](#)  
   software-based disk ownership and [62](#)

mirrored active/active configuration  
   cabling Channel A [53](#)  
   cabling Channel B [54](#)  
   cabling cluster interconnect for [56](#)  
 mirrored active/active configurations  
   about [26](#)  
   advantages of [26](#)  
   restrictions [27](#)  
   setup requirements for [27](#)  
   variations [28](#)  
 modules, disk shelf  
   about [165](#)  
   best practices for changing types [166](#)  
   hot-swapping [171](#)  
   restrictions for changing types [165](#)  
   testing [166](#)  
   upgrading [169](#)  
 Multipath Storage  
   advantages of [153](#)  
   AT-FCX module versions supported [155](#)  
   best practices [153](#)  
   cabling [156](#)  
   connection types used by [152](#)  
   description of [152](#)  
   requirements [153](#)

## N

network interfaces  
   configuration for takeover [116](#)  
   configuring for takeover [117](#)  
   emulated LAN considerations [117](#)  
   types and configurations [113](#), [116](#)  
 nondisruptive upgrades, hardware [185](#)  
 normal mode  
   managing in [127](#)  
 NVRAM adapter [44](#), [61](#)

## O

options, setting [107](#)

## P

parameters  
   change\_fsid [108](#)  
   required to be identical between nodes [108](#)  
   setting [107](#)  
 partner command [139](#)

- partner name, displaying (cf partner) [131](#)
- planning worksheet for fabric-attached MetroClusters [69](#)
- pool assignments, fabric-attached MetroClusters [91](#)
- pool rules [51](#)
- port list
  - creating for mirrored active/active configurations [52](#)
- ports
  - identifying which ones to use [46](#), [51](#)
  - mirrored active/active configurations and [51](#)
- preparing equipment racks [45](#)
- primary connections, in Multipath Storage [152](#)

## R

- redundant connections, in Multipath Storage [152](#)
- reestablishing MetroCluster configuration [179](#)
- removing an active/active configuration [95](#)
- requirements
  - adapters [61](#)
  - disk shelves [160](#)
  - documentation [42](#), [59](#)
  - equipment [44](#), [61](#)
  - Fibre Channel switches [61](#)
  - for upgrading to a fabric-attached MetroCluster using software-based disk ownership [65](#)
  - hot-swapping a disk shelf module [171](#)
  - Multipath Storage [153](#)
  - NVRAM adapter [61](#)
  - SFP modules [61](#)
  - tools [43](#), [61](#)
- restrictions
  - fabric-attached MetroCluster [36](#)
  - in active/active configurations [24](#)
  - in mirrored active/active configurations [27](#)
  - in stretch MetroClusters [32](#)
- rsh, using to access node after takeover [126](#)

## S

- setting options and parameters [107](#)
- setup, running on active/active configurations [103](#)
- SFP modules [44](#), [61](#)
- shared interfaces
  - configuring using ifconfig [117](#)
  - configuring using setup [104](#)
  - described [114](#)
  - diagram [115](#)
- single-point-of-failure (SPOF), eliminating [190](#)
- single-point-of-failure, definition of [189](#)
- SNMP protocol and takeover mode [138](#)

- software-based disk management [91](#)
- software-based disk ownership [51](#), [62](#), [65](#), [73](#), [77](#), [83](#), [86](#)
- SPOF (single-point-of-failure) [189](#)
- stand-alone operation
  - changing a node in an active/active configuration to [95](#), [96](#), [97](#), [99](#)
- standard active/active configuration
  - cabling Channel A [48](#)
  - cabling Channel B [49](#)
  - cabling cluster interconnect for [50](#)
  - contents of [23](#)
  - variations [25](#)
- standard active/active configurations
  - behavior of Data ONTAP with [23](#)
- standby connections, in Multipath Storage [152](#)
- standby interfaces
  - configuring using ifconfig [117](#)
  - configuring using setup [105](#)
  - described [114](#)
  - diagram [116](#)
- status messages, descriptions of [130](#)
- status, monitoring active/active pair [127](#)
- stretch MetroClusters
  - about [28](#)
  - advantages of [29](#)
  - behavior of Data ONTAP with [31](#)
  - cabling [67](#)
  - connections required [30](#)
  - disk ownership and [31](#)
  - illustration of [29](#), [30](#)
  - on dual-controller systems [31](#)
  - restrictions [32](#)
  - variations [32](#)
- switch configuration, for fabric-attached MetroClusters [70](#)
- system cabinets
  - installation in [42](#)
  - preparing for cabling [45](#)

## T

- takeover
  - CIFS sessions and [126](#)
  - configuring VIFs for automatic [137](#)
  - configuring when it occurs [133](#)
  - configuring with dedicated and hot standby interfaces [116](#)
  - determining why one occurred [138](#)
  - disabling [132](#)
  - disabling immediate [132](#)
  - enabling [132](#)

takeover (*continued*)

- enabling immediate [132](#)
- forcing [135](#)
- forcing for disaster recovery [176](#)
- hardware assisted [110](#)
- reasons for [133](#)
- rsh access after [126](#)
- SNMP settings and [138](#)
- specifying delay before [136](#)
- statistics [138](#)
- telnet access after [126](#)
- testing [123](#)
- troubleshooting [149](#)
- using /etc/mcrc file at takeover [123](#)
- what happens after [126](#)
- what happens during [126](#)
- when it occurs [125](#)

takeover mode

- managing in [138](#)
- statistics in [138](#)

- telnet, using to access node after takeover [126](#)
- tools, required [43](#), [61](#)

## U

- unconfiguring an active/active pair [95](#)
- upgrading
  - disk shelf modules [169](#)
  - hardware, nondisruptively [185](#)
- UPS
  - using with active/active configurations [57](#)
  - using with MetroCluster configurations [93](#)
- utility, cf-config-check.cgi [149](#)

## V

VIFs

- configuring for automatic takeover [137](#)
- using to reduce SPOF [103](#)