



Data ONTAP® 8.2

High Availability and MetroCluster Configuration Guide

For 7-Mode

October 2020 | 215-08524_B0
doccomments@netapp.com

Updated for 8.2.1

 **NetApp**®

Contents

Understanding HA pairs.....	8
What an HA pair is.....	8
How HA pairs support nondisruptive operations and fault tolerance.....	8
How the HA pair improves fault tolerance.....	9
Standard HA pair diagram.....	12
Comparison of HA pair types.....	13
Understanding mirrored HA pairs	14
Advantages of mirrored HA pairs.....	14
Asymmetrically mirrored HA pairs	14
Understanding takeover and giveback.....	15
When takeovers occur.....	15
Failover event cause-and-effect table.....	16
How hardware-assisted takeover speeds up takeover.....	20
What happens during takeover.....	21
What happens during giveback.....	21
Background disk firmware update and takeover and giveback.....	22
Planning your HA pair configuration.....	23
Best practices for HA pairs.....	23
Setup requirements and restrictions for standard HA pairs.....	24
Setup requirements and restrictions for mirrored HA pairs	25
Multipath HA requirement.....	25
What multipath HA for HA pairs is.....	25
How the connection types are used.....	26
Advantages of multipath HA.....	26
Requirements for hardware-assisted takeover.....	27
Possible storage configurations in the HA pairs	27
HA pairs and storage system model types.....	27
Single-chassis and dual-chassis HA pairs.....	28
Interconnect cabling for systems with variable HA configurations.....	28
HA configuration and the HA state PROM value.....	28
Table of storage system models and HA configuration differences.....	29
Installing and cabling an HA pair.....	32
System cabinet or equipment rack installation.....	32
HA pairs in an equipment rack.....	32
HA pairs in a system cabinet.....	32
Required documentation.....	33
Required tools.....	33
Required equipment.....	34
Preparing your equipment.....	35

- Installing the nodes in equipment racks.....35
- Installing the nodes in a system cabinet..... 35
- Cabling a standard HA pair..... 35
 - Determining which Fibre Channel ports to use for Fibre Channel disk shelf connections..... 36
 - Cabling Node A to DS14mk2 or DS14mk4 disk shelves.....37
 - Cabling Node B to DS14mk2 or DS14mk4 disk shelves..... 38
 - Cabling the HA interconnect (all systems except 32xx)..... 40
 - Cabling the HA interconnect (32xx systems in separate chassis)..... 40
- Cabling a mirrored HA pair.....41
 - Determining which Fibre Channel ports to use for Fibre Channel disk shelf connections..... 41
 - Creating your port list for mirrored HA pairs.....42
 - Cabling the Channel A DS14mk2 or DS14mk4 disk shelf loops.....42
 - Cabling the Channel B DS14mk2 or DS14mk4 disk shelf loops.....44
 - Cabling the redundant multipath HA connection for each loop.....46
 - Cabling the HA interconnect (all systems except 32xx)..... 47
 - Cabling the HA interconnect (32xx systems in separate chassis)..... 48
- Required connections for using uninterruptible power supplies with standard or mirrored HA pairs.....48

Configuring an HA pair..... 50

- Bringing up the HA pair 50
 - Considerations for HA pair setup..... 50
 - Configuring shared interfaces during setup..... 51
 - Configuring dedicated interfaces with setup..... 51
 - Configuring standby interfaces with setup..... 52
- Verifying and setting the HA state on the controller modules and chassis.....52
- Enabling HA mode capability and controller failover.....53
- Setting options and parameters.....53
 - Option types for HA pairs53
 - Setting matching node options.....54
 - Parameters that must be the same on each node.....54
 - Best practices for cf options.....55
 - Disabling the change_fsid option in MetroCluster configurations.....56
 - Configuring hardware-assisted takeover.....57
- Configuring network interfaces for HA pairs 60
 - Understanding interfaces in an HA pair 60
 - Making nondisruptive changes to the interface groups..... 65
 - Configuring network interfaces for the HA pair.....66
 - Configuring a dedicated interface in an HA pair.....66
 - Configuring a partner interface in an HA pair 67
 - Configuring partner addresses on different subnets (MetroCluster configurations only)..... 68
- Verifying the HA pair configuration.....71
- Testing takeover and giveback.....71

Managing takeover and giveback.....73

- Monitoring an HA pair in normal mode..... 73
 - Monitoring HA pair status..... 73
 - Description of HA pair status messages..... 73
 - Monitoring the hardware-assisted takeover feature..... 74
 - Displaying the partner's name.....76
 - Displaying disk and array LUN information on an HA pair 77
- What takeover and giveback are.....77
 - When takeovers occur.....77

- What happens during takeover..... 78
- What happens after takeover.....78
- What happens during giveback..... 78
- Configuring automatic takeover..... 79
 - Reasons that automatic takeover occurs..... 79
 - Commands for performing a manual takeover..... 81
 - Halting a node without takeover..... 82
 - Rebooting a node without takeover..... 82
 - Enabling and disabling takeover..... 82
 - Enabling and disabling takeover on reboot..... 82
 - Enabling and disabling automatic takeover of a panicked partner..... 83
 - Specifying the time period before takeover..... 83
 - Enabling or disabling negotiated failover for a network interface..... 84
 - Takeover of vFiler units and the vFiler limit..... 84
- Managing an HA pair in takeover mode..... 84
 - Determining why takeover occurred..... 84
 - Statistics in takeover mode..... 85
- Managing emulated nodes..... 85
 - Management exceptions for emulated nodes..... 85
 - Accessing the emulated node from the takeover node..... 85
 - Using Remote Shell..... 87
 - Emulated node command exceptions..... 87
- Performing dump and restore operations for a failed node..... 88
- Giveback operations..... 89
 - About manual giveback..... 89
 - Configuring giveback..... 92
 - Configuring automatic giveback..... 92

Reconfiguring an HA pair into two stand-alone systems..... 95

- Ensuring uniform disk ownership within disk shelves and loops in the system.....95
- Configuring a node for non-HA (stand-alone) use..... 95
- Reconfiguring nodes using disk shelves for stand-alone operation..... 97
- Requirements when changing a node using array LUNs to stand-alone..... 99
- Reconfiguring nodes using array LUNs for stand-alone operation..... 99

Managing DS14mk2 or DS14mk4 disk shelves in an HA pair..... 102

- Adding DS14mk2 or DS14mk4 disk shelves to a multipath HA loop..... 102
- Upgrading or replacing modules in an HA pair 103
- About the disk shelf modules..... 103
- Restrictions for changing module types..... 104
- Best practices for changing module types..... 104
- Testing the modules..... 104
- Determining path status for your HA pair 104
- Hot-swapping a module..... 106

Nondisruptive operations with HA pairs..... 108

- Where to find procedures for nondisruptive operations with HA pairs..... 108
- Hot-removing disk shelves in systems running Data ONTAP 8.2.1 or later 108

Overview of a MetroCluster configuration..... 114

Types of MetroCluster configurations.....	114
How mirroring works in a MetroCluster configuration.....	114
How Data ONTAP works with MetroCluster configurations.....	114
Advantages of stretch MetroCluster configurations.....	115
Configuration variations for stretch MetroCluster configurations.....	115
Advantages of fabric-attached MetroCluster configurations.....	116
Configuration variations for fabric-attached MetroCluster configurations.....	116

Implementing stretch MetroCluster configurations with disks..... 118

Planning a stretch MetroCluster configuration using disks.....	118
Setup requirements and restrictions for stretch MetroCluster configurations with disks.....	118
Hardware requirements for using disks in a MetroCluster configuration with V-Series systems.....	119
Required connections for using uninterruptible power supplies with MetroCluster configurations.....	120
Stretch MetroCluster configuration with disks	121
Required documentation, tools, and equipment.....	122
Cabling a stretch MetroCluster configuration with disks.....	125
Stretch MetroCluster configuration on single-enclosure HA pairs	126
Assigning disk pools in a stretch MetroCluster configuration.....	126
Managing the default configuration speed of a stretch MetroCluster configuration.....	128
Changing the default configuration speed of a stretch MetroCluster configuration.....	128
Resetting a stretch MetroCluster configuration to the default speed.....	131

Implementing fabric-attached MetroCluster configurations with disks..... 133

Planning a fabric-attached MetroCluster configuration using disks.....	133
Setup requirements and restrictions for fabric-attached MetroCluster configurations with disks.....	133
Configuration limitations for fabric-attached MetroCluster configurations with disks.....	135
Hardware requirements for using disks in a MetroCluster configuration with V-Series systems.....	135
Required connections for using uninterruptible power supplies with MetroCluster configurations.....	136
Requirements for a shared-switches configuration.....	136
Fabric-attached MetroCluster configuration with disks.....	137
How fabric-attached MetroCluster configurations use Brocade and Cisco Fibre Channel switches.....	137
Planning the fabric-attached MetroCluster installation.....	138
Required documentation, tools, and equipment.....	139
Converting an HA pair to a fabric-attached MetroCluster configuration.....	142
Cabling a fabric-attached MetroCluster configuration with disks.....	143
Configuring the switches.....	145
Cabling Node A.....	145
Cabling Node B.....	149
Assigning disk pools.....	152
Verifying disk paths.....	153
Setting up a shared-switches configuration	153
Fabric-attached MetroCluster configuration on single-enclosure HA pairs.....	155
Configuration differences for fabric-attached MetroCluster configurations on single-enclosure HA pairs ..	156
Setting a preferred primary port in a MetroCluster configuration.....	157
Removing the preferred primary port in a fabric-attached MetroCluster configuration.....	157
Enabling unidirectional port mirroring in a MetroCluster configuration.....	157
What unidirectional port mirroring is.....	158
Verifying the MetroCluster configuration.....	159
Gathering the information for verifying the MetroCluster configuration.....	159
Testing takeover and giveback.....	160
Testing MetroCluster configuration during a switch failure.....	160
Testing a MetroCluster configuration during FibreBridge 6500N bridge failure.....	161

Testing a MetroCluster configuration during cluster failover on disaster..... 161

Implementing a stretch or fabric-attached MetroCluster configuration with array

LUNs.....162

- Planning for a MetroCluster configuration with array LUNs..... 162
 - Types of systems that can use array LUNs in a MetroCluster configuration..... 163
 - Implementation overview for a MetroCluster configuration with array LUNs..... 163
 - Requirements for a MetroCluster configuration with array LUNs..... 164
 - Requirements for a shared-switches MetroCluster configuration with array LUNs..... 166
 - Recommended fabric-attached MetroCluster configuration with array LUNs..... 166
 - Recommended stretch MetroCluster configuration with array LUNs..... 169
 - Cabling guidelines for a MetroCluster configuration with array LUNs..... 172
 - Planning zoning for a MetroCluster configuration with array LUNs..... 173
- Connecting devices in a MetroCluster configuration with array LUNs..... 175
 - Connecting the local systems in a MetroCluster configuration with array LUNs..... 176
 - Connecting the remote systems in a MetroCluster configuration with array LUNs..... 179
 - Connecting the switch fabric in a MetroCluster configuration with array LUNs..... 182
 - Connecting the fabric and storage array in a MetroCluster configuration with array LUNs..... 183
 - Configuring zoning in a MetroCluster configuration with array LUNs..... 185
 - Setting up a shared-switches configuration 186
 - Setting a preferred primary port in a MetroCluster configuration..... 189
 - Removing the preferred primary port in a fabric-attached MetroCluster configuration..... 189
 - Enabling unidirectional port mirroring in a MetroCluster configuration..... 189
 - Changing the configuration speed of a stretch MetroCluster configuration..... 190
- Setting up Data ONTAP after connecting devices in a MetroCluster configuration with array LUNs..... 191
- Testing a MetroCluster configuration with array LUNs..... 191
 - Testing zoning of FC-VI ports in a MetroCluster configuration with array LUNs..... 191
 - Verifying proper setup at MetroCluster sites with storage arrays..... 192
 - Simulating a disaster recovery in a MetroCluster configuration with array LUNs..... 193

Implementing a MetroCluster configuration with both disks and array LUNs.... 195

- Types of systems that can use array LUNs in a MetroCluster configuration..... 195
- Planning a MetroCluster configuration with disks and array LUNs..... 195
 - Guidelines for implementing a MetroCluster configuration with disks and array LUNs..... 195
 - Supported fabric-attached MetroCluster configuration with disks and array LUNs..... 197
 - Supported stretch MetroCluster configuration with disks and array LUNs..... 199
- Testing a MetroCluster configuration with disks and array LUNs..... 201

Recovering from a disaster by using MetroCluster configurations..... 202

- Conditions that constitute a disaster..... 202
 - Ways to determine whether a disaster occurred..... 202
 - Failures that do not require disaster recovery..... 202
- Recovering from a disaster..... 203
 - Restricting access to the disaster site node..... 204
 - Forcing a node into takeover mode..... 204
 - Remounting volumes of the failed node..... 205
 - Recovering LUNs of the failed node..... 205
 - Fixing failures caused by the disaster..... 206
 - Reestablishing the MetroCluster configuration..... 206

Feature update record..... 211

Copyright.....216

Trademark.....217

How to send comments about documentation and receive update notifications...218

Index.....219

Understanding HA pairs

HA pairs provide hardware redundancy that is required for nondisruptive operations and fault tolerance and give each node in the pair the software functionality to *take over* its partner's storage and subsequently *give back* the storage.

What an HA pair is

An HA pair 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. In this configuration, one node can take over its partner's storage to provide continued data service if the partner goes down.

You can configure the HA pair 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.

The controllers are connected to each other through an HA 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). The interconnect is internal and requires no external cabling if both controllers are in the same chassis.

Takeover is the process in which a node takes over the storage of its partner. *Giveback* is the process in which that storage is returned to the partner. Both processes can be initiated manually or configured for automatic initiation.

How HA pairs support nondisruptive operations and fault tolerance

HA pairs provide fault tolerance and let you perform nondisruptive operations, including hardware and software upgrades, and hardware maintenance.

- Fault tolerance
When one node fails or becomes impaired and a takeover occurs, the partner node continues to serve the failed node's data.
- Nondisruptive software upgrades or hardware maintenance
During hardware maintenance or upgrades, when you halt one node and a takeover occurs (automatically, unless you specify otherwise), the partner node continues to serve data for the halted node while you upgrade or perform maintenance on the node you halted.
During nondisruptive upgrades of Data ONTAP, the user manually enters the `cf takeover` command to take over the partner node to allow the software upgrade to occur. The takeover node continues to serve data for both nodes during this operation.
For more information about nondisruptive software upgrades, see the *Data ONTAP Upgrade and Revert/Downgrade Guide for 7-Mode*.

The HA pair supplies nondisruptive operation and fault tolerance due to the following aspects of its configuration:

- The controllers in the HA pair are connected to each other either through an HA interconnect consisting of adapters and cables, or, in systems with two controllers in the same chassis, through an internal interconnect
The nodes use the interconnect to perform the following tasks:
 - Continually check if the other node is functioning
 - Mirror log data for each other's NVRAM or NVMEM
- The nodes use two or more disk shelf loops, or storage arrays, in which the following conditions apply:

- Each node manages its own disks or array LUNs
- In case of takeover, the surviving node provides read/write access to the partner's disks or array LUNs until the failed node becomes available again

Note: Disk ownership is established by Data ONTAP or the administrator; it is not based on which disk shelf the disk is attached to.

For more information about disk ownership, see the *Data ONTAP Storage Management Guide for 7-Mode*.

- 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 that perform 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

Related concepts

[Nondisruptive operations with HA pairs](#) on page 108

By taking advantage of an HA pair's takeover and giveback operations, you can change hardware components and perform software upgrades in your configuration without disrupting access to the system's storage.

[Where to find procedures for nondisruptive operations with HA pairs](#) on page 108

An HA pair allows you to perform an extensive range of nondisruptive system maintenance and upgrade operations. You can refer to the specific documents for the required procedures.

[Managing DS14mk2 or DS14mk4 disk shelves in an HA pair](#) on page 102

You must follow specific procedures to add disk shelves to an HA pair or a MetroCluster configuration, or to upgrade or replace disk shelf hardware in an HA pair.

How the HA pair improves fault tolerance

A storage system has a variety of single points of failure, such as certain cables or hardware components. An HA pair greatly reduces the number of single points of failure because if a failure occurs, the partner can take over and continue serving data for the affected system until the failure is fixed.

Single point of failure definition

A single point of failure represents the failure of a single hardware component that can lead to loss of data access or potential loss of data.

Single point of failure 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 such as a controller or adapter fails, you can use the controller failover function to provide continuous data availability and preserve data integrity for client applications and users.

Single point of failure analysis for HA pairs

Different individual hardware components and cables in the storage system are single points of failure, but an HA configuration can eliminate these points to improve data availability.

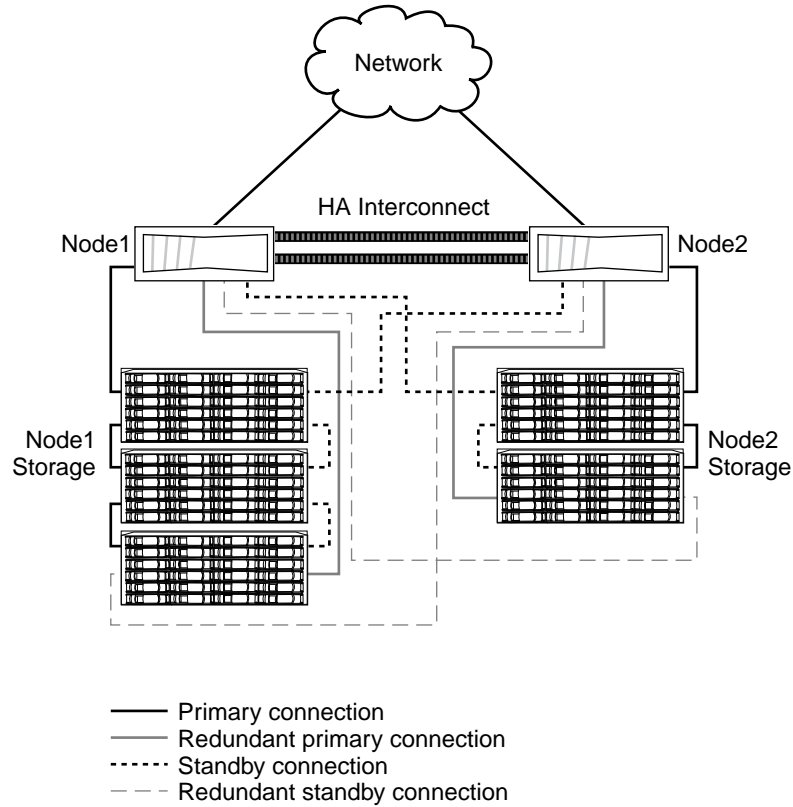
Hardware components	Single point of failure		How storage failover eliminates single point of failure
	Stand-alone	HA pair	
Controller	Yes	No	If a controller fails, the node automatically fails over to its partner node. The partner (takeover) node serves data for both of the nodes.
NVRAM	Yes	No	If an 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 automatically fails over to its partner node. The partner (takeover) node serves data for both of the nodes.
Multiple NICs with interface groups (virtual interfaces)	Maybe, if all NICs fail	No	If one of the networking links within an interface group 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 an interface group fail, the node automatically fails over to its partner node if failover is enabled for the interface group.
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 or SAS HBA	Yes	No	If an FC-AL adapter for the primary loop fails for a configuration without multipath HA, the partner node attempts a takeover at the time of failure. With multipath HA, no takeover is required. If the FC-AL adapter for the secondary loop fails for a configuration without multipath HA, the failover capability is disabled, but both nodes continue to serve data to their respective applications and users, with no impact or delay. With multipath HA, failover capability is not affected.
FC-AL or SAS cable (controller-to-shelf, shelf-to-shelf)	No, if dual-path cabling is used	No	If an FC-AL loop or SAS stack breaks in a configuration that does not have multipath HA, 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 multipath HA is used, no failover is required.
Disk shelf module	No, if dual-path cabling is used	No	If a disk shelf module fails in a configuration that does not have multipath HA, 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 multipath HA is used, there is no impact.

Hardware components	Single point of failure		How storage failover eliminates single point of failure
	Stand-alone	HA pair	
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.
Power supply	Maybe, if both power supplies fail	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)	Maybe, if both fans fail	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.
HA interconnect adapter	Not applicable	No	If an HA interconnect adapter fails, the failover capability is disabled but both nodes continue to serve data to their respective applications and users.
HA interconnect cable	Not applicable	No	The HA interconnect adapter supports dual HA 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.

Standard HA pair diagram

Each node in a standard HA pair requires a network connection, an HA interconnect between the controllers, and connections both to its own disk shelves as well as its partner node's shelves.

Standard HA pair



This diagram shows a standard HA pair with native disk shelves and multipath HA.

This diagram shows DS4243 disk shelves. For more information about cabling SAS disk shelves, see the *Universal SAS and ACP Cabling Guide* on the NetApp Support Site.

Comparison of HA pair types

The different types of HA pairs support different capabilities for data duplication, distance between nodes, and failover.

HA pair type	Data duplication?	Distance between nodes	Failover possible after loss of entire node (including storage)?	Notes
Standard HA pair	No	Up to 500 meters Note: SAS configurations are limited to 5 meters between nodes when using SAS copper cables	No	Use this configuration to provide higher availability by protecting against many hardware single-points-of-failure.
Mirrored HA pair	Yes	Up to 500 meters Note: SAS configurations are limited to 5 meters between nodes when using SAS copper cables	No	Use this configuration to add increased data protection to the benefits of a standard HA pair.
Stretch MetroCluster	Yes	Up to 500 meters (270 meters if operating at 4 Gbps) Note: SAS configurations are limited to 20 meters between nodes when using SAS copper cables to connect disk shelves to controllers directly.	Yes	Use this configuration to provide data and hardware duplication to protect against a local disaster (for example, a power outage to one node).
Fabric-attached MetroCluster	Yes	Up to 200 kilometers, depending on switch configuration. See the Interoperability Matrix at support.netapp.com/matrix .	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.

Understanding mirrored HA pairs

Mirrored HA pairs provide high availability through failover, just as standard HA pairs do. Additionally, mirrored HA pairs 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 HA pairs do not provide the capability to fail over to the partner node if one node fails disastrously or is disabled. For example, if an entire node loses power, including its storage, you cannot fail over to the partner node. For this capability, use a MetroCluster configuration.

Mirrored HA pairs use SyncMirror. For more information about SyncMirror, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

Advantages of mirrored HA pairs

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 that 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 storage array

The failure of an FC-AL adapter, SAS HBA, disk shelf loop or stack, or disk shelf module does not require a failover in a mirrored HA pair.

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

Asymmetrically mirrored HA pairs

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 HA pair .

Understanding takeover and giveback

Takeover and giveback are the operations that let you take advantage of the HA configuration to perform nondisruptive operations and avoid service interruptions. 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.

When takeovers occur

Takeovers can be initiated manually or occur automatically when a failover event happens, depending on how you configure the HA pair. In some cases, takeovers occur automatically, regardless of configuration.

Takeovers can occur under the following conditions:

- A takeover is manually initiated.
- A node is in an HA pair with the `default` configuration for immediate takeover on panic, and that node undergoes a software or system failure that leads to a panic.

Note: The default setting for the `cf.takeover.on_panic` option is `on`. If the `cf.takeover.on_panic` option is set to `on`, a node panic will cause an automatic takeover. If the `cf.takeover.on_panic` option is set to `off`, a node panic will not cause an automatic takeover. You should not turn this option to `off` unless you are instructed by technical support to do so.

By default, after the partner has recovered from the panic and booted up, the node automatically performs a giveback, returning the partner to normal operation.

- A node that is in an HA pair 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 MetroCluster configurations, you can initiate a forced takeover in this situation.

- One or more network interfaces that are configured to support failover become unavailable.
- A node does not receive heartbeat messages from its partner.
This could happen if the partner 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` parameter.
- You reboot one of the nodes without using the `-f` parameter.
- Hardware-assisted takeover is enabled and triggers a takeover when the remote management device (RLM or Service Processor) detects failure of the partner node.

Failover event cause-and-effect table

Failover events cause a controller failover in HA pairs. The storage system responds differently depending on the event and the type of HA pair.

Cause-and-effect table for standard or mirrored HA pairs

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?	
			Single storage system	Standard or mirrored HA pair
Single disk failure	No	No	Yes	Yes
Double disk failure (2 disks fail in same RAID group)	Yes, unless you are using SyncMirror or RAID-DP, then no	Maybe. If root volume has double disk failure, or if the mailbox disks are affected, no failover is possible	No, unless you are using RAID-DP or SyncMirror, then yes	No, unless you are using RAID-DP or SyncMirror, then yes
Triple disk failure (3 disks fail in same RAID group)	Maybe. If SyncMirror is being used, no takeover occurs; otherwise, yes	Maybe. If root volume has triple disk failure, no failover is possible	No	No
Single HBA (initiator) failure, Loop A	Maybe. If SyncMirror or multipath HA is in use, then no; otherwise, yes	Maybe. If root volume has double disk failure, no failover is possible	Yes, if multipath HA or SyncMirror is being used	Yes, if multipath HA or SyncMirror is being used, or if failover succeeds
Single HBA (initiator) failure, Loop B	No	Yes, unless you are using SyncMirror or multipath HA and the mailbox disks are not affected, then no	Yes, if multipath HA or SyncMirror is being used	Yes, if multipath HA or SyncMirror is being used, or if failover succeeds
Single HBA initiator failure (both loops at the same time)	Yes, unless the data is mirrored on a different (up) loop or multipath HA is in use, then no takeover needed	Maybe. If the data is mirrored or multipath HA is being used and the mailbox disks are not affected, then no; otherwise, yes	No, unless the data is mirrored or multipath HA is in use, then yes	No failover needed if data is mirrored or multipath HA is in use

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?	
			Single storage system	Standard or mirrored HA pair
AT-FCX failure (Loop A)	Only if multidisk volume failure or open loop condition occurs, and neither SyncMirror nor multipath HA is in use	Maybe. If root volume has double disk failure, no failover is possible	No	Yes, if failover succeeds
AT-FCX failure (Loop B)	No	Maybe. If SyncMirror or multipath HA is in use, then no; otherwise, yes	Yes, if multipath HA or SyncMirror is in use	Yes
IOM failure (Loop A)	Only if multidisk volume failure or open loop condition occurs, and neither SyncMirror nor multipath HA is in use	Maybe. If root volume has double disk failure, no failover is possible	No	Yes, if failover succeeds
IOM failure (Loop B)	No	Maybe. If SyncMirror or multipath HA is in use, then no; otherwise, yes	Yes, if multipath HA or SyncMirror is in use	Yes
Shelf (backplane) failure	Only if multidisk volume failure or open loop condition occurs, and data isn't mirrored	Maybe. If root volume has double disk failure or if the mailboxes are affected, no failover is possible	Maybe. If data is mirrored, then yes; otherwise, no	Maybe. If data is mirrored, then yes; otherwise, no
Shelf, single power failure	No	No	Yes	Yes
Shelf, dual power failure	Only if multidisk volume failure or open loop condition occurs and data is not mirrored	Maybe. If root volume has double disk failure, or if the mailbox disks are affected, no failover is possible	Maybe. If data is mirrored, then yes; otherwise, no	Maybe. If data is mirrored, then yes; otherwise, no
Controller, single power failure	No	No	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?	
			Single storage system	Standard or mirrored HA pair
Controller, dual power failure	Yes	Yes, until power is restored	No	Yes, if failover succeeds
HA interconnect failure (1 port)	No	No	Not applicable	Yes
HA interconnect failure (both ports)	No	Yes	Not applicable	Yes
Ethernet interface failure (primary, no interface group)	Yes, if set up to do so	No	Yes	Yes
Ethernet interface failure (primary, interface group)	Yes, if set up to do so	No	Yes	Yes
Ethernet interface failure (secondary, interface group)	Yes, if set up to do so	No	Yes	Yes
Ethernet interface failure (interface group, all ports)	Yes, if set up to do so	No	Yes	Yes
Tape interface failure	No	No	Yes	Yes
Heat exceeds permissible amount	Yes	No	No	No
Fan failures (disk shelves or controller)	No	No	Yes	Yes
Reboot	Yes	No	No	Yes, if failover occurs
Panic	Yes	No	No	Yes, if failover occurs

Cause-and-effect table for stretch and fabric-attached MetroClusters

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?	
			Stretch MetroCluster	Fabric Attached MetroCluster
Single disk failure	No	No	Yes	Yes
Double disk failure (2 disks fail in same RAID group)	No	No	Yes	Yes, with no failover necessary
Triple disk failure (3 disks fail in same RAID group)	No	No	Yes, with no failover necessary	Yes, with no failover necessary
Single HBA (initiator) failure, Loop A	No	No	Yes	Yes, with no failover necessary
Single HBA (initiator) failure, Loop B	No	No	Yes	Yes, with no failover necessary
Single HBA initiator failure, (both loops at the same time)	No	No	Yes, with no failover necessary	Yes, with no failover necessary
AT-FCX failure (Loop A)	No	No	Yes	Not applicable
AT-FCX failure (Loop B)	No	No	Yes	Not applicable
ESH4 or IOM Failure (loop A)	No	No	Yes	Yes
ESH4 or IOM Failure (loop B)	No	No	Yes	Yes
Shelf (backplane) failure	No	No	Yes	Yes, with no failover necessary
Shelf, single power failure	No	No	Yes	Yes
Shelf, dual power failure	No	No	Yes	Yes, with no failover necessary
Controller, single power failure	No	No	Yes	Yes
Controller, dual power failure	Yes	Yes, until power is restored	Yes, if failover succeeds	Yes, if failover succeeds

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?	
			Stretch MetroCluster	Fabric Attached MetroCluster
HA interconnect failure (1 port)	No	No	Yes	Yes
HA interconnect failure (both ports)	No	No; failover is possible	Yes	Yes
Ethernet interface failure (primary, no interface group)	Yes, if set up to do so	No	Yes	Yes
Ethernet interface failure (primary, interface group)	Yes, if set up to do so	No	Yes	Yes
Ethernet interface failure (secondary, interface group)	Yes, if set up to do so	No	Yes	Yes
Ethernet interface failure (interface group, all ports)	Yes, if set up to do so	No	Yes	Yes
Tape interface failure	No	No	Yes	Yes
Heat exceeds permissible amount	Yes	No	Yes, if failover occurs	Yes, if failover occurs
Fan failures (disk shelves or controller)	No	No	Yes	Yes
Reboot	Yes	No	Maybe. Depends on cause of reboot	Maybe. Depends on cause of reboot
Panic	Yes	No	Maybe. Depends on cause of panic	Maybe. Depends on cause of panic

How hardware-assisted takeover speeds up takeover

Hardware-assisted takeover speeds up the takeover process by using a node's remote management device (SP or RLM) to detect failures and quickly initiate the takeover rather than waiting for Data ONTAP to recognize that the partner's heartbeat has stopped.

Without hardware-assisted takeover, if a failure occurs, the partner waits until it notices that the node is no longer giving a heartbeat, confirms the loss of heartbeat, and then initiates the takeover.

The hardware-assisted takeover feature uses the following process to take advantage of the remote management device and avoid that wait:

1. The remote management device monitors the local system for certain types of failures.
2. If a failure is detected, the remote management device immediately sends an alert to the partner node.
3. Upon receiving the alert, the partner initiates takeover.

The hardware-assisted takeover option (`cf.hw_assist.enable`) is enabled by default.

Related tasks

[Disabling and enabling the hardware-assisted takeover option](#) on page 58

Hardware-assisted takeover is enabled by default on systems that use remote management. Hardware-assisted takeover speeds the takeover process by using the RLM or SP to quickly detect potential takeover events and alerting the partner node.

[Setting the partner address for hardware-assisted takeover](#) on page 58

By default, on systems with an e0M port, the hardware-assisted takeover feature uses the IP address of the partner's e0M port to communicate with the partner. On systems without an e0M port, the system automatically selects another configured IP address. You can use the `cf.hw_assist.partner.address` option to select a different IP address.

[Setting the partner port for hardware-assisted takeover](#) on page 59

When hardware-assisted takeover is enabled, the remote management (either RLM or SP) 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.

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.

If the node doing the takeover panics

If the node that is performing the takeover panics within 60 seconds of initiating takeover, the following events occur:

- The node that panicked reboots.
- After it reboots, the node performs self-recovery operations and is no longer in takeover mode.
- Failover is disabled.

What happens during giveback

The local node returns ownership of the aggregates and volumes to the partner node after any issues on the partner node are resolved or maintenance is complete. In addition, the local node returns ownership when the partner node has booted up and giveback is initiated either manually or automatically.

When the failed node is functioning again, the following events can occur:

- You issue a `cf giveback` command that terminates the emulated node on the partner.
- The failed node resumes normal operation, serving its own data.

- The HA pair resumes normal operation, with each node ready to take over for its partner if the partner fails.

Background disk firmware update and takeover and giveback

Background disk firmware updates affect HA pair takeover and giveback operations differently, depending on how those operations are initiated.

The following list describes how background disk firmware update affects takeover and giveback:

- If a background disk firmware update is occurring on a disk on the source (or takeover) node, takeover operations are delayed for less than one minute until the disk firmware update completes on that disk.
- If a background disk firmware update is occurring on a disk on a node and it panics, takeover of the panicked node begins immediately.

Planning your HA pair configuration

As you plan your HA pair, you must consider recommended best practices, the requirements, and the possible variations.

Best practices for HA pairs

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

- Make sure that each power supply unit in the storage system is on a different power grid so that a single power outage does not affect all power supply units.
- Use interface groups (virtual interfaces) to provide redundancy and improve availability of network communication.
- Follow the documented procedures in the *Data ONTAP Upgrade and Revert/Downgrade Guide for 7-Mode* when upgrading your HA pair.
- 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 Config Advisor tool to help ensure that failovers are successful.
- If your system supports remote management (through an RLM or Service Processor), make sure that you configure it properly, as described in the *Data ONTAP System Administration Guide for 7-Mode*.
- Follow recommended limits for FlexVol volumes, dense volumes, Snapshot copies, and LUNs to reduce the takeover or giveback time.

When adding traditional or FlexVol volumes to an HA pair, consider testing the takeover and giveback times to ensure that they fall within your requirements.

- For systems using disks, check for failed disks regularly and remove them as soon as possible, as described in the *Data ONTAP Storage Management Guide for 7-Mode*. Failed disks can extend the duration of takeover operations or prevent giveback operations.
- Multipath HA is required on all HA pairs except for some FAS22xx system configurations, which use single-path HA and lack the redundant standby connections.
- To ensure that you receive prompt notification if takeover capability becomes disabled, configure your system for automatic email notification for the `takeover impossible` EMS messages:
 - `ha.takeoverImpVersion`
 - `ha.takeoverImpLowMem`
 - `ha.takeoverImpDegraded`
 - `ha.takeoverImpUnsync`
 - `ha.takeoverImpIC`
 - `ha.takeoverImpHotShelf`
 - `ha.takeoverImpNotDef`
- Set the `cf.giveback.auto.cancel.on_network_failure` option to `true` if network-related failovers are enabled.

Related tasks

[Verifying the HA pair configuration](#) on page 71

You can go to the NetApp Support Site and download the Config Advisor tool to check for common configuration errors.

Setup requirements and restrictions for standard HA pairs

You must follow certain requirements and restrictions when setting up a new standard HA pair. These requirements help you ensure the data availability benefits of the HA pair design.

The following list specifies the requirements and restrictions you should be aware of when setting up a new standard HA pair:

- Architecture compatibility
Both nodes must have the same system model and be running the same Data ONTAP software and system firmware versions. See the *Data ONTAP Release Notes for 7-Mode* for the list of supported systems.
- Nonvolatile memory (NVRAM or NVMEM) size and version compatibility
The size and version of the system's nonvolatile memory must be identical on both nodes in an HA pair.
- Storage capacity
The number of disks or array LUNs must not exceed the maximum configuration capacity. If your system uses both native disks and array LUNs, 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, array LUNs, or both, see the *Hardware Universe* at hwu.netapp.com.
Note: After a failover, the takeover node temporarily serves data from all the storage in the HA pair.
- Disks and disk shelf compatibility
 - FC, SATA, and SAS storage are supported in standard HA pairs.
 - FC disks cannot be mixed on the same loop as SATA or SAS disks.
 - Different connection types cannot be combined in the same loop or stack.
 - Different types of storage can be used on separate stacks or loops on the same node. You can also dedicate a node to one type of storage and the partner node to a different type, if needed.
See the *Data ONTAP Storage Management Guide for 7-Mode* and the *Hardware Universe* for information on storage configuration options.
 - Multipath HA is required on all HA pairs except for some FAS22xx system configurations, which use single-path HA and lack the redundant standby connections.
- Mailbox disks or array LUNs on the root volume
 - Two disks are required if the root volume is on a disk shelf.
 - One array LUN is required if the root volume is on a storage array.

The mailbox disks and LUNs are used for the following tasks:

- Maintaining consistency between the nodes in the HA pair
- Continually checking whether the other node is running or whether it has performed a takeover
- Storing configuration information that is not specific to any particular node
- HA 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) or Network File System (NFS), 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 HA pair loses functionality after a takeover.
- For an HA pair using array LUNs, both nodes in the pair must be able to detect the same array LUNs.
However, only the node that is the configured owner of a LUN has read-and-write access to that LUN. During takeover operations, the emulated storage system maintains read-and-write access to the LUN.

Setup requirements and restrictions for mirrored HA pairs

The restrictions and requirements for mirrored HA pairs include those for a standard HA pair 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.
 - There must be sufficient spares in each pool to account for a disk or array LUN failure.
 - Both plexes of a mirror should not reside on the same disk shelf because it might result in a single point of failure.

See the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode* for more information about requirements for setting up SyncMirror with array LUNs.

- The `cf.mode` option must be set to `ha`.
- If you are using array LUNs, paths to an array LUN must be redundant.

Related concepts

[Setup requirements and restrictions for standard HA pairs](#) on page 24

You must follow certain requirements and restrictions when setting up a new standard HA pair. These requirements help you ensure the data availability benefits of the HA pair design.

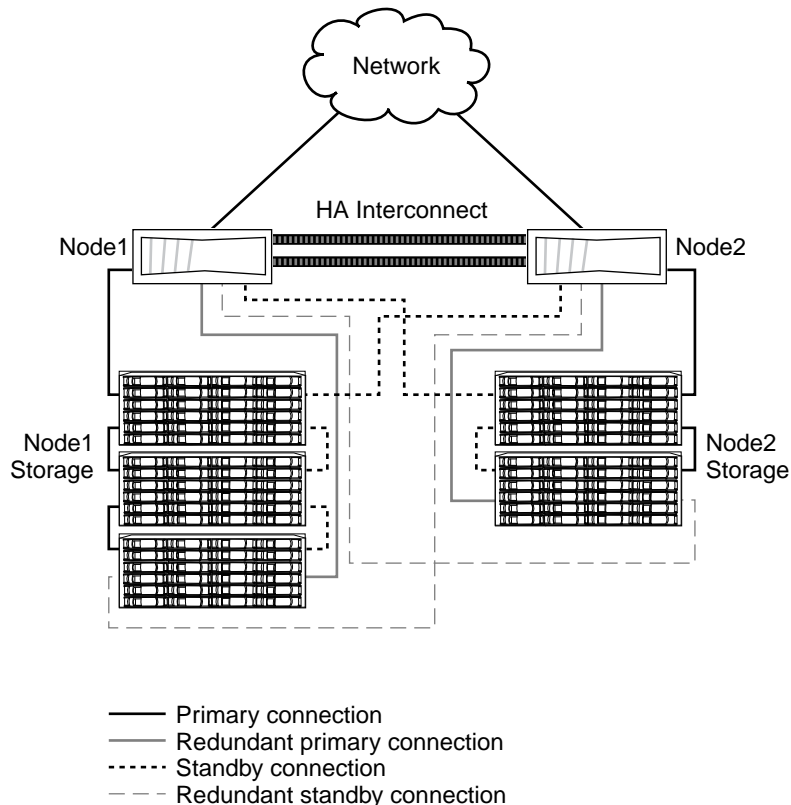
Multipath HA requirement

Multipath HA is required on all HA pairs except for some FAS22xx system configurations, which use single-path HA and lack the redundant standby connections. Multipath HA was previously referred to as *Multipath Storage*.

What multipath HA for HA pairs is

Multipath HA 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 HA pair without multipath HA has only one path from each controller to every disk, but an HA pair with multipath HA 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 HA pair using multipath HA. The redundant primary connections and the redundant standby connections are the additional connections required for multipath HA for HA pairs .



How the connection types are used

A multipath HA configuration uses primary, redundant and standby connections to ensure continued service in the event of the failure of an individual connection.

The following table outlines the connection types used for multipath HA for HA pairs, 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 HA

Multipath connections in an HA pair reduce single-points-of-failure.

By providing two paths from each controller to every disk shelf, multipath HA 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 HA adds value to a stretch MetroCluster environment, it is not necessary in a fabric MetroCluster configuration since multiple paths already exist.

Requirements for hardware-assisted takeover

The hardware-assisted takeover feature is available only on systems with an RLM or SP module configured for remote management. Remote management provides remote platform management capabilities, including remote access, monitoring, troubleshooting, logging, and alerting features.

Although a system with remote management on both nodes provides hardware-assisted takeover for both, hardware-assisted takeover is also supported on HA pairs in which only one of the two systems has remote management configured. Remote management does not have to be configured on both nodes in the HA pair. Remote management can detect failures on the system in which it is installed and provide faster takeover times if a failure occurs on the system with remote management.

See the *Data ONTAP 7-Mode System Administration Guide* for information about setting up remote management.

Possible storage configurations in the HA pairs

HA pairs can be configured symmetrically, asymmetrically, as an active/passive pair, or with shared disk shelf stacks.

Symmetrical configurations

In a symmetrical HA pair, each node has the same amount of storage.

Asymmetrical configurations

In an asymmetrical standard HA pair, 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 in addition to serving 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

In this configuration, shared loops or stacks between the nodes are particularly useful for active/passive configurations, as described in the preceding bullet.

HA pairs and storage system model types

Different model storage systems support some different HA configurations. This includes the physical configuration of the HA pair and the manner in which the system recognizes that it is in an HA pair.

Note: The physical configuration of the HA pair does not affect the cluster cabling of the nodes in the HA pair.

Single-chassis and dual-chassis HA pairs


Depending on the model of the storage system, an HA pair can consist of two controllers in a single chassis, or two controllers in two separate chassis. Some models can be configured either way, while other models can be configured only as a single-chassis HA pair or dual-chassis HA pair.

The following example shows a single-chassis HA pair:



In a single-chassis HA pair, both controllers are in the same chassis. The HA interconnect is provided by the internal backplane. No external HA interconnect cabling is required.

The following example shows a dual-chassis HA pair and the HA interconnect cables:



In a dual-chassis HA pair, the controllers are in separate chassis. The HA interconnect is provided by external cabling.

Interconnect cabling for systems with variable HA configurations

In systems that can be configured either as a single-chassis or dual-chassis HA pair, the interconnect cabling is different depending on the configuration.

The following table describes the interconnect cabling for 32xx and 62xx systems:

If the controller modules in the HA pair are...	The HA interconnect cabling is...
Both in the same chassis	Not required, since an internal interconnect is used
Each in a separate chassis	Required

HA configuration and the HA state PROM value

Some controller modules and chassis automatically record in a PROM whether they are in an HA pair or stand-alone. This record is the *HA state* and must be the same on all components within the stand-alone system or HA pair. The HA state can be manually configured if necessary.

Related tasks

[Verifying and setting the HA state on the controller modules and chassis](#) on page 52

For systems that use the HA state value, the value must be consistent in all components in the HA pair. You can use the Maintenance mode `ha-config` command to verify and, if necessary, set the HA state.

Table of storage system models and HA configuration differences

Supported storage systems can have significant differences in their HA configurations, depending on the model.

The following table lists the supported storage systems and their HA configuration differences:

Storage system model	HA configuration (single-chassis, dual-chassis, or either)	Interconnect type (internal InfiniBand, external InfiniBand, or external 10-Gb Ethernet)	Uses HA state PROM value?
FAS8060	Single-chassis	Internal InfiniBand	Yes
FAS8040	Single-chassis	Internal InfiniBand	Yes
FAS8020	Single-chassis	Internal InfiniBand	Yes
6290	Single-chassis or dual-chassis	<ul style="list-style-type: none"> Dual-chassis: External InfiniBand using NVRAM adapter Single-chassis: Internal InfiniBand 	Yes
6280	Single-chassis or dual-chassis	<ul style="list-style-type: none"> Dual-chassis: External InfiniBand using NVRAM adapter Single-chassis: Internal InfiniBand 	Yes
6250	Single-chassis or dual-chassis	<ul style="list-style-type: none"> Dual-chassis: External InfiniBand using NVRAM adapter Single-chassis: Internal InfiniBand 	Yes
6240	Single-chassis or dual-chassis	<ul style="list-style-type: none"> Dual-chassis: External InfiniBand using NVRAM adapter Single-chassis: Internal InfiniBand 	Yes
6220	Single-chassis or dual-chassis	<ul style="list-style-type: none"> Dual-chassis: External InfiniBand using NVRAM adapter Single-chassis: Internal InfiniBand 	Yes
6210	Single-chassis or dual-chassis	<ul style="list-style-type: none"> Dual-chassis: External InfiniBand using NVRAM adapter Single-chassis: Internal InfiniBand 	Yes
60xx	Dual-chassis	External InfiniBand using NVRAM adapter	No

Storage system model	HA configuration (single-chassis, dual-chassis, or either)	Interconnect type (internal InfiniBand, external InfiniBand, or external 10-Gb Ethernet)	Uses HA state PROM value?
3270	Single-chassis or dual-chassis	<ul style="list-style-type: none"> • Dual-chassis: External 10-Gb Ethernet using onboard ports c0a and c0b These ports are dedicated HA interconnect ports. Regardless of the system configuration, these ports cannot be used for data or other purposes. • Single-chassis: Internal InfiniBand 	Yes
3250	Single-chassis or dual-chassis	<ul style="list-style-type: none"> • Dual-chassis: External 10-Gb Ethernet using onboard ports c0a and c0b These ports are dedicated HA interconnect ports. Regardless of the system configuration, these ports cannot be used for data or other purposes. • Single-chassis: Internal InfiniBand 	Yes
3240	Single-chassis or dual-chassis	<ul style="list-style-type: none"> • Dual-chassis: External 10-Gb Ethernet using onboard ports c0a and c0b These ports are dedicated HA interconnect ports. Regardless of the system configuration, these ports cannot be used for data or other purposes. • Single-chassis: Internal InfiniBand 	Yes
3220	Single-chassis or dual-chassis	<ul style="list-style-type: none"> • Dual-chassis: External 10-Gb Ethernet using onboard ports c0a and c0b These ports are dedicated HA interconnect ports. Regardless of the system configuration, these ports cannot be used for data or other purposes. • Single-chassis: Internal InfiniBand 	Yes

Storage system model	HA configuration (single-chassis, dual-chassis, or either)	Interconnect type (internal InfiniBand, external InfiniBand, or external 10-Gb Ethernet)	Uses HA state PROM value?
3210	Single-chassis or dual-chassis	<ul style="list-style-type: none"> • Dual-chassis: External 10-Gb Ethernet using onboard ports c0a and c0b These ports are dedicated HA interconnect ports. Regardless of the system configuration, these ports cannot be used for data or other purposes. • Single-chassis: Internal InfiniBand 	Yes
31xx	Single-chassis	Internal InfiniBand	No
FAS22xx	Single-chassis	Internal InfiniBand	Yes

Installing and cabling an HA pair

To install and cable a new standard or mirrored HA pair, you must have the correct tools and equipment and you must connect the controllers to the disk shelves. If it is a dual-chassis HA pair, you must also cable the HA interconnect between the nodes. HA pairs can be installed in either NetApp system cabinets or in equipment racks.

The term V-Series system refers to the storage systems released prior to Data ONTAP 8.2.1 that can use array LUNs. The FAS systems released in Data ONTAP 8.2.1 and later can use array LUNs if the proper license is installed.

The specific procedure you use depends on the following aspects of your configuration:

- Whether you have a standard or mirrored HA pair
- Whether you are using FC or SAS disk shelves

Note: If your configuration includes SAS disk shelves, see the *Universal SAS and ACP Cabling Guide* on the NetApp Support Site at support.netapp.com for information about cabling. For cabling the HA interconnect between the nodes, use the procedures in this guide.

Multipath HA is required on all HA pairs except for some FAS22xx system configurations, which use single-path HA and lack the redundant standby connections.

System cabinet or equipment rack installation

You need to install your HA pair in one or more NetApp system cabinets or in standard telco equipment racks. Each of these options has different requirements.

HA pairs 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.

HA pairs in a system cabinet

Depending on the number of disk shelves, the HA pair you ordered arrives in a single system cabinet or multiple system cabinets.

The number of system cabinets you receive depends on how much storage you ordered. All internal adapters, such as networking adapters, Fibre Channel adapters, and other adapters, arrive preinstalled in the nodes.

If it comes in a single system cabinet, both the Channel A and Channel B disk shelves are cabled, and the HA adapters are also pre-cabled.

If the HA pair you ordered has more than one cabinet, 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 HA interconnects. If the HA pair 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.

Required documentation

Installing an HA pair requires the correct documentation.

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

Manual name	Description
<i>Hardware Universe</i>	This guide describes the physical requirements that your site must meet to install NetApp equipment.
The appropriate system cabinet guide	This guide describes how to install NetApp equipment into a system cabinet.
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 Network Management Guide for 7-Mode</i>	This guide describes how to perform network configuration for the storage system.
<i>Data ONTAP Upgrade and Revert/Downgrade Guide for 7-Mode</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 for 7-Mode</i>	This guide describes, among other topics, SyncMirror technology, which is used for mirrored HA pairs.
<i>Data ONTAP System Administration Guide for 7-Mode</i>	This guide describes general storage system administration.
<i>Data ONTAP Software Setup Guide for 7-Mode</i>	This guide describes how to configure the software of a new storage system for the first time.

Note: If you are installing a Data ONTAP system HA pair with array LUNs, see the *FlexArray Virtualization Installation Requirements and Reference Guide* for information about cabling Data ONTAP systems to storage arrays. You can also refer to the *FlexArray Virtualization Implementation Guide for Third-Party Storage* for information about configuring storage arrays to work with Data ONTAP systems.

Related information

[Data ONTAP Information Library](#)

Required tools

Installation of an HA pair requires the correct tools.

The following list specifies the tools you need to install the HA pair :

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

Required equipment

When you receive your HA pair, you should receive the equipment listed in the following table. See the *Hardware Universe* at hwu.netapp.com to confirm your storage system type, storage capacity, and so on.

Required equipment	Details
Storage system	Two of the same type of storage systems
Storage	See the <i>Hardware Universe</i> at hwu.netapp.com
HA interconnect adapter (for controller modules that do not share a chassis) Note: When 32xx systems are in a dual-chassis HA pair, the c0a and c0b 10-GbE ports are the HA interconnect ports. They do not require an HA interconnect adapter. Regardless of configuration, the 32xx system's c0a and c0b ports cannot be used for data. They are only for the HA interconnect.	InfiniBand (IB) HA adapter (The NVRAM adapter functions as the HA interconnect adapter on FAS900 series and later storage systems, except the 32xx systems)
For SAS disk shelves: SAS HBAs, if applicable For DS14 disk shelves: FC-AL or FC HBA (FC HBA for disk) adapters, if applicable	Minimum of two SAS HBAs, two FC-AL adapters, or their equivalent in onboard ports
Fibre Channel switches, if applicable	Not applicable
SFP (small form-factor pluggable) modules, if applicable	Not applicable
NVRAM HA adapter media converter	Only if using fiber cabling
Cables (provided with shipment unless otherwise noted)	<ul style="list-style-type: none"> For systems using FC disk shelf connections, two optical controller-to-disk shelf cables per loop For systems using SAS disk shelf connections, two SAS controller-to-disk shelf cables per stack Multiple disk shelf-to-disk shelf cables, if applicable For systems using the IB HA interconnect adapter, two 4xIB copper cables, two 4xIB optical cables, or two optical cables with media converters <p>Note: You must purchase longer optical cables separately for cabling distances greater than 30 meters.</p> <ul style="list-style-type: none"> The 32xx systems, when in a dual-chassis HA pair, require 10 GbE cables (Twinax or SR optical) for the HA interconnect.

Preparing your equipment

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

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 guide for your disk shelf, 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 disk shelf guide.
3. Label the interfaces, where appropriate.
4. Connect the nodes to the network, as described in the setup instructions for your system.

Result

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

After you finish

Proceed to cable the HA pair .

Installing the nodes in a system cabinet

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

Steps

1. Install the system cabinets, nodes, and disk shelves 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, as described in the *Installation and Setup Instructions* for your system.
3. Connect the system cabinets to an appropriate power source and apply power to the cabinets.

Result

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

After you finish

Proceed to cable the HA pair.

Cabling a standard HA pair

To cable a standard HA pair, you identify the ports you need to use on each node. You then cable the ports and cable the HA interconnect.

About this task

This procedure explains how to cable a configuration using DS14mk2 or DS14mk4 disk shelves.

For cabling SAS disk shelves in an HA pair, see the *Universal SAS and ACP Cabling Guide*.

Note: If you are installing an HA pair between nodes using array LUNs, see the *FlexArray Virtualization Installation Requirements and Reference Guide* for information about cabling V-

Series systems to storage arrays. See the *FlexArray Virtualization Implementation Guide for Third-Party Storage* for information about configuring storage arrays to work with Data ONTAP.

The sections for cabling the HA interconnect apply to all systems regardless of disk shelf type.

Steps

1. [Determining which Fibre Channel ports to use for Fibre Channel disk shelf connections](#) on page 36
2. [Cabling Node A to DS14mk2 or DS14mk4 disk shelves](#) on page 37
3. [Cabling Node B to DS14mk2 or DS14mk4 disk shelves](#) on page 38
4. [Cabling the HA interconnect \(all systems except 32xx\)](#) on page 40
5. [Cabling the HA interconnect \(32xx systems in separate chassis\)](#) on page 40

Determining which Fibre Channel ports to use for Fibre Channel disk shelf connections

Before cabling your HA pair, 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 HA pair requires two ports on the node, one for the primary connection and one for the redundant multipath HA connection.
A standard HA pair with one loop for each node uses four 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 *Hardware Universe* at hwu.netapp.com for your platform for an HA pair.
- If using Fibre Channel HBAs, insert the adapters in the same slots on both systems.

See the *Hardware Universe* at hwu.netapp.com to obtain all slot assignment information for the various adapters you use to cable your HA pair.

After identifying the ports, you should have a numbered list of Fibre Channel ports for both nodes, starting with Port 1.

Cabling guidelines for 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:

- Disk shelf loops using ESH4 modules must be cabled to the quad-port HBA first.
- Disk shelf loops using AT-FCX modules must be cabled to dual-port HBA ports or onboard ports before using ports on the quad-port HBA.
- Port A of the HBA must be cabled to the In port of Channel A of the first disk shelf in the loop. Port A of the partner node's HBA must be cabled 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.
- Additional disk shelf loops must be cabled sequentially with the HBA's ports. Port A is used for the first loop, port B for the second loop, and so on.
- If available, ports C or D must be used for the redundant multipath HA connection after cabling all remaining disk shelf loops.
- All other cabling rules described in the documentation for the HBA and the *Hardware Universe* must be observed.

Cabling Node A to DS14mk2 or DS14mk4 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.

About this task

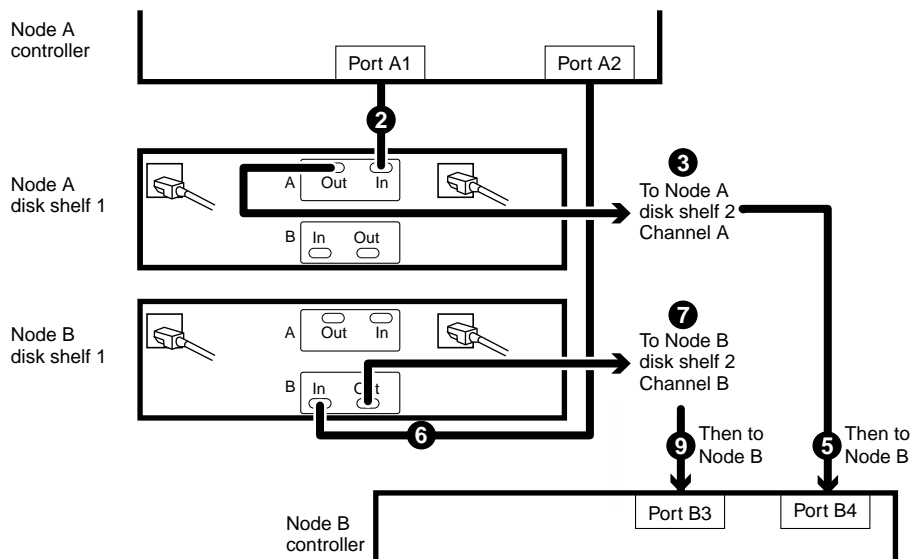
- This procedure uses multipath HA, which is required on all systems.
- This procedure does not apply to SAS disk shelves.

For cabling SAS disk shelves in an HA pair, see the *Universal SAS and ACP Cabling Guide*.

Note: You can find additional cabling diagrams in your system's *Installation and Setup Instructions* on the NetApp Support Site.

Steps

1. Review the cabling diagram before proceeding to the cabling steps.
 - 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.



2. Cable Fibre Channel port A1 of Node A to the Channel A Input port of the first disk shelf of Node A loop 1.
3. Cable the Node A disk shelf Channel A Output port to the Channel A Input port of the next disk shelf in loop 1.

4. Repeat step 3 for any remaining disk shelves in loop 1.
5. Cable the Channel A Output port of the last disk shelf in the loop to Fibre Channel port B4 of Node B.
This provides the redundant multipath HA connection for Channel A.
6. Cable Fibre Channel port A2 of Node A to the Channel B Input port of the first disk shelf of Node B loop 1.
7. Cable the Node B disk shelf Channel B Output port to the Channel B Input port of the next disk shelf in loop 1.
8. Repeat step 7 for any remaining disk shelves in loop 1.
9. Cable the Channel B Output port of the last disk shelf in the loop to Fibre Channel port B3 of Node B.
This provides the redundant multipath HA connection for Channel B.
10. Repeat steps 2 to 9 for each pair of loops in the HA pair , using ports 3 and 4 for the next loop, ports 5 and 6 for the next one, and so on.

Result

Node A is completely cabled.

After you finish

Cable Node B.

Cabling Node B to DS14mk2 or DS14mk4 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.

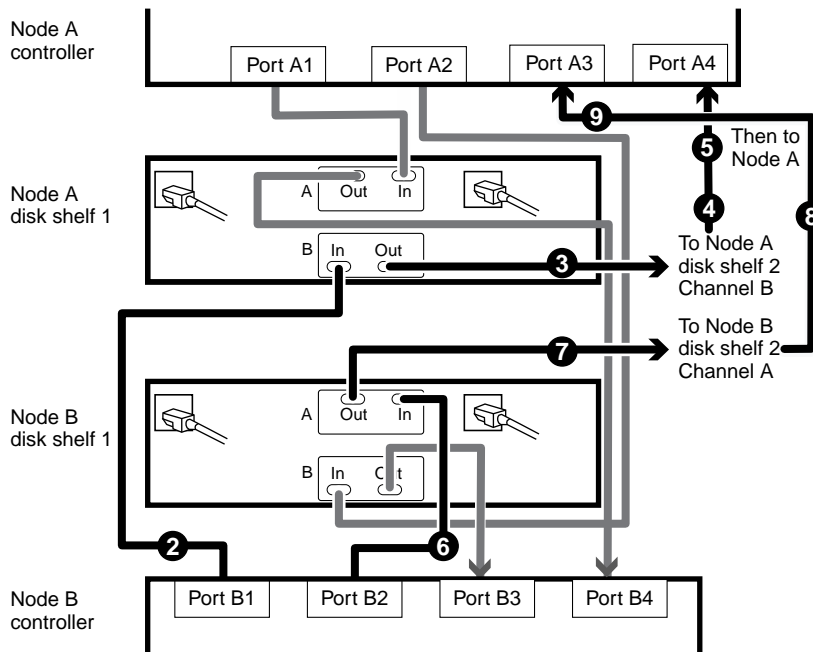
About this task

- This procedure uses multipath HA, required on all systems.
- This procedure does not apply to SAS disk shelves.
For cabling SAS disk shelves in an HA pair , see the *Universal SAS and ACP Cabling Guide*.

Note: You can find additional cabling diagrams in your system's *Installation and Setup Instructions* on the NetApp Support Site at support.netapp.com.

Steps

1. Review the cabling diagram before proceeding to the cabling steps.
 - 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.



2. Cable Port B1 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 HA pair easier to administer because the disks have the same ID on each node. This is true for Step 5 also.
3. Cable the disk shelf Channel B Output port to the Channel B Input port of the next disk shelf in loop 1.
4. Repeat Step 3 for any remaining disk shelves in loop 1.
5. Cable the Channel B Output port of the last disk shelf in the loop to Fibre Channel port A4 of Node A.
This provides the redundant multipath HA connection for Channel B.
6. Cable Fibre Channel port B2 of Node B to the Channel A Input port of the first disk shelf of Node B loop 1.
7. Cable the disk shelf Channel A Output port to the Channel A Input port of the next disk shelf in loop 1.
8. Repeat Step 7 for any remaining disk shelves in loop 1.
9. Cable the Channel A Output port of the last disk shelf in the loop to Fibre Channel port A3 of Node A.
This provides the redundant multipath HA connection for Channel A.
10. Repeat Steps 2 to 9 for each pair of loops in the HA pair , using ports 3 and 4 for the next loop, ports 5 and 6 for the next one, and so on.

Result

Node B is completely cabled.

After you finish

Cable the HA interconnect.

Cabling the HA interconnect (all systems except 32xx)

To cable the HA interconnect between the HA pair 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.

About this task

This procedure applies to all dual-chassis HA pairs (HA pairs in which the two controller modules reside in separate chassis) except 32xx systems, regardless of disk shelf type.

Steps

1. See the *Hardware Universe* at hwu.netapp.com to ensure that your interconnect adapter is in the correct slot for your system in an HA pair .

For systems that use an NVRAM adapter, the NVRAM adapter functions as the HA interconnect adapter.

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

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

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

```
HA 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 HA adapters.

Result

The nodes are connected to each other.

After you finish

Configure the system.

Cabling the HA interconnect (32xx systems in separate chassis)

To enable the HA interconnect between 32xx controller modules that reside in separate chassis, you must cable the onboard 10-GbE ports on one controller module to the onboard GbE ports on the partner.

About this task

This procedure applies to 32xx systems regardless of the type of attached disk shelves.

Steps

1. Plug one end of the 10-GbE cable to the c0a port on one controller module.
2. Plug the other end of the 10-GbE cable to the c0a port on the partner controller module.
3. Repeat the preceding steps to connect the c0b ports.

Do not cross-cable the HA interconnect adapter; cable the local node ports only to the identical ports on the partner node.

Result

The nodes are connected to each other.

After you finish

Proceed to configure the system.

Cabling a mirrored HA pair

To cable a mirrored HA pair, you identify the ports you need to use on each node, cable those ports, and then cable the HA interconnect.

About this task

This procedure explains how to cable a configuration using DS14mk2 or DS14mk4 disk shelves.

For cabling SAS disk shelves in an HA pair, see the *Universal SAS and ACP Cabling Guide*.

Note: If you are installing an HA pair between nodes using array LUNs, see the *FlexArray Virtualization Installation Requirements and Reference Guide* for information about cabling Data ONTAP systems to storage arrays. See the *FlexArray Virtualization Implementation Guide for Third-Party Storage* for information about configuring storage arrays to work with Data ONTAP.

The sections for cabling the HA interconnect apply to all systems regardless of disk shelf type.

Determining which Fibre Channel ports to use for Fibre Channel disk shelf connections

Before cabling your HA pair, 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 HA pair requires two ports on the node, one for the primary connection and one for the redundant multipath HA connection.
A standard HA pair with one loop for each node uses four 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 *Hardware Universe* at hwu.netapp.com for your platform for an HA pair.
- If using Fibre Channel HBAs, insert the adapters in the same slots on both systems.

See the *Hardware Universe* at hwu.netapp.com to obtain all slot assignment information for the various adapters you use to cable your HA pair.

After identifying the ports, you should have a numbered list of Fibre Channel ports for both nodes, starting with Port 1.

Cabling guidelines for 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:

- Disk shelf loops using ESH4 modules must be cabled to the quad-port HBA first.
- Disk shelf loops using AT-FCX modules must be cabled to dual-port HBA ports or onboard ports before using ports on the quad-port HBA.
- Port A of the HBA must be cabled to the In port of Channel A of the first disk shelf in the loop. Port A of the partner node's HBA must be cabled 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.
- Additional disk shelf loops must be cabled sequentially with the HBA's ports. Port A is used for the first loop, port B for the second loop, and so on.
- If available, ports C or D must be used for the redundant multipath HA connection after cabling all remaining disk shelf loops.
- All other cabling rules described in the documentation for the HBA and the *Hardware Universe* must be observed.

Creating your port list for mirrored HA pairs

After you determine which Fibre Channel ports to use, you can create a table that identifies which ports belong to which port pool.

About this task

Mirrored HA pairs, regardless of disk shelf type, 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. You must assign disks to the pools appropriately.

Follow the guidelines for software-based disk ownership in the *Data ONTAP Storage Management Guide for 7-Mode*.

For more information about SyncMirror, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

Step

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

For a 32xx HA pair that has two mirrored loops, the port list might 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

After you finish

Proceed to cable the Channel A loops.

Cabling the Channel A DS14mk2 or DS14mk4 disk shelf loops

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

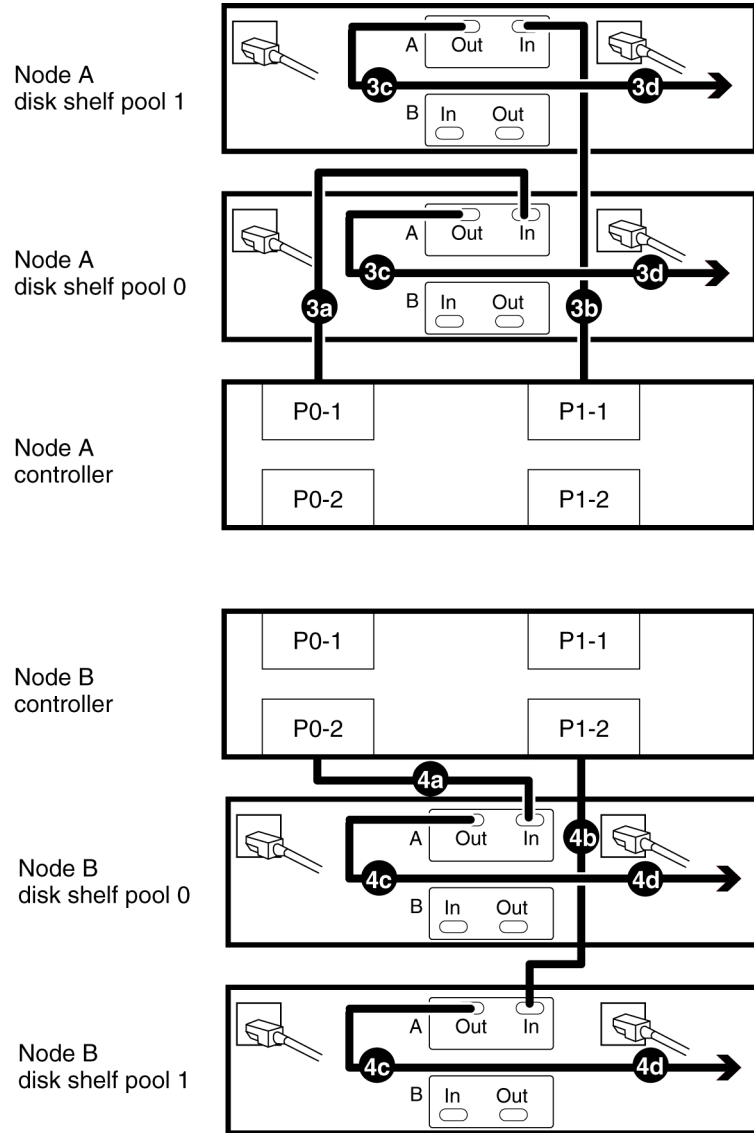
About this task

- This procedure uses multipath HA, which is required on all systems.
- This procedure does not apply to SAS disk shelves.
For cabling SAS disk shelves in an HA pair, see the *Universal SAS and ACP Cabling Guide*.

Steps

1. Complete your port list.
2. Review the cabling diagram before proceeding to the cabling steps.
 - 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.



3. Cable Channel A for Node 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.
- 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.
- 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.

- Repeat Substep 3c, connecting the next Channel A output to the next disk shelf Channel A Input port for any remaining disk shelves in this loop for each disk pool.

- e. Repeat Substep 3a through Substep 3d 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).
4. 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 4c, connecting Channel A output to input, for any remaining disk shelves in each disk pool.
 - e. Repeat Substep 4a through Substep 4d 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).

After you finish


Proceed to cable the Channel B loops.

Cabling the Channel B DS14mk2 or DS14mk4 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.

Steps

1. Review the cabling diagram before proceeding to the cabling steps.
 - 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.



2. 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 HA pair 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 2c, connecting Channel B output to input, for any remaining disk shelves in each disk pool.

- e. Repeat Substep 2a through Substep 2d 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).
3. 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 3c, connecting Channel B output to input, for any remaining disk shelves in each disk pool.
 - e. Repeat Substep 3a through Substep 3d 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).

After you finish


Proceed to cable the HA interconnect.

Cabling the redundant multipath HA connection for each loop

To complete the multipath HA cabling for the disk shelves, you must add the final connection for each channel on the final disk shelf in each loop.

Steps

1. Review the cabling diagram before proceeding to the cabling steps.
 - 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.



2. Connect the Channel A output port on the last disk shelf for each loop belonging to Node A to an available port on Node B in the same pool.
3. Connect the Channel B output port on the last disk shelf for each loop belonging to Node A to an available port on Node B in the same pool.
4. Connect the Channel A output port on the last disk shelf for each loop belonging to Node B to an available port on Node B in the same pool.
5. Connect the Channel B output port on the last disk shelf for each loop belonging to Node B to an available port on Node B in the same pool.

Cabling the HA interconnect (all systems except 32xx)

To cable the HA interconnect between the HA pair 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.

About this task

This procedure applies to all dual-chassis HA pairs (HA pairs in which the two controller modules reside in separate chassis) except 32xx systems, regardless of disk shelf type.

Steps

1. See the *Hardware Universe* at hwi.netapp.com to ensure that your interconnect adapter is in the correct slot for your system in an HA pair .

For systems that use an NVRAM adapter, the NVRAM adapter functions as the HA interconnect adapter.

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

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

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

```
HA 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 HA adapters.

Result

The nodes are connected to each other.

After you finish

Configure the system.

Cabling the HA interconnect (32xx systems in separate chassis)

To enable the HA interconnect between 32xx controller modules that reside in separate chassis, you must cable the onboard 10-GbE ports on one controller module to the onboard GbE ports on the partner.

About this task

This procedure applies to 32xx systems regardless of the type of attached disk shelves.

Steps

1. Plug one end of the 10-GbE cable to the c0a port on one controller module.
2. Plug the other end of the 10-GbE cable to the c0a port on the partner controller module.
3. Repeat the preceding steps to connect the c0b ports.

Do not cross-cable the HA interconnect adapter; cable the local node ports only to the identical ports on the partner node.

Result

The nodes are connected to each other.

After you finish

Proceed to configure the system.

Required connections for using uninterruptible power supplies with standard or mirrored HA pairs

You can use a UPS (uninterruptible power supply) with your HA pair. 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 HA pair.

For a standard HA pair, you must connect the controller, disks, and any FC switches in use.

For a mirrored HA pair, you must connect the controller and any FC switches to the UPS, as for a standard HA pair. 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 HA pair cannot shut down gracefully.

Configuring an HA pair

Bringing up and configuring a standard or mirrored HA pair for the first time can require enabling HA mode capability and failover, setting options, configuring network connections, and testing the configuration.

These tasks apply to all HA pairs regardless of disk shelf type.

Steps

1. [Bringing up the HA pair](#) on page 50
2. [Verifying and setting the HA state on the controller modules and chassis](#) on page 52
3. [Enabling HA mode capability and controller failover](#) on page 53
4. [Setting options and parameters](#) on page 53
5. [Configuring network interfaces for HA pairs](#) on page 60
6. [Verifying the HA pair configuration](#) on page 71
7. [Testing takeover and giveback](#) on page 71

Bringing up the HA pair

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

Considerations for HA pair setup

When the setup program runs on a storage system in an HA pair, it prompts you to answer some questions specific for HA pairs.

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 interface groups for your network interfaces?
For information about interface groups, see the *Data ONTAP Network Management Guide for 7-Mode*.

Note: You are advised to use interface groups with HA pairs 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 HA pair 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 tasks

[Configuring shared interfaces during setup](#) on page 51

When you set up the storage system, you can assign an IP address to a network interface and assign a partner IP address that the network interface takes over if a failover occurs.

[Configuring dedicated interfaces with setup](#) on page 51

You can assign a dedicated IP address to a network interface, so that the interface does not have a partner IP address. Although most interfaces are configured to take over a partner IP address during failover events, a dedicated interface enables isolation of the interface to a single IP address, even in takeover. This option can be used, for example, if network throughput performance of the interface is critical.

[Configuring standby interfaces with setup](#) on page 52

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

Configuring shared interfaces during setup

When you set up the storage system, you can assign an IP address to a network interface and assign a partner IP address that the network interface takes over if a failover occurs.

Steps

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

```
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 network mask, or netmask, for the interface you are configuring, or press Return if the default value is correct:

```
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:

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

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

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

Note: If the partner is an interface group, you must use the interface name.

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. Although most interfaces are configured to take over a partner IP address during failover events, a dedicated interface enables isolation of the interface to a single IP address, even in takeover. This option can be used, for example, if network throughput performance of the interface is critical.

About this task

This procedure is performed during setup of the storage system.

Note: If you configure a dedicated interface on the local node and want to maintain availability of the IP address during failover events, you must configure an interface on the partner node to take over the IP address of the dedicated interface.

Steps

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

```
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 Enter if the default value is correct:

```
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:

```
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]:
```

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
```

Verifying and setting the HA state on the controller modules and chassis

For systems that use the HA state value, the value must be consistent in all components in the HA pair. You can use the Maintenance mode `ha-config` command to verify and, if necessary, set the HA state.

About this task

The `ha-config` command only applies to the local controller module and, in the case of a dual-chassis HA pair, the local chassis. To ensure consistent HA state information throughout the system, you must also run these commands on the partner controller module and chassis, if necessary.

The HA state is recorded in the hardware PROM in the chassis and in the controller module. It must be consistent across all components of the system, as shown in the following table:

If the system or systems are in a...	The HA state is recorded on these components...	The HA state on the components must be...
Stand-alone configuration (not in an HA pair)	<ul style="list-style-type: none"> • The chassis • Controller module A 	non-ha
A single-chassis HA pair	<ul style="list-style-type: none"> • The chassis • Controller module A • Controller module B 	ha
A dual-chassis HA pair	<ul style="list-style-type: none"> • Chassis A • Controller module A • Chassis B • Controller module B 	ha

Use the following steps to verify the HA state is appropriate and, if not, to change it:

Steps

1. Reboot the current controller module and press Ctrl-C when prompted to display the boot menu.
2. At the boot menu, select the option for Maintenance mode boot.
3. After the system boots into Maintenance mode, enter the following command to display the HA state of the local controller module and chassis:

```
ha-config show
```

The HA state should be **ha** for all components if the system is in an HA pair.

4. If necessary, enter the following command to set the HA state of the controller:

```
ha-config modify controller ha-state
```

5. If necessary, enter the following command to set the HA state of the chassis:

```
ha-config modify chassis ha-state
```

6. Exit Maintenance mode by entering the following command:

```
halt
```

7. Boot the system by entering the following command at the boot loader prompt:

```
boot_ontap
```

8. If necessary, repeat the preceding steps on the partner controller module.

Enabling HA mode capability and controller failover

The HA license is no longer required in Data ONTAP 8.2. You must manually configure each node to enable or disable HA pair (high-availability) mode capability and controller failover.

Steps

1. Enter the following command on each of the node consoles to enable HA mode capability:

```
options cf.mode ha
```

2. Enter the following command on each of the node consoles to reboot the nodes:

```
reboot
```

3. Enter the following command on either of the node consoles to enable controller failover:

```
cf enable
```

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

```
cf status
```

The system displays the following output if controller failover is enabled:

```
Controller Failover enabled, Node2 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.

Option types for HA pairs

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

In an HA pair, options are one of the following types:

- Options that must be the same on both nodes for the HA pair 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 HA pair 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 must have the same settings 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 by entering 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 any options with the following comments in parentheses are set to the same value for both nodes:

```
(Value might be overwritten in takeover)
(Same value required in local+partner)
(Same value in local+partner recommended)
```

3. Correct any mismatched options by entering the `options option_name option_value` command.

Note: See the `na_options` man page for more information about the options.

Parameters that must be the same on each node

Certain parameters must be the same on each node so that takeover is smooth and data is transferred correctly between the nodes.

The parameters corresponding to the settings listed in the following table must be the same on both nodes:

Parameter..	Setting for..
Date	date, rdate
NDMP (on or off)	ndmp (on or off)
Route table published	route
Routed enabled	routed (on or off)
Time zone	timezone

Best practices for cf options

You should generally use the default options, since they represent the best choice for most configurations.

The following table lists best practices for the `cf` options available in Data ONTAP. These settings are changed with the `options` command. You use the `man options` command to display the man page with detailed descriptions of the options.

Option	Recommended value and notes
<code>cf.giveback.auto.cancel.on_network_failure</code>	ON Leave this option on to avoid automatic giveback without assurance that the partner's network interfaces are fully operational.
<code>cf.giveback.auto.cifs.terminate.minutes</code>	5 minutes (default value)
<code>cf.giveback.auto.delay.seconds</code>	600 seconds (default value)
<code>cf.giveback.auto.enable</code>	OFF (default value) Leave this option off so that in cases other than takeover due to reboot or panic, you can verify that the failure that caused the takeover is fixed before performing a giveback. Each node in an HA pair can have a different setting for this option.
<code>cf.giveback.auto.after.panic.takeover</code>	ON (default value)
<code>cf.giveback.auto.override.vetoes</code>	OFF (default value) Each node in an HA pair can have a different setting for this option.
<code>cf.giveback.check.partner</code>	ON (default value) Leave on to reduce downtime caused by a unsuccessful giveback. Each node in an HA pair can have a different setting for this option.
<code>cf.hw_assist.enable</code>	ON (default value) The node must support remote management through a Remote LAN Management card or Service Processor to enable the hardware-assisted takeover.
<code>cf.hw_assist.partner.address</code>	When using hardware-assisted takeover, the value of this option should be equal to partner's node management IP address.
<code>cf.hw_assist.partner.port</code>	4444 (default value)
<code>cf.takeover.change_fsid</code>	ON (default value) Each node in an HA pair can have a different setting for this option.
<code>cf.takeover.detection.seconds</code>	15 seconds (default value) If <code>sk.process.timeout.override</code> has been manually set, it is strongly advised that this option is set to a value larger than or equal to the value of <code>sk.process.timeout.override+5</code> .
<code>cf.takeover.on_failure</code>	ON (default value) Changing the value on one node automatically changes the value on the partner.

Option	Recommended value and notes
<code>cf.takeover.on_network_interface_failure</code>	OFF (default value) Changing the value on one node automatically changes the value on the partner.
<code>cf.takeover.on_network_interface_failure.policy</code>	Use <code>all_nics</code> to avoid spurious takeovers due to any one network interface failure.
<code>cf.takeover.on_panic</code>	ON (default value) You should not turn this option <code>off</code> unless you are instructed by technical support to do so. If the <code>cf.takeover.on_panic</code> option is set to <code>off</code> , a node panic will not cause an automatic takeover.
<code>cf.takeover.on_reboot</code>	For systems not licensed for FC or iSCSI, the takeover on reboot outage and subsequent giveback outage are shorter than the longer outage incurred by the reboot. You should use the default value of ON for these systems. For systems licensed for FC or iSCSI, you should use the default value of OFF. In these systems, the reboot outage is shorter than the combined takeover and giveback outage. Changing the value on one node automatically changes the value on the partner.
<code>cf.takeover.on_short_uptime</code>	ON (default value) Leave this option on to avoid outages due to failures occurring early in the node's boot process. Changing the value on one node automatically changes the value on the partner.
<code>cf.takeover.use_mcr_file</code>	OFF (default value) Use only in MetroCluster environments.

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.



Attention: 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

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

[Recovering from a disaster by using MetroCluster configurations](#) on page 202

In situations such as prolonged power outages or natural disasters, you can use the 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.

When data loss can occur if 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 offline, 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 that 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_fsid` option is set to `on`, the system changes the fsids when the `cf forcetakeover -d` command is issued so clients are forced to remount their volumes and check the integrity of the data before proceeding.

Configuring hardware-assisted takeover

You can configure hardware-assisted takeover to speed up takeover times. Hardware-assisted takeover uses the remote management device to quickly communicate local status changes to the partner node.

How hardware-assisted takeover speeds up takeover

Hardware-assisted takeover speeds up the takeover process by using a node's remote management device (SP or RLM) to detect failures and quickly initiate the takeover rather than waiting for Data ONTAP to recognize that the partner's heartbeat has stopped.

Without hardware-assisted takeover, if a failure occurs, the partner waits until it notices that the node is no longer giving a heartbeat, confirms the loss of heartbeat, and then initiates the takeover.

The hardware-assisted takeover feature uses the following process to take advantage of the remote management device and avoid that wait:

1. The remote management device monitors the local system for certain types of failures.
2. If a failure is detected, the remote management device immediately sends an alert to the partner node.
3. Upon receiving the alert, the partner initiates takeover.

The hardware-assisted takeover option (`cf.hw_assist.enable`) is enabled by default.

Related tasks

[Disabling and enabling the hardware-assisted takeover option](#) on page 58

Hardware-assisted takeover is enabled by default on systems that use remote management. Hardware-assisted takeover speeds the takeover process by using the RLM or SP to quickly detect potential takeover events and alerting the partner node.

[Setting the partner address for hardware-assisted takeover](#) on page 58

By default, on systems with an e0M port, the hardware-assisted takeover feature uses the IP address of the partner's e0M port to communicate with the partner. On systems without an e0M port, the system automatically selects another configured IP address. You can use the `cf.hw_assist.partner.address` option to select a different IP address.

[Setting the partner port for hardware-assisted takeover](#) on page 59

When hardware-assisted takeover is enabled, the remote management (either RLM or SP) 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.

Disabling and enabling the hardware-assisted takeover option

Hardware-assisted takeover is enabled by default on systems that use remote management. Hardware-assisted takeover speeds the takeover process by using the RLM or SP to quickly detect potential takeover events and alerting the partner node.

Step

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
```

Related concepts

[How hardware-assisted takeover speeds up takeover](#) on page 20

Hardware-assisted takeover speeds up the takeover process by using a node's remote management device (SP or RLM) to detect failures and quickly initiate the takeover rather than waiting for Data ONTAP to recognize that the partner's heartbeat has stopped.

Related tasks

[Setting the partner address for hardware-assisted takeover](#) on page 58

By default, on systems with an e0M port, the hardware-assisted takeover feature uses the IP address of the partner's e0M port to communicate with the partner. On systems without an e0M port, the system automatically selects another configured IP address. You can use the `cf.hw_assist.partner.address` option to select a different IP address.

[Setting the partner port for hardware-assisted takeover](#) on page 59

When hardware-assisted takeover is enabled, the remote management (either RLM or SP) 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.

Setting the partner address for hardware-assisted takeover

By default, on systems with an e0M port, the hardware-assisted takeover feature uses the IP address of the partner's e0M port to communicate with the partner. On systems without an e0M port, the system automatically selects another configured IP address. You can use the `cf.hw_assist.partner.address` option to select a different IP address.

Step

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
```

Related concepts

[How hardware-assisted takeover speeds up takeover](#) on page 20

Hardware-assisted takeover speeds up the takeover process by using a node's remote management device (SP or RLM) to detect failures and quickly initiate the takeover rather than waiting for Data ONTAP to recognize that the partner's heartbeat has stopped.

Related tasks

[Disabling and enabling the hardware-assisted takeover option](#) on page 58

Hardware-assisted takeover is enabled by default on systems that use remote management. Hardware-assisted takeover speeds the takeover process by using the RLM or SP to quickly detect potential takeover events and alerting the partner node.

[Setting the partner port for hardware-assisted takeover](#) on page 59

When hardware-assisted takeover is enabled, the remote management (either RLM or SP) 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.

Setting the partner port for hardware-assisted takeover

When hardware-assisted takeover is enabled, the remote management (either RLM or SP) 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

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
```

Related concepts

[How hardware-assisted takeover speeds up takeover](#) on page 20

Hardware-assisted takeover speeds up the takeover process by using a node's remote management device (SP or RLM) to detect failures and quickly initiate the takeover rather than waiting for Data ONTAP to recognize that the partner's heartbeat has stopped.

Related tasks

[Disabling and enabling the hardware-assisted takeover option](#) on page 58

Hardware-assisted takeover is enabled by default on systems that use remote management. Hardware-assisted takeover speeds the takeover process by using the RLM or SP to quickly detect potential takeover events and alerting the partner node.

[Setting the partner address for hardware-assisted takeover](#) on page 58

By default, on systems with an e0M port, the hardware-assisted takeover feature uses the IP address of the partner's e0M port to communicate with the partner. On systems without an e0M

port, the system automatically selects another configured IP address. You can use the `cf.hw_assist.partner.address` option to select a different IP address.

Configuring network interfaces for HA pairs

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.

Understanding interfaces in an HA pair

You can configure three types of interfaces on nodes in an HA pair .

What the networking interfaces do

When a node in an HA pair 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 HA pair to maintain communication with the network if the partner fails.

See the *Data ONTAP Network Management Guide for 7-Mode* for a description of available options and the function that each option performs.

Note: You should always use multiple NICs with interface groups to improve networking availability for both stand-alone storage systems and systems in an HA pair .

Related concepts

[Shared, dedicated, and standby interfaces](#) on page 60

These different types of interfaces have different roles in normal and takeover mode.

[Interface roles in normal and takeover modes](#) on page 61

You can configure shared, dedicated, and standby interfaces in an HA pair . Each type has a different role in normal and takeover mode.

[Takeover configuration with shared interfaces](#) on page 62

You can configure two NICs, one on each node, to provide two shared interfaces to each node.

[Takeover configuration with dedicated and standby interfaces](#) on page 63

With two NICs on each node, one can provide a dedicated interface and the other can act as a standby interface.

[Interface types and configurations](#) on page 64

This table lists the configurations supported by each type of interface in an HA pair .

[IPv6 considerations in an HA pair](#) on page 64

When enabled, IPv6 provides features such as address autoconfiguration. Using these IPv6 features requires an understanding of how these features work with the HA pair functionality.

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 HA pair :

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.

Interface type	Description
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 HA pair interfaces are configured as shared interfaces because they do not require an extra NIC.

Related concepts

[What the networking interfaces do](#) on page 60

When a node in an HA pair 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 HA pair to maintain communication with the network if the partner fails.

[Interface roles in normal and takeover modes](#) on page 61

You can configure shared, dedicated, and standby interfaces in an HA pair . Each type has a different role in normal and takeover mode.

[Takeover configuration with shared interfaces](#) on page 62

You can configure two NICs, one on each node, to provide two shared interfaces to each node.

[Takeover configuration with dedicated and standby interfaces](#) on page 63

With two NICs on each node, one can provide a dedicated interface and the other can act as a standby interface.

[Interface types and configurations](#) on page 64

This table lists the configurations supported by each type of interface in an HA pair .

[IPv6 considerations in an HA pair](#) on page 64

When enabled, IPv6 provides features such as address autoconfiguration. Using these IPv6 features requires an understanding of how these features work with the HA pair functionality.

Interface roles in normal and takeover modes

You can configure shared, dedicated, and standby interfaces in an HA pair . 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

Related concepts

[What the networking interfaces do](#) on page 60

When a node in an HA pair 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 HA pair to maintain communication with the network if the partner fails.

[Shared, dedicated, and standby interfaces](#) on page 60

These different types of interfaces have different roles in normal and takeover mode.

[Takeover configuration with shared interfaces](#) on page 62

You can configure two NICs, one on each node, to provide two shared interfaces to each node.

[Takeover configuration with dedicated and standby interfaces](#) on page 63

With two NICs on each node, one can provide a dedicated interface and the other can act as a standby interface.

[Interface types and configurations](#) on page 64

This table lists the configurations supported by each type of interface in an HA pair .


[IPv6 considerations in an HA pair](#) on page 64

When enabled, IPv6 provides features such as address autoconfiguration. Using these IPv6 features requires an understanding of how these features work with the HA pair functionality.

Takeover configuration with shared interfaces

You can configure two NICs, one on each node, to provide two shared interfaces to each node.

In the following configuration illustration, you use two NICs to provide the two 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.

Related concepts

[What the networking interfaces do](#) on page 60

When a node in an HA pair 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 HA pair to maintain communication with the network if the partner fails.

[Shared, dedicated, and standby interfaces](#) on page 60

These different types of interfaces have different roles in normal and takeover mode.

[Interface roles in normal and takeover modes](#) on page 61

You can configure shared, dedicated, and standby interfaces in an HA pair . Each type has a different role in normal and takeover mode.

Takeover configuration with dedicated and standby interfaces on page 63

With two NICs on each node, one can provide a dedicated interface and the other can act as a standby interface.

Interface types and configurations on page 64

This table lists the configurations supported by each type of interface in an HA pair .


IPv6 considerations in an HA pair on page 64

When enabled, IPv6 provides features such as address autoconfiguration. Using these IPv6 features requires an understanding of how these features work with the HA pair functionality.

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.



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.

Related concepts

What the networking interfaces do on page 60

When a node in an HA pair 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 HA pair to maintain communication with the network if the partner fails.

Shared, dedicated, and standby interfaces on page 60

These different types of interfaces have different roles in normal and takeover mode.

Interface roles in normal and takeover modes on page 61

You can configure shared, dedicated, and standby interfaces in an HA pair . Each type has a different role in normal and takeover mode.

Takeover configuration with shared interfaces on page 62

You can configure two NICs, one on each node, to provide two shared interfaces to each node.

Interface types and configurations on page 64

This table lists the configurations supported by each type of interface in an HA pair .

[IPv6 considerations in an HA pair](#) on page 64

When enabled, IPv6 provides features such as address autoconfiguration. Using these IPv6 features requires an understanding of how these features work with the HA pair functionality.

Interface types and configurations

This table lists the configurations supported by each type of interface in an HA pair .

Interface	Shared	Dedicated	Standby	Partner parameter
Ethernet	Supported	Supported	Supported	IP address or interface name
Gigabit Ethernet	Supported	Supported	Supported	IP address or interface name
Logical interface (LIF)	Supported	Supported	Supported	LIF name
VLAN interface	Supported	Supported	Supported	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 HA pair in the same way as a regular Ethernet interface.

Related concepts

[What the networking interfaces do](#) on page 60

When a node in an HA pair 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 HA pair to maintain communication with the network if the partner fails.

[Shared, dedicated, and standby interfaces](#) on page 60

These different types of interfaces have different roles in normal and takeover mode.

[Interface roles in normal and takeover modes](#) on page 61

You can configure shared, dedicated, and standby interfaces in an HA pair . Each type has a different role in normal and takeover mode.

[Takeover configuration with shared interfaces](#) on page 62

You can configure two NICs, one on each node, to provide two shared interfaces to each node.

[Takeover configuration with dedicated and standby interfaces](#) on page 63

With two NICs on each node, one can provide a dedicated interface and the other can act as a standby interface.

[IPv6 considerations in an HA pair](#) on page 64

When enabled, IPv6 provides features such as address autoconfiguration. Using these IPv6 features requires an understanding of how these features work with the HA pair functionality.

IPv6 considerations in an HA pair

When enabled, IPv6 provides features such as address autoconfiguration. Using these IPv6 features requires an understanding of how these features work with the HA pair functionality.

For additional information about IPv6, see the *Data ONTAP Network Management Guide for 7-Mode*.

Configuration requirements for using IPv6

To use IPv6 in an HA pair, 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.

Related concepts

[What the networking interfaces do](#) on page 60

When a node in an HA pair 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 HA pair to maintain communication with the network if the partner fails.

[Shared, dedicated, and standby interfaces](#) on page 60

These different types of interfaces have different roles in normal and takeover mode.

[Interface roles in normal and takeover modes](#) on page 61

You can configure shared, dedicated, and standby interfaces in an HA pair. Each type has a different role in normal and takeover mode.

[Takeover configuration with shared interfaces](#) on page 62

You can configure two NICs, one on each node, to provide two shared interfaces to each node.

[Takeover configuration with dedicated and standby interfaces](#) on page 63

With two NICs on each node, one can provide a dedicated interface and the other can act as a standby interface.

[Interface types and configurations](#) on page 64

This table lists the configurations supported by each type of interface in an HA pair.

Making nondisruptive changes to the interface groups

You can use the `cf takeover` and `cf giveback` commands to make changes to interface groups in the HA pair in a nondisruptive manner.

About this task

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 HA pair, 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 interface groups.

See the *Data ONTAP Network Management Guide for 7-Mode* for more information about configuring interface groups.

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 interface groups.

4. Repeat these steps, making any required changes to the `/etc/rc` file on the partner node.

Configuring network interfaces for the HA pair

You must configure network interfaces so that if takeover occurs, interfaces on the operating node takes over interfaces on the failed-over node and hosts can still reach data over the network.

Before you begin

Both nodes in the HA pair 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.
- 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 interface groups to improve networking availability for both stand-alone storage systems and systems in an HA pair.

About this task

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 HA pair to use FC, see the *Data ONTAP SAN Administration Guide for 7-Mode*.

Steps

1. Determine whether to use shared interfaces or dedicated and standby interfaces.
2. Configure your interfaces on one node.
3. Repeat Step 2 on the other node.
4. If desired, configure automatic takeover for the interfaces or interface groups.

Configuring an IP address for a network interface

You can configure IP addresses for your network interface during system setup. To configure the IP addresses later, you should use the `ifconfig` command. You can configure both IPv4 and IPv6 addresses for a network interface.

About this task

- Network configuration changes made by using the `ifconfig` command are not automatically included in the `/etc/rc` file. To make the configuration changes persistent after reboots, include the `ifconfig` command in the `/etc/rc` file.

- When you configure an IP address, your storage system creates a network mask based on the class of the address (Class A, B, C, or D) by default.

Step

To configure an IP address for a network interface, enter the following command:

```
ifconfig interface_name IP_address
```

interface_name is the name of the network interface.

IP_address is the IP address that you want to assign to the network interface.

To configure a quad-port Ethernet interface e3a to use the IPv4 address 192.0.2.10, enter the following command:

```
ifconfig e3a 192.0.2.10
```

To configure a quad-port Ethernet interface e3a to use the IPv6 address 2001:0db8:35ab:0:8a2e:0:0370:85, enter the following command:

```
ifconfig e3a 2001:0db8:35ab:0:8a2e:0:0370:85
```

Configuring a partner interface in an HA pair

To prepare for a successful takeover in an HA configuration, you can map a network interface to an IP address or to another network interface on the partner node. During a takeover, the network interface on the surviving node assumes the identity of the partner interface.

Before you begin

When specifying the partner IP address, both the local network interface and the partner's network interface must be attached to the same network segment or network switch.

About this task

- If the network interface is an interface group, the partner interface must be denoted by an interface name and not an IP address.
The partner interface can be an interface group or a physical network interface.
- You cannot specify the underlying physical ports of an interface group in a partner configuration.
- If IPv6 addresses are to be taken over, you must specify the partner interface, and not an IP address.
Address to address mapping is not supported for IPv6 addresses.
- For the partner configuration to be persistent across reboots, you must include the `ifconfig` command in the `/etc/rc` file.
For a successful takeover in both directions, you must repeat the partner configuration in the `/etc/rc` files of each node.
- When specifying the partner interface name, you can configure the interfaces symmetrically, for example map interface e1 on one node to interface e1 on the partner node.
Though symmetrical configuration is not mandatory, it simplifies administration and troubleshooting tasks.

Step

Depending on the partner configuration that you want to specify, enter the following command:

If you want specify a...	Enter the following command...
Partner IP address	<pre>ifconfig interface_name partner address</pre> <p><i>interface_name</i> is the name of the network interface. <i>address</i> is the partner IP address.</p>
Partner interface name	<pre>ifconfig interface_name partner partner_interface</pre> <p><i>partner_interface</i> is the name of the partner network interface.</p>

Example: Specifying a partner IP address and partner interface name

Consider node1 and node2 are two storage systems in an HA configuration.

If the IP address of the interface e8 on node2 is 198.9.200.38, the following command allows the interface e1 of node1 to take over the IP address of node2 for the duration of the takeover:

```
node1> ifconfig e1 partner 198.9.200.38
```

Instead of specifying the IP address, you can also specify the partner interface name. The following command allows the interface e1 of node1 to assume the identity of e8 of node2 for the duration of the takeover:

```
node1> ifconfig e1 partner e8
```

Configuring partner addresses on different subnets (MetroCluster configurations 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.

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.

Examples of `/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 mycompany.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 mycompany.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 mycompany.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 mycompany.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 its 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 the 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.

The following example shows nodeA's `/etc/mcrd` 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 mycompany.com
options dns.enable on
options nis.enable off
savecore
```

3. Create an `/etc/mcrd` file on the other node (nodeB) and place it in the `/etc` directory.

The `/etc/mcrd` 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/mcrd` file by copying the `/etc/rc` file.

Note: The `/etc/mcrd` file must be configured manually. It is not updated automatically. It must include all commands necessary to configure the interfaces.

4. Enter the following commands in nodeB's `/etc/mcrd` 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 the 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 other options that are needed to correctly configure the interface in your network environment.

The following example shows nodeB's `/etc/mcrd` 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 mycompany.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 a MetroCluster configuration.

Step

Enter the following command on both nodes:

```
options cf.takeover.use_mcrc_file on
```

The default is `off`.

Verifying the HA pair configuration

You can go to the NetApp Support Site and download the Config Advisor tool to check for common configuration errors.

About this task

Config Advisor is a configuration validation and health check tool for NetApp systems. It can be deployed at both secure sites and non-secure sites for data collection and system analysis.

Note: Support for Config Advisor is limited, and available only online.

Steps

1. Log in to the NetApp Support Site at support.netapp.com and go to **Downloads > Utility ToolChest**.
2. Click **Config Advisor (WireGauge renamed)**.
3. Follow the directions on the web page for downloading, installing, and running the utility.
4. After running Config Advisor, review the tool's output and follow the recommendations to address any issues discovered.

Testing takeover and giveback

After you configure all aspects of your HA pair, you need to verify that it is operating as expected in maintaining uninterrupted access to both nodes' storage during takeover and giveback operations. Throughout the takeover process, the local (or takeover) node should continue serving the data normally provided by the partner node. During giveback, control and delivery of the partner's storage should return transparently to the partner node.

Steps

1. Check the cabling on the HA 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
```

See the man page for command details.

The local node takes over the partner node and gives the following output:

```
Failover monitor: takeover completed
```

4. Use the `sysconfig -r` command to ensure that the local (takeover) node can access its partner's disks.
5. Give back the partner node's data service after it displays the `Waiting for giveback` message 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 saw the message that giveback was completed successfully:

If takeover and giveback...	Then...
Is completed successfully	Repeat Step 2 through Step 5 on the partner node.
Fails	Correct the takeover or giveback failure and then repeat this procedure.

Managing takeover and giveback

An HA pair 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 HA pair differs depending on whether one partner has taken over the other, and the takeover and giveback operations themselves have different options.

This information applies to all HA pair s regardless of disk shelf type.

Monitoring an HA pair in normal mode

You can display information about the status and configuration of HA pair in normal mode (when neither node has taken over the other).

Monitoring HA pair 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 HA pair is up.

Step

Enter the following command:

```
cf status
```

The following example shows that the HA pair is enabled and the interconnect is up and working correctly:

```
node1>cf status
```

```
Controller Failover enabled, node2 is up.  
RDMA Interconnect is up (Link 0 up).
```

If the output shows that one link is down, the HA pair is degraded and you must configure the link so that it is up while the other link is still active.

Note: Depending on the storage system model, the output might display either `RDMA interconnect` or `VIA interconnect` in the last line.

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 failover is disabled.

Description of HA pair status messages

The `cf status` command displays information about the status of the HA pair.

The following table shows messages that the `cf status` command can display:

Message	Meaning
HA enabled, partner_name is up.	The HA pair is operating normally.
local_node has taken over partner_node.	One node has taken over the other node.
Interconnect not present.	The system does not recognize the existence of a HA interconnect adapter.
Interconnect is down.	The HA interconnect adapter cannot access the partner. This might be due to cabling problems or the partner might be down.

Message	Meaning
Interconnect is up.	The HA interconnect adapter is active and can transmit data to the partner.
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, takeover disabled because of reason (partner halted in notakeover mode) partner_name_1 has disabled takeover by partner_name_2 (interconnect error) interconnect_type Interconnect is down (Link 0 down).	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 has 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.

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.

Checking the hardware-assisted takeover status of the local and partner node

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

Enter the following command to display the hardware-assisted takeover status:

```
cf hw_assist status
```

Example of hardware-assisted takeover status

The following example shows output from the `cf hw_assist status` command:

```
Local Node Status - hal
```

```
Active: Monitoring alerts from partner(ha2)
port 4004 IP address 172.27.1.14
```

Partner Node Status - ha2

```
Active: Monitoring alerts from partner(ha1)
port 4005 IP address 172.27.1.15
```

Related tasks

[Testing the hardware-assisted takeover configuration](#) on page 75

You can test the hardware-assisted takeover configuration with the `cf hw_assist test` command.

[Checking hardware-assisted takeover statistics](#) on page 75

You can display statistics about hardware-assisted takeovers to determine how many alert events of each type have been received from the partner.

Testing the hardware-assisted takeover configuration

You can test the hardware-assisted takeover configuration with the `cf hw_assist test` command.

About this task

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

Enter the following command to test the hardware-assisted takeover configuration:

```
cf hw_assist test
```

After you finish

Depending on the message received from the `cf hw_assist test` command, you might need to reconfigure options so that the HA pair and the remote management card are operating.

Related tasks

[Checking the hardware-assisted takeover status of the local and partner node](#) on page 74

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.

[Checking hardware-assisted takeover statistics](#) on page 75

You can display statistics about hardware-assisted takeovers to determine how many alert events of each type have been received from the partner.

Checking hardware-assisted takeover statistics

You can display statistics about hardware-assisted takeovers to determine how many alert events of each type have been received from the partner.

Step

Enter the following command to display or clear the hardware-assisted takeover statistics, respectively:

```
cf hw_assist stats
```

```
cf hw_assist stats clear
```

Example of 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
```

Related tasks

[Checking the hardware-assisted takeover status of the local and partner node](#) on page 74

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.

[Testing the hardware-assisted takeover configuration](#) on page 75

You can test the hardware-assisted takeover configuration with the `cf hw_assist test` command.

Displaying the partner's name

You can display the name of the other node with the `cf partner` command.

Step

Enter the following command:

```
cf partner
```

Note: If the node does not yet know the name of its partner because the HA pair is new, this command returns `partner`.

Displaying disk and array LUN information on an HA pair

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 channel A and channel B:

- The information about disks on channel A is the same as for storage systems not in an HA pair .
- The information about disks on channel 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 channel B is a file system disk, spare disk, or parity disk.

Step

Enter one of the following commands:

```
sysconfig -r  
or  
aggr status -r
```

What takeover and giveback are

Takeover is the process in which a node takes over the storage of its partner. Giveback is the process in which that storage is returned to the partner. Both processes can be initiated manually or configured for automatic initiation.

When takeovers occur

Takeovers can be initiated manually or occur automatically when a failover event happens, depending on how you configure the HA pair. In some cases, takeovers occur automatically, regardless of configuration.

Takeovers can occur under the following conditions:

- A takeover is manually initiated.
- A node is in an HA pair with the `default` configuration for immediate takeover on panic, and that node undergoes a software or system failure that leads to a panic.

Note: The default setting for the `cf.takeover.on_panic` option is `on`. If the `cf.takeover.on_panic` option is set to `on`, a node panic will cause an automatic takeover. If the `cf.takeover.on_panic` option is set to `off`, a node panic will not cause an automatic takeover. You should not turn this option to `off` unless you are instructed by technical support to do so.

By default, after the partner has recovered from the panic and booted up, the node automatically performs a giveback, returning the partner to normal operation.

- A node that is in an HA pair 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 MetroCluster configurations, you can initiate a forced takeover in this situation.

- One or more network interfaces that are configured to support failover become unavailable.
- A node does not receive heartbeat messages from its partner.

This could happen if the partner 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` parameter.
- You reboot one of the nodes without using the `-f` parameter.
- Hardware-assisted takeover is enabled and triggers a takeover when the remote management device (RLM or Service Processor) detects failure of the partner node.

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.

If the node doing the takeover panics

If the node that is performing the takeover panics within 60 seconds of initiating takeover, the following events occur:

- The node that panicked reboots.
- After it reboots, the node performs self-recovery operations and is no longer in takeover mode.
- Failover is disabled.

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 will 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

The local node returns ownership of the aggregates and volumes to the partner node after any issues on the partner node are resolved or maintenance is complete. In addition, the local node returns ownership when the partner node has booted up and giveback is initiated either manually or automatically.

When the failed node is functioning again, the following events can occur:

- You issue a `cf giveback` command that terminates the emulated node on the partner.
- The failed node resumes normal operation, serving its own data.

- The HA pair resumes normal operation, with each node ready to take over for its partner if the partner fails.

Configuring automatic takeover

You can control when automatic takeovers occur by setting the appropriate options.

Reasons that automatic takeover occurs

You can set options to control whether automatic takeovers occur due to different system errors. In some cases, automatic takeover occurs by default unless you disable the option, and in some cases automatic takeover cannot be prevented.

Takeovers can happen for several reasons. Some system errors cause a takeover; for example, when a system in an HA pair 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 HA pair. The following table outlines which system errors can cause a takeover to occur, and whether you can configure the HA pair 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	On	If the <code>cf.takeover.on_panic</code> option is set to on , a node panic will cause an automatic takeover. If the <code>cf.takeover.on_panic</code> option is set to off , a node panic will not cause an automatic takeover. You should not turn this option off unless you are instructed by technical support to do so.

System error	Option used to configure	Default value	Notes
A node reboots.	<code>cf.takeover.on_reboot</code> set to on	On (unless FC or iSCSI is licensed) Note: In releases prior to Data ONTAP 8.0, the system would only take over the partner after reboot if the partner took longer than 90 seconds to boot.	If a takeover is performed because of the partner node rebooting, then an automatic giveback will be performed, regardless of the setting of the <code>cf.giveback.auto.enable</code> option.
All the network interface cards (NICs) or interface groups 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>	Takeover on network failure disabled (by default)	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 for 7-Mode</i> .
One or more of the NICs or interface groups enabled for negotiated failover failed. Note: If interfaces fail on both nodes in the HA pair, 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>	Takeover on network failure disabled (by default)	To enable a network interface or interface group for negotiated failover, you use the <code>ifconfig if_name nfo</code> command. For more information, see the <i>Data ONTAP MultiStore Management Guide for 7-Mode</i> .
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.	Not applicable	Not applicable	You cannot prevent this condition from causing a takeover.
You halt one of the nodes <i>without</i> using the <code>-f</code> flag.	Not applicable	Not applicable	You cannot prevent this condition from causing a takeover. If you include the <code>-f</code> flag, the takeover is prevented.

System error	Option used to configure	Default value	Notes
You initiate a takeover manually using the <code>cf takeover</code> command.	Not applicable	Not applicable	You cannot prevent this condition from causing a takeover.


Related tasks

[Enabling and disabling automatic takeover of a panicked partner](#) on page 83

Data ONTAP is configured by default to initiate a takeover immediately if the partner node panics. This shortens the time between the initial failure and the time that service is fully restored because the takeover can be quicker than the recovery from the panic, although the subsequent giveback causes another brief outage.

Commands for performing a manual takeover

You must know the commands you can use when initiating a takeover. You can initiate a takeover on a node in an HA pair to perform maintenance on that node while still serving the data on its disks, array LUNs, or both to users.

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 <code>off</code>). 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 partner node is dumping its core The partner node is not allowed to halt gracefully
<code>cf forcetakeover</code>	Tells the HA monitor to ignore some configuration problems that would otherwise prevent a takeover (such as unsynchronized NVRAM due to a faulty HA interconnect connection), and then initiates a takeover of the partner of the local node
<code>cf forcetakeover -d</code>	Initiates a takeover of the local partner even in the absence of a quorum of partner mailbox disks or partner mailbox LUNs The <code>cf forcetakeover -d</code> command is valid only if the <code>remote_syncmirror</code> option is enabled.  Attention: You should only use the <code>-d</code> option after you verify that the partner is down. 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 for 7-Mode</i> .
<code>cf takeover -n</code>	Initiates a takeover for a nondisruptive upgrade For more information, see the <i>Data ONTAP Upgrade and Revert/Downgrade Guide for 7-Mode</i>

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

Enter the following command:

```
halt -f
```

Rebooting a node without takeover

You can reboot the node and prevent its partner from taking over, overriding the `cf.takeover.on_reboot` option.

Step

Enter the following command:

```
reboot -f
```

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

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 takeover on reboot

The takeover on reboot option enables you to control whether an automatic takeover occurs when a node reboots. This automatic takeover, and the automatic giveback that follows after the reboot is complete, can reduce the outage during which the storage belonging to the rebooting system is unavailable.

About this task

If this option is enabled and a takeover occurs because of a reboot, then an automatic giveback is performed after the partner has booted. This giveback occurs even if the `cf.giveback.auto.enable` option is set to `off`. However, if a node takes over its partner due to a reboot and that node itself reboots before it can execute a giveback, it performs automatic giveback only if `cf.giveback.auto.enable` is set to `on`.

If the `cf.takeover.on_reboot` is off and a node is rebooted then the partner will not take over immediately. But the partner could take over later if the node takes more than 180 seconds to boot.

Note: If the `reboot -f` command is used, then the partner does not take over under any circumstances, even if the reboot timer expires.

Step

Enter the following command:

```
options cf.takeover.on_reboot on
```

The default is `on`, unless FC or iSCSI is licensed, in which case the default is `off`.

Note: If you enter this command on one node, the value applies to both nodes.

This option is persistent across reboots.

Enabling and disabling automatic takeover of a panicked partner

Data ONTAP is configured by default to initiate a takeover immediately if the partner node panics. This shortens the time between the initial failure and the time that service is fully restored because the takeover can be quicker than the recovery from the panic, although the subsequent giveback causes another brief outage.

About this task

- If you enter this command on one node, the value applies to both nodes.
- The setting of this option is persistent across reboots.
- By default, Data ONTAP will initiate an automatic giveback after a takeover on panic. The `cf.giveback.auto.after.panic.takeover` option can be used to disable this automatic giveback.

Steps

1. Verify that controller takeover is enabled by entering the following command:

```
cf enable
```

2. Enable or disable automatic takeover on panic by entering the following command:

```
options cf.takeover.on_panic {on|off}
```

- If the `cf.takeover.on_panic` option is set to `on`, a node panic will cause an automatic takeover.
- If the `cf.takeover.on_panic` option is set to `off`, a node panic will not cause an automatic takeover.

You should not turn this option `off` unless you are instructed by technical support to do so.

Specifying the time period before takeover

You can specify how long (in seconds) a partner in an HA pair 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 HA pair 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

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.

Enabling or disabling negotiated failover for a network interface

You can enable or disable negotiated failover for a network interface to trigger automatic takeover if the interface experiences a persistent failure. You can use the `nfo` option of the `ifconfig` command to enable or disable negotiated failover.

About this task

You can specify the `nfo` option for an interface group. However, you cannot specify the `nfo` option for any underlying physical interface of the interface group.

Steps

1. To enable takeover during interface failure, enter the following command:

```
options cf.takeover.on_network_interface_failure on
```

2. To enable or disable negotiated failover, enter the following command:

```
ifconfig interface_name {nfo|-nfo}
```

interface_name is the name of the network interface.

`nfo` enables negotiated failover.

`-nfo` disables negotiated failover.

To enable negotiated failover on the interface `e8` of an HA configuration, enter the following command:

```
ifconfig e8 nfo
```

Note: The `nfo` option is persistent across reboots after it is enabled on an interface.

Takeover of vFiler units and the vFiler limit

The vFiler limit, set with the `vfiler limit` command, determines how many vFiler units can exist on a system. In an HA pair, if the two systems have different vFiler limits, some vFiler units might not be taken over if a takeover occurs.

When performing a takeover, a system can take over only the number of vFiler units that were specified by that system's vFiler 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 has a higher vFiler limit, some vFiler units cannot be taken over successfully.

For more information about setting the vFiler limit, see the *Data ONTAP MultiStore Management Guide for 7-Mode*.

Managing an HA pair in takeover mode

You manage an HA pair in takeover mode by performing a number of management actions.

Determining why takeover occurred

You can use the `cf status` command to determine why a takeover occurred.

Step

At the takeover prompt, enter the following command:

```
cf status
```

Result

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.

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.

Accessing the remote node using the partner command without arguments

You can use the `partner` command to toggle between the partner mode, in which commands are executed on the partner node, and takeover mode.

Step

From the takeover prompt, enter the following command:

```
partner
```

Result

The prompt changes to the partner-mode prompt, which has the following form:
`emulated_node/takeover_node>`

Example of 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)>
```

Related tasks

[Accessing the takeover node with the partner command containing arguments](#) on page 86

You use the `partner` command with a Data ONTAP command as an argument, so you can execute single commands on the takeover node without entering partner mode.

Accessing the takeover node with the partner command containing arguments

You use the `partner` command with a Data ONTAP command as an argument, so you can execute single commands on the takeover node without entering partner mode.

Step

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)>
```

Related tasks

[Accessing the remote node using the partner command without arguments](#) on page 86

You can use the `partner` command to toggle between the partner mode, in which commands are executed on the partner node, and takeover mode.

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

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 on a takeover node are available on an 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`
- `rdate`
- `reboot`
- `timezone`

Commands with different behaviors

Command	Difference in emulated mode
<code>ifconfig interface</code>	<ul style="list-style-type: none"> Displays emulated interface mappings based on the failed node's <code>/etc/rc</code> file rather than the takeover node interface mappings Displays emulated interface names rather than the interface names of the takeover node Displays only interfaces that have been configured, unlike the takeover node, which displays all interfaces, configured or not
<code>mt</code>	Uses the tape devices on the takeover node because the failed node has no access to its tape devices
<code>netstat -i</code>	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 HA pair
<code>sysconfig</code>	Displays information only about the hardware that is attached to the takeover node 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.
<code>uptime</code>	Displays how long the failed node has been down and the host name of the takeover node
<code>aggr status</code>	Displays information only about the hardware that is attached to the takeover node 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 dump and restore operations for a failed node

You can use the emulated node and peripheral devices attached to the takeover node to perform dump and restore operations for the failed node.

About this task

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 HA pair.
- 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.

Because the peripheral devices for a failed node are inaccessible, you perform dump and restore operations 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 dump and restore operations, see the *Data ONTAP Data Protection Tape Backup and Recovery Guide for 7-Mode*.

Step

Issue the `backup` or `restore` command, either in `partner` mode or as an argument in the `partner` command.

Issuing a `restore` command in `partner` mode is shown in the following example:

```
node1 (takeover)> partner
node1/node2> restore [options [arguments]]
node1 (takeover)> partner
```

Issuing a `restore` command as an argument in the `partner` command:

```
node1 (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.

About manual 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 the failed drives in the taken-over system as described in the *Data ONTAP Storage Management Guide for 7-Mode*.

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`.

Related concepts

[If giveback is interrupted](#) on page 92

If the takeover node experiences a failure or a power outage during the giveback process, that process stops and the takeover node returns to takeover mode until the failure is repaired or the power is restored.

Related tasks

[Removing failed disks prior to attempting giveback](#) on page 90

For taken-over systems that use disks, you must remove the failed disk or disks prior to attempting to implement giveback.

[Initiating normal giveback](#) on page 90

You can return control to a taken-over partner with the `cf giveback` command.

[Forcing giveback](#) on page 91

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.

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

Remove the failed disks, as described in the *Data ONTAP Storage Management Guide for 7-Mode*.

After you finish

When all failed disks are removed or replaced, proceed with the giveback operation.

Related concepts

[Option for shortening giveback time](#) on page 89

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`.

[If giveback is interrupted](#) on page 92

If the takeover node experiences a failure or a power outage during the giveback process, that process stops and the takeover node returns to takeover mode until the failure is repaired or the power is restored.

Related tasks

[Initiating normal giveback](#) on page 90

You can return control to a taken-over partner with the `cf giveback` command.

[Forcing giveback](#) on page 91

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.

Initiating normal giveback

You can return control to a taken-over partner with the `cf giveback` command.

Before you begin

For fabric-attached MetroCluster configurations, the aggregates on the surviving node and the partner node must already be rejoined to reestablish the MetroCluster configuration.

Step

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, or there might be failed disks in the system. Remove any failed disks and retry the command. For processes that prevent giveback, you can wait and repeat the command. You can also initiate giveback using the `-f` option to override the checks that prevent giveback.

After a giveback, the takeover node's ability to take over its partner automatically is not reenabled until the partner successfully reboots. If the partner fails to reboot, you can enter the `cf takeover` command to initiate a manual takeover of the partner.

Related concepts

[Option for shortening giveback time](#) on page 89

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`.

If giveback is interrupted on page 92

If the takeover node experiences a failure or a power outage during the giveback process, that process stops and the takeover node returns to takeover mode until the failure is repaired or the power is restored.

Related tasks

Removing failed disks prior to attempting giveback on page 90

For taken-over systems that use disks, you must remove the failed disk or disks prior to attempting to implement giveback.

Forcing giveback on page 91

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.

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.

About this task

You can 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.

Note: The `cf forcegiveback` command should be used with caution because it can cause a loss of data. If you cannot risk loss of data and are unable to complete the giveback, contact technical support.

Steps

1. On the takeover node, enter the following command:

```
cf giveback -f
```

The `-f` parameter allows giveback to proceed as long as it would not result in data corruption or an error on the storage system.

2. If giveback is still not successful, and if you can risk possible loss of data, enter the following command on the takeover node:

```
cf forcegiveback
```



Attention: Use `cf forcegiveback` only when you cannot get `cf giveback -f` 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.

Related concepts

Option for shortening giveback time on page 89

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`.

If giveback is interrupted on page 92

If the takeover node experiences a failure or a power outage during the giveback process, that process stops and the takeover node returns to takeover mode until the failure is repaired or the power is restored.

Related tasks

[Removing failed disks prior to attempting giveback](#) on page 90

For taken-over systems that use disks, you must remove the failed disk or disks prior to attempting to implement giveback.

[Initiating normal giveback](#) on page 90

You can return control to a taken-over partner with the `cf giveback` command.

If giveback is interrupted

If the takeover node experiences a failure or a power outage during the giveback process, that process stops and the takeover node returns to takeover mode until the failure is repaired or the power is restored.

However, this depends upon the stage of giveback in which the failure occurred. If the node encountered failure or a power outage during partial giveback state (after it has given back the root aggregate), it will not return to takeover mode. Instead, the node returns to partial-giveback mode. If this occurs, complete the process by repeating the giveback operation.

Related concepts

[Option for shortening giveback time](#) on page 89

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`.

Related tasks

[Removing failed disks prior to attempting giveback](#) on page 90

For taken-over systems that use disks, you must remove the failed disk or disks prior to attempting to implement giveback.

[Initiating normal giveback](#) on page 90

You can return control to a taken-over partner with the `cf giveback` command.

[Forcing giveback](#) on page 91

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.

Configuring giveback

You can configure how giveback occurs, setting different Data ONTAP options to improve the speed and timing of giveback.

Configuring automatic giveback

You can enable automatic giveback by using the `cf.giveback.auto.enable` command.

About this task

You should 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.

- If an automatic takeover occurred because the partner node panicked, the default behaviour is that an automatic giveback occurs even if this option is set to **off**.

Step

Enter the following command to enable automatic giveback:

```
option cf.giveback.auto.enable on
```

The **on** value enables automatic giveback. The **off** value disables automatic giveback. This option is **off** by default.

Adjusting the giveback delay time for automatic giveback

By default, there is a 600-second minimum time that a node stays in the takeover state before performing an automatic giveback. This delay reduces the overall outage that can occur while the taken-over partner reboots. Instead of a single longer outage, there are two brief outages (first when the partner is taken over, the second when giveback occurs). This option affects all types of automatic giveback but does not affect manual giveback.

Step

Enter the following command:

```
options cf.giveback.auto.delay.seconds number of seconds
```

The valid values for *number_of_seconds* are 0 to 600. The default is 600.



Attention: If `cf.giveback.auto.delay.seconds` is set to 0, the combined outage during takeover and giveback results in a long total client outage.

Setting giveback delay time for CIFS clients

You can specify the number of minutes to delay an automatic giveback before the system terminates CIFS clients that have open files.

About this task

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 enabled.

Step

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.

Terminating long-running processes to speed automatic giveback

You can use the `cf.giveback.auto.override.vetoes` option to speed implementation of automatic giveback.

When enabled, the `cf.giveback.auto.override.vetoes` option specifies how automatic giveback handles long-running operations such as dump or restore. If the option is set to **on**, long-running operations are terminated immediately when automatic giveback is initiated. If the option is set to **off**, automatic giveback is deferred until the long-running operations are complete. This option is used only if automatic giveback is enabled.

Setting giveback to terminate long-running processes and override vetos

You can set the automatic giveback process to terminate long-running processes and override vetos that might prevent the giveback.

Step

Enter the following command:

```
options cf.giveback.auto.override.vetoes {on|off}
```

The **on** argument enables this option. The **off** argument disables this option. This option is **off** by default.

Enabling and disabling automatic giveback after takeover due to partner panicking

Data ONTAP is configured by default to initiate an automatic giveback after an automatic takeover that occurred due to the partner node panicking. This shortens the time between the initial failure and the full restoration of service.

About this task

- If you enter this command on one node, the value applies to both nodes.
- The setting of this option is persistent across reboots.
- This option is not affected by the setting of the `cf.giveback.auto.enable` option.
If `cf.giveback.auto.enable` is set to **OFF**, automatic giveback after takeover due to panic still occurs if `cf.giveback.auto.after.panic.takeover` is set to **ON**.

Steps

1. Ensure that you enabled controller takeover by entering the following command:

```
cf enable
```

2. Enable or disable automatic giveback after takeover on panic by entering the following command:

```
options cf.giveback.auto.after.panic.takeover {on|off}
```

on enables automatic giveback to the partner. This is the default value.

off disables automatic giveback after takeover on panic.

Reconfiguring an HA pair into two stand-alone systems

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

About this task

This procedure applies to all HA pairs regardless of disk shelf type.

Steps

1. *Ensuring uniform disk ownership within disk shelves and loops in the system* on page 95
2. *Configuring a node for non-HA (stand-alone) use* on page 95
3. *Reconfiguring nodes using disk shelves for stand-alone operation* on page 97
4. *Requirements when changing a node using array LUNs to stand-alone* on page 99
5. *Reconfiguring nodes using array LUNs for stand-alone operation* on page 99

Ensuring uniform disk ownership within disk shelves and loops in the system

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.

Before you begin

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

About this task

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 HA 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 HA software.

Configuring a node for non-HA (stand-alone) use

By default, storage controllers are configured for use in HA mode. To use a controller in stand-alone mode, you must disable the controller failover functionality and change the node to non-HA mode.

Before you begin

You must determine the current configuration of the storage controller because the controller failover and HA mode states can vary. You can use the `cf status` command to determine the current configuration.

You must also confirm that all loops and disk shelves in the system contain disks that belong to only one of the two nodes you intend to isolate. If any disk shelves or loops contain a mix of disks belonging to both nodes, you must move data.

About this task

When a storage controller is shipped from the factory or when Data ONTAP is reinstalled using option four of the Data ONTAP boot menu (Clean configuration and initialize all disks), HA mode is enabled by default, and the system's nonvolatile memory (NVRAM or NVMEM) is split. If you plan to use the controller in standalone mode, you must configure the node as non-HA. Reconfiguring the node as non-HA mode enables full use of the system's nonvolatile memory.

Note: Configuring the node as standalone removes the availability benefits of the HA configuration and creates a single point of failure.

For information on using the boot menu to manage the storage system, see the *Data ONTAP System Administration Guide for 7-Mode*.

Choices

- If the `cf status` output displays `Non-HA mode`, then the node is configured for non-HA mode and you are finished:

```
node > cf status
Non-HA mode.
```

- If the `cf status` output directs you to reboot, you must reboot the node to enable full use of the system's nonvolatile memory:

```
node> cf status
Non-HA mode. Reboot to use full NVRAM.
```

1. Reboot the node using the following command:

```
node> reboot
```

After the node reboots, you are finished.

- If the `cf status` output displays `Controller Failover enabled`, you must disable both controller failover and HA mode and then reboot the node to enable full use of the system's nonvolatile memory:

```
node> cf status
Controller Failover enabled
```

1. Disable controller failover using the following command:

```
node> cf disable
```

2. Set the mode to non-HA by using the following command:

```
node> options cf.mode non_ha
```

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

Original entry:

```
ifconfig e0 199.9.204.254 partner 199.9.204.255
```

Edited entry:

```
ifconfig e0 199.9.204.254
```

4. Reboot the node by using the following command:

```
node> reboot
```


After the node reboots, you are finished.

- If the `cf status` output displays `Controller Failover disabled`, then the HA mode is still enabled, so you must set the HA mode to non-HA and reboot the node to enable full use of the system's nonvolatile memory:

```
node> cf status
Controller Failover disabled
```

1. Set the mode to non-HA by using the following command:

```
node> options cf.mode non_ha
```

2. Reboot the node by using the following command:

```
node> reboot
```

After the node reboots, you are finished.

After you finish

Reconfigure the hardware.

Related concepts

[Requirements when changing a node using array LUNs to stand-alone](#) on page 99

After uncoupling storage systems using array LUNs in an HA pair, you might need to perform additional reconfiguration related to Data ONTAP ownership of array LUNs.

Related tasks

[Reconfiguring nodes using array LUNs for stand-alone operation](#) on page 99

After uncoupling the nodes in an HA pair, you have three options: each node can continue to own its assigned array LUNs, you can reassign its array LUNs to another storage system that can use array LUNs, or you can release the persistent reservations on the array LUNs so that the LUNs can be used by a non Data ONTAP host.

Reconfiguring nodes using disk shelves for stand-alone operation

You can reconfigure the hardware to return to a single-controller configuration.

Before you begin

The HA pair software must be disabled.

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 that belongs to the node that owns the disks.

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 Snapshot technology rather than physically moving the disks. See the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

After moving the data from the disk, you can zero the disk and use the `disk remove_ownership` command to erase the disk's ownership information. See the *Data ONTAP Storage Management Guide for 7-Mode*.

3. If you are completely removing one node so that all the disk shelves belong to a single stand-alone node, complete the following substeps:
 - a. Boot the node that you are removing into Maintenance mode as described in the *Data ONTAP System Administration Guide for 7-Mode*.
 - b. Enter the following command to reassign all disk shelves so that they all belong to the remaining node:

```
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 HA interconnect cables from both nodes.  
See the hardware documentation for your system for details.
  6. Move or remove the adapter used for the HA interconnect:

| If your system uses...                         | Then...                                                                                                                                             |
|------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| An HA interconnect adapter or an FC-VI adapter | Remove the adapter from the system.                                                                                                                 |
| An NVRAM5 or NVRAM6 adapter                    | You might need to change the slot position of the adapter. See the <i>Hardware Universe</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_ontap
```

12. Check HA pair status by entering the following command:

```
cf status
```

If the HA pair is disabled, you see the following message:

```
Non-HA mode.
```

13. Repeat Step 1 through Step 12 for the partner node.

**Related concepts**

[Requirements when changing a node using array LUNs to stand-alone](#) on page 99

After uncoupling storage systems using array LUNs in an HA pair, you might need to perform additional reconfiguration related to Data ONTAP ownership of array LUNs.

**Related tasks**

[Reconfiguring nodes using array LUNs for stand-alone operation](#) on page 99

After uncoupling the nodes in an HA pair, you have three options: each node can continue to own its assigned array LUNs, you can reassign its array LUNs to another storage system that can use array LUNs, or you can release the persistent reservations on the array LUNs so that the LUNs can be used by a non Data ONTAP host.

[Configuring a node for non-HA \(stand-alone\) use](#) on page 95

By default, storage controllers are configured for use in HA mode. To use a controller in stand-alone mode, you must disable the controller failover functionality and change the node to non-HA mode.

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

After uncoupling storage systems using array LUNs in an HA pair, 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 HA pair.

| 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 HA 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 HA 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> |

**Related tasks**

[Configuring a node for non-HA \(stand-alone\) use](#) on page 95

By default, storage controllers are configured for use in HA mode. To use a controller in stand-alone mode, you must disable the controller failover functionality and change the node to non-HA mode.

[Reconfiguring nodes using array LUNs for stand-alone operation](#) on page 99

After uncoupling the nodes in an HA pair, you have three options: each node can continue to own its assigned array LUNs, you can reassign its array LUNs to another storage system that can use array LUNs, or you can release the persistent reservations on the array LUNs so that the LUNs can be used by a non Data ONTAP host.

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

After uncoupling the nodes in an HA pair, you have three options: each node can continue to own its assigned array LUNs, you can reassign its array LUNs to another storage system that can use

array LUNs, or you can release the persistent reservations on the array LUNs so that the LUNs can be used by a non Data ONTAP host.

### Before you begin

The HA pair software must be disabled.

### About this task

If you want both nodes in the HA pair 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, continues to see the array LUNs owned by the other system. 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 HA interconnect cables from both nodes. See the hardware documentation for your system for details.
4. Move or remove the adapter used for the HA interconnect.

---

| <b>If your system uses...</b>                  | <b>Then...</b>                                                                                                                                      |
|------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| An HA interconnect adapter or an FC-VI adapter | Remove the adapter from the system.                                                                                                                 |
| An NVRAM5 or NVRAM6 adapter                    | You might need to change the slot position of the adapter. See the <i>Hardware Universe</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:

---

| <b>If you want to...</b>                                                                                                                    | <b>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 | Boot both systems by entering the following command on each system:<br><pre>boot_ontap</pre> |

---

---

| If you want to...                                                                                                                       | Then...                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Remove one of the systems from service but still use the storage that was assigned to that system for Data ONTAP                        | <p>a. Boot the node being removed into Maintenance mode, as described in the <i>Data ONTAP System Administration Guide for 7-Mode</i>.</p> <p>b. Use the <code>disk reassign</code> command to reassign all the array LUNs so that they all belong to the node that remains.</p> <p>The <code>disk reassign</code> command has the following syntax:</p> <pre>disk reassign [-o &lt;old_name&gt;   -s &lt;old_sysid&gt;] [-n &lt;new_name&gt;] -d &lt;new_sysid&gt;</pre> <p>c. Remove the node from service.</p> <p>d. Boot the node you are keeping in service by entering the following command:</p> <pre>boot_ontap</pre> |
| 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 | <p>Release the persistent reservations that Data ONTAP placed on those array LUNs so that the storage administrator can use those LUNs for other hosts.</p> <p>See the <i>Data ONTAP Storage Management Guide for 7-Mode</i> for information about what you need to do to prepare for taking a system using array LUNs out of service.</p>                                                                                                                                                                                                                                                                                    |

---

### Related tasks

[Configuring a node for non-HA \(stand-alone\) use](#) on page 95

By default, storage controllers are configured for use in HA mode. To use a controller in stand-alone mode, you must disable the controller failover functionality and change the node to non-HA mode.

## Managing DS14mk2 or DS14mk4 disk shelves in an HA pair

---

You must follow specific procedures to add disk shelves to an HA pair or a MetroCluster configuration, or to upgrade or replace disk shelf hardware in an HA pair.

If your configuration includes SAS disk shelves, see the following documents on the NetApp Support Site:

- For SAS disk shelf management, see the *Installation and Service Guide* for your disk shelf model.
- For cabling SAS disk shelves in an HA pair, see the *Universal SAS and ACP Cabling Guide*.
- For cabling SAS disk shelves in a MetroCluster configuration using FibreBridge 6500N bridges, see *Configuring a MetroCluster system with SAS disk shelves and FibreBridge 6500N bridges*.
- For cabling SAS disk shelves in a stretch MetroCluster configuration using SAS optical cables, see *Configuring a stretch MetroCluster system with SAS disk shelves and SAS optical cables*.

### Adding DS14mk2 or DS14mk4 disk shelves to a multipath HA loop

To add supported DS14mk2 or DS14mk4 disk shelves to an HA pair configured for multipath HA, you must 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.

#### About this task

This procedure does not apply to SAS disk shelves.

#### Steps

1. Confirm that there are two paths to every disk by entering the following command:

```
storage show disk -p
```

**Note:** If two paths are not 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 DiskShelf14mk2 and DiskShelf14mk4 FC, or DiskShelf14mk2 AT *Hardware Service Guide*.
3. Find the last disk shelf in the loop to which you want to add the new disk shelf. The Channel A Output port of the last disk shelf in the loop is connected back to one of the controllers.

**Note:** In Step 4 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.

4. 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.

5. 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.
6. Connect the cable and SFP you removed in Step 4 to the Channel A Output port of the new disk shelf.
7. If you disabled the adapter in Step 3, reenable the adapter by entering the following command:

```
fcadmin config -e adapter
```

8. Repeat Step 4 through Step 7 for Channel B.

**Note:** The Channel B Output port is connected to the other controller.

9. Confirm that there are two paths to every disk by entering the following command:

```
storage show disk -p
```

Two paths should be listed for every disk.

#### Related tasks

[Determining path status for your HA pair](#) on page 104

If you want to remove a module from your HA pair, you need to know if the path you plan to disrupt is redundant. You can use the `storage show disk -p` command to indicate whether the disks have redundant paths.

## Upgrading or replacing modules in an HA pair

In an HA pair with redundant pathing, you can upgrade or replace disk shelf modules without interrupting access to storage.

#### About this task

These procedures are for DS14mk2 or DS14mk4 disk shelves.

If your configuration includes SAS disk shelves, refer to the following documents on the NetApp Support Site:

- For SAS disk shelf management, see the *Installation and Service Guide* for your disk shelf model.
- For cabling SAS disk shelves in an HA pair, see the *Universal SAS and ACP Cabling Guide*.
- For cabling SAS disk shelves in a MetroCluster configuration using FibreBridge 6500N bridges, see *Configuring a MetroCluster system with SAS disk shelves and FibreBridge 6500N bridges*.
- For cabling SAS disk shelves in a stretch MetroCluster configuration using SAS optical cables, see *Configuring a stretch MetroCluster system with SAS disk shelves and SAS optical cables*.

## About the disk shelf modules

A disk shelf module (ESH4 or AT-FCX) in a 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 DS14mk2, DS14mk4 FC or DS14mk2 AT disk shelves support the 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 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 HA pair, make sure that you understand the restrictions.

You cannot mix ESH4 modules in the same loop with AT-FCX modules.

## Best practices for changing module types

If you plan to change the type of any module in your HA pair, make sure that you review the best practice guidelines.

- Whenever you remove a module from an HA pair, 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 HA pair, 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:** ESH4 modules are self-terminating; this guideline does not apply to 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.

## 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.

## Determining path status for your HA pair

If you want to remove a module from your HA pair, you need to know if the path you plan to disrupt is redundant. You can use the `storage show disk -p` command to indicate whether the disks have redundant paths.

### About this task

If the disks have redundant paths, 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 of the disks in your HA pair, you must take action to ensure that you do not incur system downtime.

### Step

Enter the `storage show disk -p` command on your system console.



This command displays the following information for every disk in the HA pair :

- 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 HA pair consisting of Data ONTAP 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  |
| 0c.124  | A    | 0b.124    | B    | 7     | 12  |
| 0b.125  | B    | 0c.125    | A    | 7     | 13  |

Notice that every disk (for example, 0c.112/0b.112) has two active ports: one for A and one for B. The presence of the redundant path means that you do not need to fail over a node before removing modules from the system.



**Attention:** Make sure that every disk has two paths. Even in an HA pair configured for redundant paths, a hardware or configuration problem can cause one or more disks to have only one path. If any disk in your HA pair has only one path, you must treat that loop as if it were in a single-path HA pair when removing modules.

The following example shows what the `storage show disk -p` command output might look like for an HA pair consisting of Data ONTAP systems that do 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 |

For this HA pair, there is only one path to each disk. This means that you cannot remove a module from the configuration and disable that path without first performing a takeover.

## 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; disabling that single path 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. This firmware update causes a service interruption on non-multipath HA AT-FCX installations, multipath HA configurations running versions of Data ONTAP prior to 7.3.1, and systems with non-RoHS AT-FCX modules.

### 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 *DiskShelf14mk2 and DiskShelf14mk4 FC*, or *DiskShelf14mk2 AT Hardware Service 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. Complete the following steps if any disks use this loop as their only path to a controller:
  - a. Follow the cables from the module you want to replace back to one of the nodes, called NodeA.
  - b. Enter the following command at the NodeB console:  

```
cf takeover
```
  - c. Wait for takeover to be complete and make sure that the partner node, or NodeA, reboots and is waiting for giveback.  
Any module in the loop that is attached to NodeA can now be replaced.
4. Put on the antistatic wrist strap and grounding leash.
5. Disconnect the module that you are removing from the Fibre Channel cabling.
6. Using the thumb and index fingers of both hands, press the levers on the CAM mechanism on the module to release it and pull it out of the disk shelf.
7. Slide the replacement module into the slot at the rear of the disk shelf and push the levers of the cam mechanism into place.



**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.

8. Recable the disk shelf to its original location.
9. 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.

10. Complete the following steps if you performed a takeover previously:
  - a. Return control of NodeA's disk shelves by entering the following command at the console of the takeover node:

```
cf giveback
```

- b. Wait for the giveback to be completed before proceeding to the next step.

11. Test the replacement module.

12. Test the configuration.

### **Related concepts**

[Best practices for changing module types](#) on page 104

If you plan to change the type of any module in your HA pair, make sure that you review the best practice guidelines.

### **Related tasks**

[Determining path status for your HA pair](#) on page 104

If you want to remove a module from your HA pair, you need to know if the path you plan to disrupt is redundant. You can use the `storage show disk -p` command to indicate whether the disks have redundant paths.

[Hot-removing disk shelves in systems running Data ONTAP 8.2.1 or later](#) on page 108

If your system is running Data ONTAP 8.2.1 or later, you can hot-remove disk shelves—physically remove disk shelves that have had the aggregates removed from the disk drives—in a 7-Mode multipath HA configuration with DS14 disk shelves that is up and serving data. You can hot-remove one or more disk shelves from anywhere within a loop of disk shelves or remove a loop of disk shelves.

## Nondisruptive operations with HA pairs

---

By taking advantage of an HA pair's takeover and giveback operations, you can change hardware components and perform software upgrades in your configuration without disrupting access to the system's storage.

You can perform nondisruptive operations on a system by having its partner take over the system's storage, performing maintenance, and then giving back the storage.

### Where to find procedures for nondisruptive operations with HA pairs

An HA pair allows you to perform an extensive range of nondisruptive system maintenance and upgrade operations. You can refer to the specific documents for the required procedures.

The following table lists where you can find information on nondisruptive operations:

| If you want to perform this task nondisruptively... | See the...                                                      |
|-----------------------------------------------------|-----------------------------------------------------------------|
| Upgrade Data ONTAP                                  | <i>Data ONTAP Upgrade and Revert/Downgrade Guide for 7-Mode</i> |
| Replace a hardware FRU component                    | FRU procedures for your platform                                |

### Hot-removing disk shelves in systems running Data ONTAP 8.2.1 or later

If your system is running Data ONTAP 8.2.1 or later, you can hot-remove disk shelves—physically remove disk shelves that have had the aggregates removed from the disk drives—in a 7-Mode multipath HA configuration with DS14 disk shelves that is up and serving data. You can hot-remove one or more disk shelves from anywhere within a loop of disk shelves or remove a loop of disk shelves.

#### Before you begin

**Note:** If you have a MetroCluster configuration with SAS disk shelves, see the appropriate MetroCluster document for hot-removing SAS disk shelves.

For stretch MetroCluster systems using SAS optical cables, see the *Configuring a stretch MetroCluster system with SAS disk shelves and SAS optical cables* document. For fabric and stretch MetroCluster systems using FibreBridge 6500N bridges and SAS copper cables, see the *Configuring a MetroCluster system with SAS disk shelves and FibreBridge 6500N bridges* document.

These documents can be found on the NetApp Support Site at [support.netapp.com](http://support.netapp.com).

- Your storage system must be running Data ONTAP 8.2.1 or later.
- Your storage system must be a multipath HA system.
- For FAS2240 configurations, the external storage must be cabled as multipath HA.
- You must have already removed all aggregates from the disk drives in the disk shelves you are removing.



**Attention:** If you attempt this procedure on non MetroCluster configurations with aggregates on the disk shelf you are removing, you could fail the system with a multi-disk panic.

For MetroCluster configurations, removing aggregates involves splitting the mirrored aggregates on the disk shelves you are removing and then re-creating the mirrored aggregates

with another set of disk drives. For procedures, see "Splitting a mirrored aggregate" and "Rejoining split aggregates" in the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

For non MetroCluster configurations, see "Taking an aggregate offline" in the *Storage Management Guide for 7-Mode* for information. This document is available on the NetApp Support Site at [support.netapp.com](http://support.netapp.com).

- As a best practice, you should remove disk drive ownership after you remove the aggregates from the disk drives in the disk shelves you are removing.

**Note:** This procedure follows the best practice of removing disk drive ownership; therefore, steps are written with the assumption that you have removed disk drive ownership.

For information about removing disk drive ownership, see the "Removing ownership from a disk" procedure in the *Storage Management Guide for 7-Mode*. This document is available on the NetApp Support Site at [support.netapp.com](http://support.netapp.com).

**Note:** The procedure for removing ownership from disk drives requires you to disable disk autoassignment. You reenable disk autoassignment when prompted at the end of this shelf hot-remove procedure.

- Multipath HA configurations cannot be in a takeover state.
- If you are removing one or more disk shelves from within a loop, you must have factored the distance to bypass the disk shelves you are removing; therefore, if the current cables are not long enough, you need to have longer cables available.  
The *Hardware Universe* at [hwu.netapp.com](http://hwu.netapp.com) contains information about supported cables.

### About this task

- This procedure follows cabling best practices; therefore, references to modules and module input and output ports align with the best practices. If your storage system is cabled differently from what is prescribed as best practice, the modules and/or module input and output ports might be different.
- *Path A* refers to the A-side disk shelf module (module A) located in the top of the disk shelf.
- *Path B* refers to the B-side disk shelf module or module B located in the bottom location in the disk shelf.
- The *first* disk shelf in the loop is the disk shelf with the input ports directly connected to the controllers or switches.
- The *interim* disk shelf in the loop is the disk shelf directly connected to other disk shelves in the loop.
- The *last* disk shelf in the loop is the disk shelf with output ports directly connected to the controllers, except for fabric MetroCluster systems.  
For fabric MetroCluster systems, the output ports of the last disk shelf are not connected; they are not used.
- The *next* disk shelf is the disk shelf downstream of the disk shelf being removed, in depth order.
- The *previous* disk shelf is the disk shelf upstream of the disk shelf being removed, in depth order.
- For MetroCluster configurations, after completing this procedure for the local site, you are instructed to repeat this procedure for the remote site if you are also removing disk shelves for the remote site.

### Steps

1. Verify that your system configuration is `Multi-Path HA` by entering the following command from the console of either controller:

```
sysconfig
```

It might take up to a minute for the system to complete discovery.

The configuration is listed in the `System Storage Configuration` field.

**Note:** For FAS2240 systems with external storage, the output is displayed as `Mixed-Path HA` because the internal storage is cabled as `single-path HA` and the external storage is cabled as `multipath HA`.



**Attention:** If your non FAS2240 system configuration is shown as something other than `Multi-Path HA`, you cannot continue with this procedure. Your system must meet the prerequisites stated in the "Before you begin" section of this procedure.

2. Verify that no aggregates are on the disk drives in the disk shelves you are removing, by completing the following substeps:

- a. Enter the following command at the console of either controller:

```
aggr status -r
```

If you have traditional volumes, you enter the `vol status -r` command.

- b. Check the output `SHELF` column to verify that the disk shelves from which you have removed aggregates do not appear.



**Attention:** If any disk drives in the disk shelves you are removing are anything other than spares or failed, you cannot continue with this procedure. Your system must meet the prerequisites stated in the "Before you begin" section of this procedure.

The following `aggr status -r` output is for a 7-Mode system with disk shelf 3 being removed. Because aggregates have been removed from disk shelf 3, only disk shelves in the loop with aggregates (disk shelves 2 and 4) appear in the `SHELF` column:

```
fas6200> aggr status -r
RAID Disk Device HA SHELF BAY CHAN Pool Type RPM Used (MB/blks) Phys (MB/blks)

...
dparity 0c.33 0c 2 1 FC:A 0 FCAL 15000 136000/278528000
137104/280790184
parity 0c.38 0c 2 6 FC:A 0 FCAL 15000 136000/278528000
137104/280790184
data 0c.64 0c 4 0 FC:A 0 FCAL 15000 136000/278528000
137104/280790185
...
```

3. Verify that ownership is removed from the disk drives on the disk shelves you are removing:

- a. Enter the following command at the console of either controller:

```
disk show -v
```

- b. Check the output and take the appropriate action:

| If the output shows... | Then...                                                                                                |
|------------------------|--------------------------------------------------------------------------------------------------------|
| Not Owned or FAILED    | Go to the next step.<br>The disk drives in the disk shelves you are removing are in the correct state. |

| If the output shows...                                              | Then...                                                                                                                    |
|---------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------|
| Any disk drives in the disk shelves you are removing have ownership | You can use the "Removing ownership from a disk" procedure referenced in the "Before you begin" section of this procedure. |

The following output for the `disk show -v` command shows that disk drives on the disk shelf being removed (disk shelf 3) are in a correct state for removal. Two disk drives have the ownership removed; therefore, `Not Owned` appears in the `OWNER` column. Two disk drives are failed; therefore, `FAILED` appears in the `POOL` column:

```
fas6200> disk show -v
```

| DISK  | OWNER     | POOL   | SERIAL NUMBER | DR HOME | CHKSUM |
|-------|-----------|--------|---------------|---------|--------|
| ...   |           |        |               |         |        |
| 4b.5  | Not Owned | NONE   | YHGLB7PB      | Block   |        |
| 4b.11 | Not Owned | NONE   | YHGL33ZA      | Block   |        |
| 4b.15 |           | FAILED |               | Block   |        |
| 4b.13 |           | FAILED |               | Block   |        |
| ...   |           |        |               |         |        |

4. Turn on the LEDs for each disk drive in the disk shelves you are removing so that the disk shelves are physically identifiable by completing the following substeps:

You need to be certain of which disk shelves you are removing so that you can correctly recable path A and path B later in this procedure.

You can enter the commands from the console of either controller.

- a. Identify the disk drives in each disk shelf you are removing:

**fcadmin device\_map**

In this output, the shelf mapping shows three disk shelves in a loop and their respective 14 disk drives. If disk shelf 3 is being removed, disk drives 45 44 43 42 41 40 39 38 37 36 35 34 33 32 are applicable.

```
fas6200> fcadmin device_map
```

```
Loop Map for channel 0c:
```

```
...
```

```
Shelf mapping:
```

```
Shelf 3: 45 44 43 42 41 40 39 38 37 36 35 34 33 32
```

```
Shelf 4: 77 76 75 74 73 72 71 70 69 68 67 66 65 64
```

```
Shelf 5: 93 92 91 90 89 88 87 86 85 84 83 82 81 80
```

```
...
```

- b. Turn on the LEDs for the disk drives you identified in Substep a:

**led\_on disk\_name**

You must be in advanced privilege level to enter this command.

The fault LED on the front of the disk drive illuminates—solid. Additionally, if you have any failed disk drives in the disk shelves you are removing, the activity LED on the front of those disk drives blinks.

It is recommended that you turn on the LED for a minimum of four disk drives so that the disk shelves you are removing can be visually identified. You must repeat the command for each disk drive.

To turn on the fault LED for disk drive 0c.45 in disk shelf 3 identified in Substep a, you enter **led\_on 0c.45**

5. If you are removing an entire loop of disk shelves, complete the following substeps; otherwise, go to the next step:

- a. Remove all cables on path A and path B.  
 This includes controller-to-shelf cables, switch-to-shelf cables, and shelf-to-shelf cables for all disk shelves in the loop you are removing
- b. Go to Step 9.

**6.** If you are removing one or more disk shelves from a loop (but keeping the loop), recable the applicable path A loop connections to bypass the disk shelves you are removing by completing the applicable set of substeps:

If you are removing more than one disk shelf, complete the applicable set of substeps one disk shelf at a time.

If you need a graphical system cabling reference, use the platform specific *Installation and Setup Instructions* document that ships with each platform, or access these documents on the NetApp Support Site at [support.netapp.com](http://support.netapp.com) by searching on your specific platform. For example, to find the *Installation and Setup Instructions* document for FAS3200 systems, search on "FAS3200 series".

| If you are removing...          | Then...                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| The first disk shelf in a loop  | <ul style="list-style-type: none"> <li>a. Remove the cable connecting the module A output port of the first disk shelf and the module A input port of the second disk shelf in the loop and set it aside.</li> <li>b. Move the cable connecting the controller or switch to the module A input port of the first disk shelf to the module A input port of the second disk shelf in the loop</li> </ul>                                                                                                                                                                                                                                                  |
| An interim disk shelf in a loop | <ul style="list-style-type: none"> <li>a. Remove the cable connecting the module A output port of the disk shelf being removed and the module A input port of the next disk shelf in the loop and set it aside.</li> <li>b. Move the cable connecting the module A input port of the disk shelf being removed to the module A input port of the next disk shelf in the loop</li> </ul>                                                                                                                                                                                                                                                                  |
| The last disk shelf in a loop   | <ul style="list-style-type: none"> <li>a. Remove the cable connecting the module A input port of the last disk shelf and the module A output port of the previous disk shelf in the loop and set it aside.</li> <li>b. If you have a non fabric MetroCluster system, move the cable connecting the controller to the module A output port of the last disk shelf to the module A output port of the previous disk shelf in the loop</li> </ul> <p><b>Note:</b> For fabric MetroCluster systems, the last disk shelf in a loop is not connected to the switch; therefore, removing the last disk shelf involves removing shelf-to-shelf cables only.</p> |

**7.** Verify that the cabling on path A has successfully bypassed the disk shelves you are removing and all disk drives on the disk shelves you are removing are still connected through path B, by entering the following command from the console of either controller:

```
storage show disk -p
```

It might take up to a minute for the system to complete discovery.

In this example of how the output should appear for a stretch MetroCluster system, the disk shelf being removed is disk shelf 3. One line item appears for each disk drive connected through path B (now the primary path); therefore, the disk drives are listed in the PRIMARY column and B appears in the first PORT column. There is no connectivity through path A for any of the disk drives in the disk shelf being removed; therefore, no information is shown in the SECONDARY or second PORT columns:



```
fas6200> storage show disk -p
PRIMARY PORT SECONDARY PORT SHELF BAY

...
0d.64 B 3 0
0d.65 B 3 1
0d.66 B 3 2
0d.67 B 3 3
0d.68 B 3 4
0d.69 B 3 5
0d.70 B 3 6
0d.71 B 3 7
...
```



**Attention:** If the output shows anything other than all the disk drives connected only through path B, you must correct the cabling by repeating Step 6.

8. Complete the following substeps:
  - a. Repeat Step 6 and Step 7 for path B.
  - b. Repeat Step 1 to confirm that your system configuration is the same as before you began this procedure.
  - c. Go to the next step.
9. If, when you removed ownership from the disk drives as part of the preparation for this procedure, you disabled disk autoassignment, then reenable disk autoassignment by entering the following command; otherwise, go to the next step:

```
options disk.auto_assign on
```

Enter the applicable command from the console of each controller.

10. Power off the disk shelves you disconnected and unplug the power cords from the disk shelves.
11. Remove the disk shelves from the rack or cabinet.

To make disk shelves lighter and easier to maneuver, remove the power supplies and modules. Avoid removing the disk drives if possible, because excessive handling can lead to internal damage.
12. If you are removing disk shelves at the remote site, repeat this procedure.

### Related concepts

[Best practices for changing module types](#) on page 104

If you plan to change the type of any module in your HA pair, make sure that you review the best practice guidelines.

### Related tasks

[Hot-swapping a module](#) on page 106

You can hot-swap a faulty disk shelf module, removing the faulty module and replacing it without disrupting data availability.

[Determining path status for your HA pair](#) on page 104

If you want to remove a module from your HA pair, you need to know if the path you plan to disrupt is redundant. You can use the `storage show disk -p` command to indicate whether the disks have redundant paths.

## Overview of a MetroCluster configuration

---

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.

### Related reference

[Cause and effect for a MetroCluster configuration to failover to another site](#) on page 16  
Failover events cause a controller failover in HA pairs. The storage system responds differently depending on the event and the type of HA pair.

## Types of MetroCluster configurations

The MetroCluster configuration is of two types—Stretch MetroCluster configuration and Fabric-attached MetroCluster configuration.

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

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

## How mirroring works in a MetroCluster configuration

MetroCluster configurations provide data mirroring and the additional ability to initiate a failover if an entire site becomes lost or unavailable. The MetroCluster configuration uses SyncMirror to build a system that can continue to serve data even after complete loss of one of the sites.

Data consistency is retained, even when the data is contained in more than one aggregate.

Like mirrored HA pairs, MetroCluster configurations contain two complete copies of the specified data volumes or file systems that you indicated as being mirrored volumes or file systems in your HA pair. 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 or array LUNs.

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

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

**Note:** If you are using non-mirrored aggregates in a MetroCluster configuration, the plex must consist of disks from Pool0. This is because using Pool1 disks for non-mirrored aggregate might lead to unexpected takeovers due to loss of connectivity between the two sites.

For more information about using the SyncMirror to mirror data, see the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.

## How Data ONTAP works with MetroCluster configurations

Data ONTAP divides storage across physically separated pools of disks. During configuration, Data ONTAP identifies spare disks and divides them into separate groups 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 the 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.

## Advantages of stretch MetroCluster configurations

MetroCluster configurations provide the same advantages of mirroring as mirrored HA pairs, with the additional ability to initiate failover if an entire site becomes lost or unavailable.

For MetroCluster configuration on disks, the advantages of a stretch MetroCluster configuration are as follows:

- Your data is protected if there is a failure or loss of two or more disks in a RAID 4 aggregate or three or more disks in a RAID-DP aggregate.
- The failure of an FC-AL adapter, loop, or IOM module does not require a failover.

For MetroCluster configurations with systems using array LUNs, the advantages of a stretch MetroCluster configuration are as follows:

- For RAID 0 aggregates, your data is protected if there is a failure of a storage array or loss of an array LUN on one of the storage arrays in the MetroCluster configuration.

In addition, a MetroCluster configuration 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

[Recovering from a disaster by using MetroCluster configurations](#) on page 202

In situations such as prolonged power outages or natural disasters, you can use the 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.

### Related reference

[Failover event cause-and-effect table](#) on page 16

Failover events cause a controller failover in HA pairs. The storage system responds differently depending on the event and the type of HA pair.

## Configuration variations for stretch MetroCluster configurations

Stretch MetroCluster configurations have asymmetrical and active/passive variations.

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

- 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) or failure of a single array LUN in a RAID 0 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 HA pair .

- Active/passive MetroCluster configurations

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 licenses for a MetroCluster configuration installed so that remote takeover is possible.

## Advantages of fabric-attached MetroCluster configurations

Fabric-attached MetroCluster configurations 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 configuration over a stretch MetroCluster configuration include the following:

- The two halves of the configuration can be more than 500 meters apart, which provides increased disaster protection.
- For fabric-attached MetroCluster configurations with FAS systems, 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.
- For fabric-attached MetroCluster configurations with systems using array LUNs, the storage arrays 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.

### Related concepts

[Recovering from a disaster by using MetroCluster configurations](#) on page 202

In situations such as prolonged power outages or natural disasters, you can use the 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.

### Related reference

[Failover event cause-and-effect table](#) on page 16

Failover events cause a controller failover in HA pairs. The storage system responds differently depending on the event and the type of HA pair.

## Configuration variations for fabric-attached MetroCluster configurations

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

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

- 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.
- If the system contains unmirrored aggregates using disks from remote site, issues in ISLs can lead to multiple disk failure. Therefore, if you see ISL instability, such as ISL going down, switch failures and so on, you must take such unmirrored aggregates offline.

**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 HA configuration.
  - If the ISL links are being taken down for maintenance and your system contains unmirrored aggregates, you must first take all unmirrored aggregates offline.
- Active/passive MetroCluster configurations  
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 licenses for MetroCluster configuration installed so that remote takeover is possible.

## Implementing stretch MetroCluster configurations with disks

Implementing a stretch MetroCluster configuration involves careful planning and systematic cabling. You must be aware of requirements and configuration procedures that are specific to stretch MetroCluster configuration.

Implementing the stretch MetroCluster configuration involves the following tasks:

- Planning your stretch MetroCluster configuration
- Cabling the stretch MetroCluster configuration
- Managing the default configuration speed

### Planning a stretch MetroCluster configuration using disks

Creating a detailed plan for your MetroCluster configuration helps you understand the unique requirements for configuring a stretch MetroCluster configuration with disk. For installing a stretch MetroCluster configuration, you should be aware of the requirements (such as the tools and equipment) and restrictions for a stretch MetroCluster configuration.

### Setup requirements and restrictions for stretch MetroCluster configurations with disks

You must follow certain requirements and restrictions when setting up a new stretch MetroCluster configuration with disks.

The requirements and restrictions for stretch MetroCluster configurations include those for a standard HA pair and those for a mirrored HA pair. In addition, the following requirements apply:

- Starting with Data ONTAP 8.1 for FAS systems and 8.2 for V-Series systems, stretch MetroCluster configurations support SAS disk shelves when used with FibreBridge 6500N bridges.
- Your storage system must meet all the compatibility requirements for FibreBridge 6500N bridges in the Interoperability Matrix at [support.netapp.com/matrix](http://support.netapp.com/matrix).
- The stretch MetroCluster configurations support SAS disk shelves using SAS optical cables. For more information about Data ONTAP versions supported for SAS optical cables with MetroCluster configurations, see the *Hardware Universe* at [hwu.netapp.com](http://hwu.netapp.com).
- Your storage system must meet all the compatibility requirements for SAS optical cables in the Interoperability Matrix at [support.netapp.com/matrix](http://support.netapp.com/matrix).
- SAS, SATA, and Fibre Channel storage is supported on stretch MetroCluster configurations, but both plexes of the same aggregate must use the same type of storage.
- For a stretch MetroCluster configuration that uses SAS disk shelves, SAS copper cables are supported up to a distance of 20 meters.
- For a stretch MetroCluster configuration that uses SAS disk shelves, SAS optical cables are supported up to a distance of 0 to 500 meters.
- For stretch MetroCluster configurations using SAS disk shelves, either FibreBridge 6500N or SAS optical cables are required.

**Note:** For stretch MetroCluster configurations using SAS disk shelves up to a distance of 20 meters, you can use SAS copper cables to connect the disk shelves directly to the controller.

- For FAS systems, stacks of SAS disk shelves can be added to a MetroCluster configuration that has DS14mk4 FC disk shelves.
- A stack of SAS disk shelves can contain shelves of SAS disk drives and shelves of SATA disk drives, but each SAS disk shelf can contain only SAS or SATA disk drives; you cannot mix SAS and SATA disk drives in the same disk shelf.

To know about the number and the type of disk shelves supported in a stack, see the Interoperability Matrix on [support.netapp.com/matrix](http://support.netapp.com/matrix).

- For stretch MetroCluster configurations using SAS disk shelves and FibreBridge 6500N bridge, each stack requires two Fibre Channel ports on each controller.
- While using SAS optical cables (when used with SAS disk shelves), you must ensure that there are sufficient SAS ports available on the storage system.
- For the number of SAS disk shelves and the types of SAS disk shelves supported in a stack, see the Interoperability Matrix on [support.netapp.com/matrix](http://support.netapp.com/matrix).
- For a stretch MetroCluster configuration, if you have a 16-Gb FC-VI adapter with 16-Gb long range SFP connected by a 9/125µm single-mode fiber (SMF) cable, you can go up to 500m distance between nodes at the default 16-Gbps adapter speed.
- The following distance limitations dictate the default speed you can set:
  - If the distance between the nodes is 150m and you have an 8-Gb FC-VI adapter, the default speed is set to 8 Gb.  
If you want to increase the distance to 270m or 500m, you can set the default speed to 4 Gb or 2 Gb, respectively.
  - If the distance between nodes is between 150m and 270m and you have an 8-Gb FC-VI adapter, you can set the default speed to 4 Gb.
  - If the distance between nodes is between 270m and 500m and you have an 8-Gb FC-VI or 4-Gb FC-VI adapter, you can set the default speed to 2 Gb.

For more information about the distance limitation and the FC-VI adapter, see TR-3548, *Best Practices for MetroCluster Design and Implementation*, at the NetApp Support Site.

- The following Data ONTAP options must be enabled on both nodes:
  - `cf.mode`  
This option must be set to **ha**.
  - `cf.remote_syncmirror.enable`  
This option must be set to **on**.

See the Interoperability Matrix on [support.netapp.com/matrix](http://support.netapp.com/matrix) for more information about hardware and firmware requirements for this configuration.

### Related concepts

[Setup requirements and restrictions for standard HA pairs](#) on page 24

You must follow certain requirements and restrictions when setting up a new standard HA pair. These requirements help you ensure the data availability benefits of the HA pair design.

[Setup requirements and restrictions for mirrored HA pairs](#) on page 25

The restrictions and requirements for mirrored HA pairs include those for a standard HA pair with these additional requirements for disk pool assignments and cabling.

## Hardware requirements for using disks in a MetroCluster configuration with V-Series systems

Starting in Data ONTAP 8.2, a MetroCluster configuration with V-Series systems can include native disks only, array LUNs on storage arrays only, or both. There are specific hardware requirements to set up a MetroCluster with V-Series systems with native disks.

### Hardware requirements if you are using FibreBridge 6500N bridges

The following hardware is required for using native disks in a MetroCluster configuration with V-Series systems:

- A pair of V-Series systems
- FibreBridge 6500N bridges (Data ONTAP 8.2 and later) or, for stretch MetroCluster configurations only, SAS optical cables (Data ONTAP 8.2.1 and later)

**Note:** For V-Series systems, support for SAS optical cables does not start until Data ONTAP 8.2.1.

Prior to Data ONTAP 8.2, MetroCluster configurations with V-Series systems could not include V-Series systems that used disks. In Data ONTAP 8.2, for a MetroCluster configuration to include both disks and array LUNs, all systems in the configuration must be V-Series systems. Platforms released in 8.2.1 and later can include both array LUNs and disks in a MetroCluster configuration. However, no MetroCluster configurations can include both V-Series systems and FAS systems.

The following hardware is required for using native disks in a MetroCluster configuration with V-Series systems:

- Four FibreBridge 6500N bridges: two FibreBridge 6500N bridges per stack of disk shelves, distributed across two sites  
The disk shelves in the MetroCluster configuration are those supported by the FibreBridge 6500N bridges.
- Supported firmware for the FibreBridge 6500N bridges, as shown in the Interoperability Matrix at [support.netapp.com](http://support.netapp.com)
- For a fabric-attached MetroCluster configuration, four switches  
The switches must be designated in the Interoperability Matrix at [support.netapp.com](http://support.netapp.com) as being supported for a MetroCluster configuration.

**Note:** If the fabric-attached MetroCluster configuration also accesses LUNs on the storage arrays, these same switches are used for array LUN access as well. A stretch MetroCluster configuration that also uses array LUNs would use four switches to access the array LUNs, but the switches would not be used to access the disks.

See the *Configuring a MetroCluster system with SAS disk shelves and FibreBridge 6500N bridges* for information about how to connect the FibreBridge 6500N and the disk shelves.

### Required types of connections between a V-Series system and a FibreBridge 6500N bridge

| For this MetroCluster type... | The connection type is... |
|-------------------------------|---------------------------|
| Fabric-attached               | Switches                  |
| Stretch                       | Direct FC connections     |

### Hardware requirements if you are using SAS optical cables

SAS optical cables can be used to cable SAS disk shelves in a stretch MetroCluster system to achieve greater connectivity over distance than is possible with SAS copper cables. SAS optical cables can be used in a new system installation or in an existing system in which you want to replace the SAS copper cables and FibreBridge 6500N bridges. SAS optical cables are supported for use with V-Series systems starting in Data ONTAP 8.2.1.

See *Configuring a stretch MetroCluster system with SAS disk shelves and SAS optical cables* on the NetApp Support Site at [support.netapp.com](http://support.netapp.com) for hardware requirements and procedures for installing a new system with SAS optical cables or replacing FibreBridge 6500N bridges with SAS optical cables.

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

You can use a UPS (Uninterruptible Power Supply) with your MetroCluster configurations. 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.

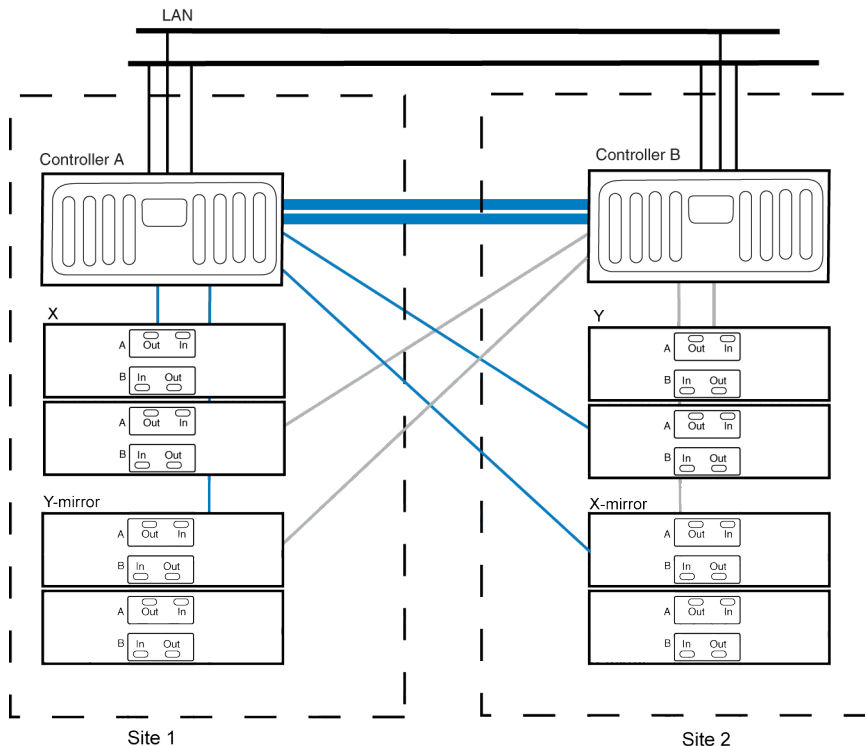
You can leave the disks on the regular power supply. In this case, if power is interrupted to one site, the controller can access the other plex until it shuts down or power is restored. If, however, power is interrupted to both sites at the same time and the disks are not connected to the UPS, the MetroCluster configuration cannot shut down gracefully.

## Stretch MetroCluster configuration with disks

You can configure a stretch MetroCluster configuration 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 using DS14mk2 or DS14mk4 disk shelves. For an example of a stretch MetroCluster configuration using SAS disk shelves and FibreBridge 6500N bridge, see *Configuring a MetroCluster system with SAS disk shelves and FibreBridge 6500N bridges* on the NetApp Support Site. For an example of a stretch MetroCluster configuration using SAS disk shelves and SAS optical cables, see the *Configuring a stretch MetroCluster system with SAS disk shelves and SAS optical cables*. 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.

**Related information**

[NetApp Support Site: support.netapp.com](http://support.netapp.com)

**Required documentation, tools, and equipment**

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

**Required documentation for MetroCluster configurations**

You must refer to some of the flyers and guides that are required to install a new MetroCluster configuration, or convert two stand-alone systems into a MetroCluster 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 NetApp Support 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:

| Manual name                                                          | Description                                                                                              |
|----------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------|
| The appropriate system cabinet guide                                 | Describes how to install NetApp equipment into a system cabinet                                          |
| <i>Site Requirements Guide</i>                                       | Describes the physical requirements your site must meet to install NetApp equipment                      |
| The appropriate disk shelf guide                                     | Describes how to cable a disk shelf to a storage system                                                  |
| The appropriate hardware documentation for your storage system model | Describes how to install the storage system, connect it to a network, and bring it up for the first time |

| Manual name                                                                                   | Description                                                                                                           |
|-----------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------|
| <i>Diagnostics Guide</i>                                                                      | Describes the diagnostics tests that you can run on the storage system                                                |
| <i>Data ONTAP Upgrade and Revert/Downgrade Guide for 7-Mode</i>                               | Describes how to upgrade the storage system and disk firmware, and how to upgrade the storage system software         |
| <i>Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode</i>                 | Describes, among other topics, SyncMirror technology, which is used for mirrored HA pairs                             |
| <i>Data ONTAP System Administration Guide for 7-Mode</i>                                      | Describes general storage system administration                                                                       |
| <i>Data ONTAP Software Setup Guide for 7-Mode</i>                                             | Describes how to configure the software of a new storage system for the first time                                    |
| <i>Fabric-attached MetroCluster Systems Brocade Switch Configuration Guide</i>                | Describes how to configure Brocade switches for a fabric-attached MetroCluster configuration                          |
| <i>Configuring a MetroCluster system with SAS disk shelves and FibreBridge 6500N bridges</i>  | Describes how to install FibreBridge 6500N bridges as part of your MetroCluster configuration                         |
| <i>Fabric-attached MetroCluster Systems Cisco Switch Configuration Guide</i>                  | Describes how to configure Cisco switches for a fabric-attached MetroCluster configuration                            |
| <i>FlexArray Virtualization Installation Requirements and Reference Guide</i>                 | Describes how to prepare array LUNs on storage arrays for use by Data ONTAP                                           |
| <i>Configuring a stretch MetroCluster system with SAS disk shelves and SAS optical cables</i> | Describes how to install SAS disk shelves using SAS optical cables as part of your stretch MetroCluster configuration |

**Related information**

*NetApp Support Site: [support.netapp.com](http://support.netapp.com)*

**Required tools**

There are some tools that you need to install the MetroCluster configuration.

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

**Required equipment for MetroCluster configurations**

You should receive the required equipment, such as storage system, HA interconnect adapter, and so on, to install a MetroCluster configuration.

See the *Hardware Universe* at [hwu.netapp.com](http://hwu.netapp.com) to confirm your storage system type, storage capacity, and so on.

**Note:** For fabric-attached MetroCluster configurations, use the information in the *Hardware Universe* labeled for MetroCluster configurations. For stretch MetroCluster configurations, use the information in the *Hardware Universe* labeled “for HA Environments.”

| Required equipment | Stretch MetroCluster configuration                                                       | Fabric-attached MetroCluster configuration |
|--------------------|------------------------------------------------------------------------------------------|--------------------------------------------|
| Storage system     | Two of the same type of storage systems                                                  |                                            |
| Storage            | See the <i>Hardware Universe</i> at <a href="http://hwu.netapp.com">hwu.netapp.com</a> . |                                            |

| Required equipment                                                | Stretch MetroCluster configuration                                                                                                                                                                                                                                                                                                                                                            | Fabric-attached MetroCluster configuration                                                                                                                                                                                  |
|-------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FibreBridge 6500N bridges (if you are attaching SAS disk shelves) | Two FibreBridges are required for each stack of SAS disk shelves.                                                                                                                                                                                                                                                                                                                             |                                                                                                                                                                                                                             |
| HA interconnect adapter                                           | InfiniBand adapter (Required only for systems that do not use an NVRAM5 or NVRAM6 adapter, which functions as the HA interconnect adapter.)<br><br>FC-VI adapter (Required for the 31xx, 32xx, and 62xx dual-controller systems.)<br><br><b>Note:</b> When the FC-VI adapter is installed in a 31xx, 32xx, or 62xx system, the internal InfiniBand interconnect is automatically deactivated. | FC-VI adapter                                                                                                                                                                                                               |
| FC-AL or FC HBA (FC HBA for Disk) adapters                        | Two or four Fibre Channel HBAs<br><br><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.                                                                                                                           |                                                                                                                                                                                                                             |
| Fibre Channel switches                                            | Stretch MetroCluster configurations with array LUNs require these switches.                                                                                                                                                                                                                                                                                                                   | Two pairs of Brocade or Cisco switches<br><br><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                                | Not applicable                                                                                                                                                                                                                                                                                                                                                                                | Two or four SFPs are required per fabric, depending on whether you are using single or 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                                                                                                                                                                                                                                                                                                                                                                   | Not applicable                                                                                                                                                                                                              |

| Required equipment                                     | Stretch MetroCluster configuration                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Fabric-attached MetroCluster configuration                                                                                                                                                                                                                                                                                             |
|--------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 or FibreBridge 6500N</li> <li>• Two SC/LC IB HA adapter cables</li> <li>• Four SC/LC or LC/LC cables</li> <li>• If you are using SAS optical cables with your SAS disk shelves, you would have received the cable qualities and lengths you ordered for your configuration.</li> </ul> <p><b>Note:</b> For information about required cables, see the Interoperability Matrix at <a href="http://support.netapp.com/matrix">support.netapp.com/matrix</a>.</p> | <ul style="list-style-type: none"> <li>• LC/LC controller-to-switch cables</li> <li>• SC/LC (for DS14) or LC/LC (for DS14mk4 FC) disk shelf-to-switch cables or switch to FibreBridge 6500N</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> |

**Related information**

*NetApp Support Site: [support.netapp.com](http://support.netapp.com)*

## Cabling a stretch MetroCluster configuration with disks

The process of cabling a stretch MetroCluster configuration is the same as a mirrored HA pair . However, your systems must meet the requirements for a stretch MetroCluster configuration.

**About this task**

If you are installing the FibreBridge 6500N bridge as part of your MetroCluster configuration, see the *Configuring a MetroCluster system with SAS disk shelves and FibreBridge 6500N bridges* on the NetApp Support Site at [support.netapp.com](http://support.netapp.com) for cabling procedures.

If you are installing SAS disk shelves with SAS optical cables as part of your stretch MetroCluster configuration, see the *Configuring a stretch MetroCluster system with SAS disk shelves and SAS optical cables* on the NetApp Support Site at [support.netapp.com](http://support.netapp.com) for cabling procedures.

**Related concepts**

*Configuring an HA pair* on page 50

Bringing up and configuring a standard or mirrored HA pair for the first time can require enabling HA mode capability and failover, setting options, configuring network connections, and testing the configuration.

*Setup requirements and restrictions for stretch MetroCluster configurations with disks* on page 118

You must follow certain requirements and restrictions when setting up a new stretch MetroCluster configuration with disks.

*Recovering from a disaster by using MetroCluster configurations* on page 202

In situations such as prolonged power outages or natural disasters, you can use the 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.

**Related tasks**

*Cabling a mirrored HA pair* on page 41

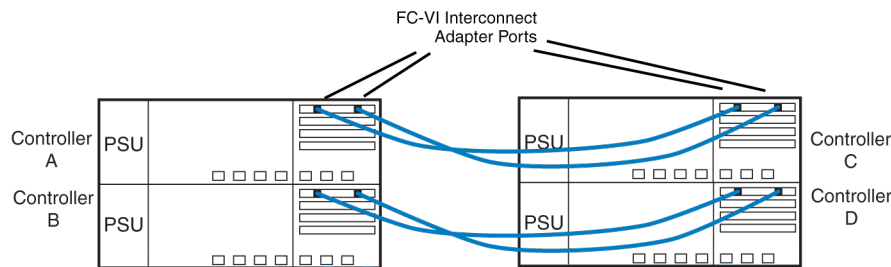
To cable a mirrored HA pair, you identify the ports you need to use on each node, cable those ports, and then cable the HA interconnect.

### Stretch MetroCluster configuration on single-enclosure HA pairs

You can configure two stretch MetroCluster configurations between a pair of single-enclosure HA pair systems. In this configuration, the HA pair between the two controllers in each chassis is deactivated, and two separate, side-by-side stretch MetroCluster configurations are formed between the four controllers.

To implement the stretch MetroCluster configuration, you must install an FC-VI adapter in each controller to provide the HA 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 MetroCluster configurations, which use NVRAM adapters to provide the interconnect.

The following figure shows a stretch MetroCluster configuration on single-enclosure HA pair systems:



#### Related tasks

[Assigning disk pools in a stretch MetroCluster configuration](#) on page 126

After completing the cabling of a stretch MetroCluster configuration, you must assign the attached disk shelves to the appropriate disk pools.

### Assigning disk pools in a stretch MetroCluster configuration

After completing the cabling of a stretch MetroCluster configuration, you must assign the attached disk shelves to the appropriate disk pools.

#### About this task

You can explicitly assign disks on the attached disk shelves to the appropriate pool with the `disk assign` command.

For the purpose of this procedure, Node A and Node B constitute a stretch MetroCluster configuration having one stack of four disk shelves at each site. Each shelf stack is cabled in multipath HA using quad-port fibre channel adapter in slot 2 of each node. 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... | At site... | Belongs to... | And is assigned to that node's... |
|---------------|------------|---------------|-----------------------------------|
| 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.
2. Assign the local disks to Node A pool 0 by entering the following command at the console:
 

```
disk assign -shelf 2a.shelf1 -p 0
disk assign -shelf 2a.shelf2 -p 0
```

This indicates that the disks in disk shelf 1 and disk shelf 2 attached to port 2a are assigned to pool 0 of Node A.

**Note:** The `-shelf` option allows you to specify a particular shelf connected to a channel for assigning all disks present in that particular shelf.
3. Assign the remote disks to Node A pool 1 by entering the following command at the console:
 

```
disk assign -shelf 2b.shelf7 -p 1
disk assign -shelf 2b.shelf8 -p 1
```

This indicates that all the disks in disk shelf 7 and disk shelf 8 attached to port 2b are assigned to pool 1 of Node A.
4. Boot Node B into Maintenance mode.
5. Assign the local disks to Node B pool 0 by entering the following command at the console:
 

```
disk assign -shelf 2a.shelf5 -p 0
disk assign -shelf 2a.shelf6 -p 0
```

This indicates that all the disks in disk shelf 5 and disk shelf 6 attached to port 2a are assigned to pool 0 of Node B
6. Assign the remote disks to Node B pool 1 by entering the following command at the console:
 

```
disk assign -shelf 2b.shelf3 -p 1
disk assign -shelf 2b.shelf4 -p 1
```

This indicates that all the disks in disk shelf 3 and disk shelf 4 attached to port 2b are assigned to pool 1 of Node B.

### After you finish

Proceed to verify the disk paths on the system.

### Related concepts

[Stretch MetroCluster configuration on single-enclosure HA pairs](#) on page 126

You can configure two stretch MetroCluster configurations between a pair of single-enclosure HA pair systems. In this configuration, the HA pair between the two controllers in each chassis is

deactivated, and two separate, side-by-side stretch MetroCluster configurations are formed between the four controllers.

## Managing the default configuration speed of a stretch MetroCluster configuration

You can manage the default configuration speed of a stretch MetroCluster configuration either by changing the default speed or by resetting the default configuration speed.

### Changing the default configuration speed of a stretch MetroCluster configuration

The distance between your nodes and the FC-VI adapter speed dictates the default configuration speed of your stretch MetroCluster configuration. If the distance between the nodes is greater than what is supported with the default configuration speed, you must change the default configuration speed of the FC-VI adapter.

#### Before you begin

The stretch MetroCluster system's default configuration speed must conform to the stretch MetroCluster system setup requirements and restrictions.

#### About this task

You must perform these steps at both the nodes in a MetroCluster configuration if they are configured at different speeds.

#### Steps

1. At the storage console prompt, halt the storage system by entering the following command:

```
halt -f
```

2. Reset the configuration speed by entering the following commands.



---

| <b>If you want to set the speed to...</b>                                                            | <b>Then...</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 8 Gbps<br><br><b>Note:</b> You can set the speed to 8 Gbps if you have 16G FC-VI card with 16 G SFP. | <p>a. Enter the following command:</p> <pre>setenv ispfcvi-force-8G-only? True</pre> <p>b. Verify that your storage system is configured for 8 Gbps by entering the following command:</p> <pre>printenv ispfcvi-force-8G-only?</pre> <p>The storage system console displays output similar to the following:</p> <pre>Variable Name          Value ----- ispfcvi-force-8G-only? true</pre> <p>c. If you have modified the speed to 4 or 2 Gbps, undo the modified speed setting by entering the following command:</p> <pre>unsetenv ispfcvi-force-4G-only? unsetenv ispfcvi-force-2G-only?</pre> <p>d. Verify that your system is unconfigured for 4 or 2 Gbps by entering the following command:</p> <pre>printenv ispfcvi-force-4G-only? printenv ispfcvi-force-2G-only?</pre> <p>The storage system console displays output similar to the following:</p> <pre>Variable Name          Value ----- ispfcvi-force-4G-only? *** Undefined ***</pre> |

---

---

| If you want to set the speed to... | Then... |
|------------------------------------|---------|
|------------------------------------|---------|

---

4 Gbps

- a. Enter the following command:  
`setenv ispf cvi-force-4G-only? True`
- b. Verify that your storage system is configured for 4 Gbps by entering the following command:

`printenv ispf cvi-force-4G-only?`

The storage system console displays output similar to the following:

```
Variable Name Value

ispf cvi-force-4G-only? true
```

- c. If you previously modified the speed to 2 or 8 Gbps, undo the modified speed setting entering the following command:

`unsetenv ispf cvi-force-2G-only?`

`unsetenv ispf cvi-force-8G-only?`

- d. Verify that your system is unconfigured for 2 or 8 Gbps by entering the following command:

`printenv ispf cvi-force-2G-only?`

`printenv ispf cvi-force-8G-only?`

The storage system console displays output similar to the following:

```
Variable Name Value

ispf cvi-force-2G-only? *** Undefined ***
```

---

---

| If you want to set the speed to... | Then...                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |               |       |                        |      |               |       |                        |                   |
|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|-------|------------------------|------|---------------|-------|------------------------|-------------------|
| 2 Gbps                             | <p data-bbox="483 317 846 344">a. Enter the following command:</p> <pre data-bbox="521 365 1008 392">setenv ispfcvi-force-2G-only? True</pre> <p data-bbox="483 401 1398 459">b. Verify that your storage system is configured for 2 Gbps by entering the following command:</p> <pre data-bbox="521 480 967 508">printenv ispfcvi-force-2G-only?</pre> <p data-bbox="521 516 1471 575">If your storage system is configured correctly, the system console displays output similar to the following:</p> <table border="1" data-bbox="521 600 1455 701"><thead><tr><th>Variable Name</th><th>Value</th></tr></thead><tbody><tr><td>ispfcvi-force-2G-only?</td><td>true</td></tr></tbody></table> <p data-bbox="483 709 1430 768">c. If you previously modified the default speed to 4 or 8 Gbps, undo the modified speed setting by entering the following command:</p> <pre data-bbox="521 789 967 816">unsetenv ispfcvi-force-4G-only?</pre> <pre data-bbox="521 837 967 865">unsetenv ispfcvi-force-8G-only?</pre> <p data-bbox="483 873 1471 932">d. Verify that your storage system is unconfigured for 4 or 8 Gbps by entering the following command:</p> <pre data-bbox="521 953 967 980">printenv ispfcvi-force-4G-only?</pre> <pre data-bbox="521 1001 967 1029">printenv ispfcvi-force-8G-only?</pre> <p data-bbox="521 1037 1247 1064">The storage system console displays output similar to the following:</p> <table border="1" data-bbox="521 1089 1455 1190"><thead><tr><th>Variable Name</th><th>Value</th></tr></thead><tbody><tr><td>ispfcvi-force-4G-only?</td><td>*** Undefined ***</td></tr></tbody></table> | Variable Name | Value | ispfcvi-force-2G-only? | true | Variable Name | Value | ispfcvi-force-4G-only? | *** Undefined *** |
| Variable Name                      | Value                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |               |       |                        |      |               |       |                        |                   |
| ispfcvi-force-2G-only?             | true                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |               |       |                        |      |               |       |                        |                   |
| Variable Name                      | Value                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |               |       |                        |      |               |       |                        |                   |
| ispfcvi-force-4G-only?             | *** Undefined ***                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |               |       |                        |      |               |       |                        |                   |

---

3. Boot the storage system by entering the following command:

```
boot_ontap
```

### Resetting a stretch MetroCluster configuration to the default speed

If you modified the default configuration speed in a stretch MetroCluster configuration by using an FC-VI adapter, you can reset the speed to the default configuration speed by using the `unsetenv` command at the boot environment prompt.

#### About this task

- When you undo the previously configured speed, the speed is set to default automatically.
- You must perform the steps at both the nodes in a MetroCluster configuration.

#### Steps

1. At the storage prompt, halt the system by entering the following command:

```
halt -f
```

2. Reset the configuration speed.

---

| If you want to reset the speed from... | Then...                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 8 Gbps                                 | <p>a. Enter the following command:</p> <pre>unsetenv ispfcvi-force-8G-only?</pre> <p>b. Verify that your system is unconfigured for 8 Gbps by entering the following command:</p> <pre>printenv ispfcvi-force-8G-only?</pre> <p>The system console displays output similar to the following:</p> <pre>Variable Name      Value ----- ispfcvi-force-8G-only? *** Undefined ***</pre> |
| 4 Gbps                                 | <p>a. Enter the following command:</p> <pre>unsetenv ispfcvi-force-4G-only?</pre> <p>b. Verify that your system is unconfigured for 4 Gbps by entering the following command:</p> <pre>printenv ispfcvi-force-4G-only?</pre> <p>The system console displays output similar to the following:</p> <pre>Variable Name      Value ----- ispfcvi-force-4G-only? *** Undefined ***</pre> |
| 2 Gbps                                 | <p>a. Enter the following command:</p> <pre>unsetenv ispfcvi-force-2G-only?</pre> <p>b. Verify that your system is unconfigured for 2 Gbps by entering the following command:</p> <pre>printenv ispfcvi-force-2G-only?</pre> <p>The system console displays output similar to the following:</p> <pre>Variable Name      Value ----- ispfcvi-force-2G-only? *** Undefined ***</pre> |

---

3. Boot the storage system by entering the following command:

```
boot_ontap
```

## Implementing fabric-attached MetroCluster configurations with disks

---

Implementing a fabric-attached MetroCluster configuration involves careful planning and systematic cabling.

Implementing a fabric-attached MetroCluster configuration involves the following tasks:

- Planning your implementation
- Cabling the disk shelves and various other components such as switches and bridges
- Setting the preferred primary ports
- Verifying the implementation

### Planning a fabric-attached MetroCluster configuration using disks

Creating a detailed plan for your fabric-attached MetroCluster configuration helps you understand the unique requirements for configuring a fabric-attached MetroCluster configuration with disk. Installing a fabric-attached MetroCluster configuration involves connecting and configuring a number of devices, which might be done by different people.

### Setup requirements and restrictions for fabric-attached MetroCluster configurations with disks

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

The setup requirements for a fabric-attached MetroCluster configuration with disks include those for standard and mirrored HA configurations, with the following exceptions:

**Note:** See the Interoperability Matrix at [support.netapp.com/matrix](http://support.netapp.com/matrix) 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.
  - 31xx systems
  - 32xx systems
  - 62xx systems
  - 6080 systems
  - 6040 systems
  - FAS80xx
- Each node requires an FC-VI (Fibre Channel-Virtual Interface) adapter; the type and the slot position is dependent on the controller model.
- The 8-Gbps FC-VI adapter is supported only on the 32xx, 62xx, and FAS80xx systems.
- Starting with Data ONTAP 8.2.1, 16-Gbps FC-VI adapter is supported for the following systems:
  - 3220
  - 3250
  - 62xx
  - FAS80xx

**Note:** For information about supported cards and slot placement, see the *Hardware Universe* on the NetApp Support Site.

- If you want to convert a stretch MetroCluster configuration to a fabric-attached MetroCluster configuration, and you have modified the default configuration speed of FC-VI adapters by setting the boot environment variables to **True**, you must reset it to default configuration speed before the conversion.

### Disk and disk shelf requirements

- The only Fibre Channel disk shelves supported are DS14mk2 or DS14mk4.
- Starting with Data ONTAP 8.1, fabric-attached MetroCluster configurations also support SAS disk shelves when used with FibreBridge 6500N.
- Each stack of SAS disk shelves requires two FibreBridges.
- Your storage system must meet all the compatibility requirements for FibreBridge 6500N in the Interoperability Matrix at [support.netapp.com/matrix](http://support.netapp.com/matrix).
- DS14mk2 or DS14mk4 disk shelves using SATA drives and AT-FCX storage are not supported.
- You can connect a maximum of two DS14mk2 or DS14mk4 disk shelves to each loop.
- A stack of SAS disk shelves can contain shelves of SAS disk drives and shelves of SATA disk drives, but each SAS disk shelf can only contain SAS or SATA disk drives; you cannot have SAS and SATA disk drives in the same disk shelf.

For more information about the number and the types of disk shelves supported in a stack, see the Interoperability Matrix at [support.netapp.com/matrix](http://support.netapp.com/matrix).

- Stacks of SAS disk shelves can be added to a fabric-attached MetroCluster configuration that has DS14mk2 or DS14mk4 FC disk shelves.
- For the number of SAS disk shelves and the types of SAS disk shelves supported in a stack, see the Interoperability Matrix at [support.netapp.com/matrix](http://support.netapp.com/matrix).

### Capacity limits

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

- The maximum storage capacity for the node  
For the maximum storage capacity, see the *Hardware Universe* on the NetApp Support Site at [support.netapp.com](http://support.netapp.com).
- 840 Fibre Channel disks (60 disk shelves)
- In a MetroCluster configuration (using both FC and SAS disk shelves), the total disk drives count must not exceed storage capacity.

### Fibre Channel switch requirements

- Each site of the MetroCluster configuration requires two switches.
- Switches must be a supported Brocade or Cisco model supplied by NetApp. Customer-supplied switches are not supported.
- The two switches within the fabric must be the same model and must be licensed for the same number of ports.
- All four switches for a particular fabric-attached MetroCluster configuration must support the same maximum speed.
- Switches must be running the correct firmware version.
- Configurations using SAS storage and FibreBridge 6500N bridges supports Brocade or Cisco switch models supplied by NetApp.

See the Interoperability Matrix at [support.netapp.com/matrix](http://support.netapp.com/matrix) for more information about supported switch models.

### Related concepts

[Setup requirements and restrictions for standard HA pairs](#) on page 24

You must follow certain requirements and restrictions when setting up a new standard HA pair. These requirements help you ensure the data availability benefits of the HA pair design.

[Setup requirements and restrictions for mirrored HA pairs](#) on page 25

The restrictions and requirements for mirrored HA pairs include those for a standard HA pair with these additional requirements for disk pool assignments and cabling.

## Configuration limitations for fabric-attached MetroCluster configurations with disks

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

The fabric-attached MetroCluster configuration has the following limitations:

- You cannot use DS14mk4 FC disk shelves that are using SATA and AT-FCX storage.
- You cannot use the switches of the MetroCluster configuration to connect Fibre Channel tape devices, or for FC traffic of any kind; you can connect only system controllers and disk shelves to those switches.

The switches of the MetroCluster configuration either connect DS14mk4 FC disk shelves to the controllers or FibreBridge 6500N bridge (connected to SAS or SATA disk shelves) to the disk shelves.

- You can connect a tape storage area network (SAN) to either of the nodes, but the tape SAN must not use the switches used in a MetroCluster configuration.

## Hardware requirements for using disks in a MetroCluster configuration with V-Series systems

Starting in Data ONTAP 8.2, a MetroCluster configuration with V-Series systems can include native disks only, array LUNs on storage arrays only, or both. There are specific hardware requirements to set up a MetroCluster with V-Series systems with native disks.

### Hardware requirements if you are using FibreBridge 6500N bridges

The following hardware is required for using native disks in a MetroCluster configuration with V-Series systems:

- A pair of V-Series systems
- FibreBridge 6500N bridges (Data ONTAP 8.2 and later) or, for stretch MetroCluster configurations only, SAS optical cables (Data ONTAP 8.2.1 and later)

**Note:** For V-Series systems, support for SAS optical cables does not start until Data ONTAP 8.2.1.

Prior to Data ONTAP 8.2, MetroCluster configurations with V-Series systems could not include V-Series systems that used disks. In Data ONTAP 8.2, for a MetroCluster configuration to include both disks and array LUNs, all systems in the configuration must be V-Series systems. Platforms released in 8.2.1 and later can include both array LUNs and disks in a MetroCluster configuration. However, no MetroCluster configurations can include both V-Series systems and FAS systems.

The following hardware is required for using native disks in a MetroCluster configuration with V-Series systems:

- Four FibreBridge 6500N bridges: two FibreBridge 6500N bridges per stack of disk shelves, distributed across two sites

The disk shelves in the MetroCluster configuration are those supported by the FibreBridge 6500N bridges.

- Supported firmware for the FibreBridge 6500N bridges, as shown in the Interoperability Matrix at [support.netapp.com](http://support.netapp.com)
- For a fabric-attached MetroCluster configuration, four switches  
The switches must be designated in the Interoperability Matrix at [support.netapp.com](http://support.netapp.com) as being supported for a MetroCluster configuration.

**Note:** If the fabric-attached MetroCluster configuration also accesses LUNs on the storage arrays, these same switches are used for array LUN access as well. A stretch MetroCluster configuration that also uses array LUNs would use four switches to access the array LUNs, but the switches would not be used to access the disks.

See the *Configuring a MetroCluster system with SAS disk shelves and FibreBridge 6500N bridges* for information about how to connect the FibreBridge 6500N and the disk shelves.

### Required types of connections between a V-Series system and a FibreBridge 6500N bridge

| For this MetroCluster type... | The connection type is... |
|-------------------------------|---------------------------|
| Fabric-attached               | Switches                  |
| Stretch                       | Direct FC connections     |

### Hardware requirements if you are using SAS optical cables

SAS optical cables can be used to cable SAS disk shelves in a stretch MetroCluster system to achieve greater connectivity over distance than is possible with SAS copper cables. SAS optical cables can be used in a new system installation or in an existing system in which you want to replace the SAS copper cables and FibreBridge 6500N bridges. SAS optical cables are supported for use with V-Series systems starting in Data ONTAP 8.2.1.

See *Configuring a stretch MetroCluster system with SAS disk shelves and SAS optical cables* on the NetApp Support Site at [support.netapp.com](http://support.netapp.com) for hardware requirements and procedures for installing a new system with SAS optical cables or replacing FibreBridge 6500N bridges with SAS optical cables.

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

You can use a UPS (Uninterruptible Power Supply) with your MetroCluster configurations. 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 configuration cannot shut down gracefully.

### Requirements for a shared-switches configuration

You must follow certain requirements when creating a shared-switches configuration.

- Fabric-attached MetroCluster configuration
- Brocade 5100 or 6510 switch
- SAS disk shelves with FibreBridge 6500N



To know how to install and configure FibreBridge 6500N as part of your MetroCluster configuration, see *Configuring a MetroCluster system with SAS disk shelves and FibreBridge 6500N bridges* on the NetApp Support Site at [support.netapp.com](http://support.netapp.com).


- A minimum of two ISLs per fabric with TI enabled.

For more information about configuring a switch in a shared-switches configuration, see the *Fabric-attached MetroCluster Systems Brocade Switch Configuration Guide* on the NetApp Support Site at [support.netapp.com](http://support.netapp.com).

## Fabric-attached MetroCluster configuration with disks

A fabric-attached MetroCluster configuration includes two pairs Brocade or Cisco Fibre Channel switch fabrics that provide long distance connectivity between the nodes. Through the 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 with disks:



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

## How fabric-attached MetroCluster configurations use Brocade and Cisco Fibre Channel switches

A MetroCluster configuration for distances greater than 500 meters connects the two nodes by using four Brocade or Cisco 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 HA pairs. For more information about the switches, see the *Fabric-attached MetroCluster Systems Brocade Switch Configuration Guide* for Brocade switches and *Fabric-attached MetroCluster Systems Cisco Switch Configuration Guide* for Cisco switches, available at the NetApp Support Site.

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 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.

Starting with Data ONTAP 8.2.1, the fabric-attached MetroCluster configuration supports ISL encryption on Brocade 65 10 switches. To learn more about enabling ISL encryption on your Brocade switch, see the *Fabric-attached MetroCluster Systems Brocade Switch Configuration Guide*.

**Related information**

*NetApp Support Site: [support.netapp.com](http://support.netapp.com)*

**Planning the fabric-attached MetroCluster installation**

You must fill out the fabric-attached MetroCluster configuration worksheet to record specific cabling information about your fabric-attached MetroCluster configuration. You must identify several pieces of information that you use during configuration procedures. Recording this information can reduce configuration errors.

**Step**

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 an 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                            |                |
| FC HBA 2        | A          | 1                            |                | 2                            |                |
|                 | B          | 3                            |                | 4                            |                |

Disk shelves or FibreBridge 6500N bridges (if you are using SAS disk shelves) at each site connect to the Fibre Channel switches. Use the following table to record which switch port the disk shelves or FibreBridge 6500N bridges are connected to.

**Note:** If you are using SAS or SATA disk shelves, the FibreBridges connect to the disk shelves. Each FibreBridge 6500N bridge needs to be connected to only one switch through one of the FC ports on the FibreBridge bridge. For cabling examples, see *Configuring a MetroCluster system with SAS disk shelves and FibreBridge 6500N bridges* at the NetApp Support Site.

| Disk shelf or FibreBridge 6500N... | 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                                  |            |                 |                         |                   |

### Required documentation, tools, and equipment

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

### Required documentation for MetroCluster configurations

You must refer to some of the flyers and guides that are required to install a new MetroCluster configuration, or convert two stand-alone systems into a MetroCluster 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 NetApp Support 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:

| Manual name                                                                   | Description                                                                                                   |
|-------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------|
| The appropriate system cabinet guide                                          | Describes how to install NetApp equipment into a system cabinet                                               |
| <i>Site Requirements Guide</i>                                                | Describes the physical requirements your site must meet to install NetApp equipment                           |
| The appropriate disk shelf guide                                              | Describes how to cable a disk shelf to a storage system                                                       |
| The appropriate hardware documentation for your storage system model          | Describes how to install the storage system, connect it to a network, and bring it up for the first time      |
| <i>Diagnostics Guide</i>                                                      | Describes the diagnostics tests that you can run on the storage system                                        |
| <i>Data ONTAP Upgrade and Revert/Downgrade Guide for 7-Mode</i>               | Describes how to upgrade the storage system and disk firmware, and how to upgrade the storage system software |
| <i>Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode</i> | Describes, among other topics, SyncMirror technology, which is used for mirrored HA pairs                     |
| <i>Data ONTAP System Administration Guide for 7-Mode</i>                      | Describes general storage system administration                                                               |
| <i>Data ONTAP Software Setup Guide for 7-Mode</i>                             | Describes how to configure the software of a new storage system for the first time                            |

| Manual name                                                                                   | Description                                                                                                           |
|-----------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------|
| <i>Fabric-attached MetroCluster Systems Brocade Switch Configuration Guide</i>                | Describes how to configure Brocade switches for a fabric-attached MetroCluster configuration                          |
| <i>Configuring a MetroCluster system with SAS disk shelves and FibreBridge 6500N bridges</i>  | Describes how to install FibreBridge 6500N bridges as part of your MetroCluster configuration                         |
| <i>Fabric-attached MetroCluster Systems Cisco Switch Configuration Guide</i>                  | Describes how to configure Cisco switches for a fabric-attached MetroCluster configuration                            |
| <i>FlexArray Virtualization Installation Requirements and Reference Guide</i>                 | Describes how to prepare array LUNs on storage arrays for use by Data ONTAP                                           |
| <i>Configuring a stretch MetroCluster system with SAS disk shelves and SAS optical cables</i> | Describes how to install SAS disk shelves using SAS optical cables as part of your stretch MetroCluster configuration |

**Related information**

*NetApp Support Site: [support.netapp.com](http://support.netapp.com)*

**Required tools**

There are some tools that you need to install the MetroCluster configuration.

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

**Required equipment for MetroCluster configurations**

You should receive the required equipment, such as storage system, HA interconnect adapter, and so on, to install a MetroCluster configuration.

See the *Hardware Universe* at [hwu.netapp.com](http://hwu.netapp.com) to confirm your storage system type, storage capacity, and so on.

**Note:** For fabric-attached MetroCluster configurations, use the information in the *Hardware Universe* labeled for MetroCluster configurations. For stretch MetroCluster configurations, use the information in the *Hardware Universe* labeled “for HA Environments.”

| Required equipment                                                | Stretch MetroCluster configuration                                                       | Fabric-attached MetroCluster configuration |
|-------------------------------------------------------------------|------------------------------------------------------------------------------------------|--------------------------------------------|
| Storage system                                                    | Two of the same type of storage systems                                                  |                                            |
| Storage                                                           | See the <i>Hardware Universe</i> at <a href="http://hwu.netapp.com">hwu.netapp.com</a> . |                                            |
| FibreBridge 6500N bridges (if you are attaching SAS disk shelves) | Two FibreBridges are required for each stack of SAS disk shelves.                        |                                            |

| Required equipment                                     | Stretch MetroCluster configuration                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Fabric-attached MetroCluster configuration                                                                                                                                                                                                                                                                                             |
|--------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| HA interconnect adapter                                | InfiniBand adapter (Required only for systems that do not use an NVRAM5 or NVRAM6 adapter, which functions as the HA interconnect adapter.)<br>FC-VI adapter (Required for the 31xx, 32xx, and 62xx dual-controller systems.)<br><b>Note:</b> When the FC-VI adapter is installed in a 31xx, 32xx, or 62xx system, the internal InfiniBand interconnect is automatically deactivated.                                                                                                                                                                                                      | FC-VI adapter                                                                                                                                                                                                                                                                                                                          |
| FC-AL or FC HBA (FC HBA for Disk) adapters             | Two or four Fibre Channel HBAs<br><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.                                                                                                                                                                                                                                                                                                                            |                                                                                                                                                                                                                                                                                                                                        |
| Fibre Channel switches                                 | Stretch MetroCluster configurations with array LUNs require these switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Two pairs of Brocade or Cisco switches<br><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                     | Not applicable                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Two or four SFPs are required per fabric, depending on whether you are using single or 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                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Not applicable                                                                                                                                                                                                                                                                                                                         |
| 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 or FibreBridge 6500N</li> <li>• Two SC/LC IB HA adapter cables</li> <li>• Four SC/LC or LC/LC cables</li> <li>• If you are using SAS optical cables with your SAS disk shelves, you would have received the cable qualities and lengths you ordered for your configuration.</li> </ul> <b>Note:</b> For information about required cables, see the Interoperability Matrix at <a href="http://support.netapp.com/matrix">support.netapp.com/matrix</a> . | <ul style="list-style-type: none"> <li>• LC/LC controller-to-switch cables</li> <li>• SC/LC (for DS14) or LC/LC (for DS14mk4 FC) disk shelf-to-switch cables or switch to FibreBridge 6500N</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> |

### Related information

[NetApp Support Site: support.netapp.com](http://support.netapp.com)

## Converting an HA pair to a fabric-attached MetroCluster configuration

With the correct hardware, you can reconfigure an HA pair with DS14mk4 FC disk shelves to a fabric-attached MetroCluster configuration.

### About this task

- If you are upgrading an existing HA pair 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 configuration. You can download the latest disk firmware from the NetApp Support Site.
- If you are upgrading a 31xx, 32xx, and 62xx dual controller systems, 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 configuration. You must also obtain and install the FC-VI interconnect card on both systems.
- For details about this conversion process, see TR-3548, *Best Practices for MetroCluster Design and Implementation*, at the NetApp Support Site.
- If you are converting an HA pair that has SAS disk shelves to a fabric-attached MetroCluster configuration, see *Configuring a MetroCluster system with SAS disk shelves and FibreBridge 6500N bridges* at the NetApp Support Site.

### Steps

1. Update Data ONTAP, storage system firmware, and disk firmware, as described in the *Data ONTAP Upgrade and Revert/Downgrade Guide for 7-Mode*, 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 *Hardware Universe* at the NetApp Support 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 *Fabric-attached MetroCluster Systems Brocade Switch Configuration Guide* and *Fabric-attached MetroCluster Systems Cisco Switch Configuration Guide*.

**Note:** The configuration and firmware requirements for Brocade and Cisco 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 switches for a MetroCluster configuration.

6. Cable the local node.
7. Enable the Data ONTAP licenses in the following order:
  - a. `cf.mode`  
Set the mode to **ha**
  - b. `cf.remote_syncmirror.enable`  
Set the option to **on**
8. Configure the local node depending on the type of HA pair:

| If you are converting a...         | Then...                                        |
|------------------------------------|------------------------------------------------|
| Standard HA pair                   | Set up mirroring and configure the local node. |
| Stretch MetroCluster configuration | 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 configuration.

#### Related concepts

[Configuring an HA pair](#) on page 50

Bringing up and configuring a standard or mirrored HA pair for the first time can require enabling HA mode capability and failover, setting options, configuring network connections, and testing the configuration.

[Recovering from a disaster by using MetroCluster configurations](#) on page 202

In situations such as prolonged power outages or natural disasters, you can use the 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.

#### Related tasks

[Cabling Node A](#) on page 145

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

[Cabling Node B](#) on page 149

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

[Disabling the change\\_fsid option in MetroCluster configurations](#) on page 56

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.

#### Related information

[NetApp Support Site: support.netapp.com](http://support.netapp.com)

## Cabling a fabric-attached MetroCluster configuration with disks

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

### Before you begin

- To cable a fabric-attached MetroCluster configuration, you must be familiar with HA configuration, the Brocade or Cisco command-line interface, and synchronous mirroring.
- You must also be familiar with the characteristics of fabric-attached MetroCluster configurations.
- You must also have the following information:
  - Correct Brocade or Cisco 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.

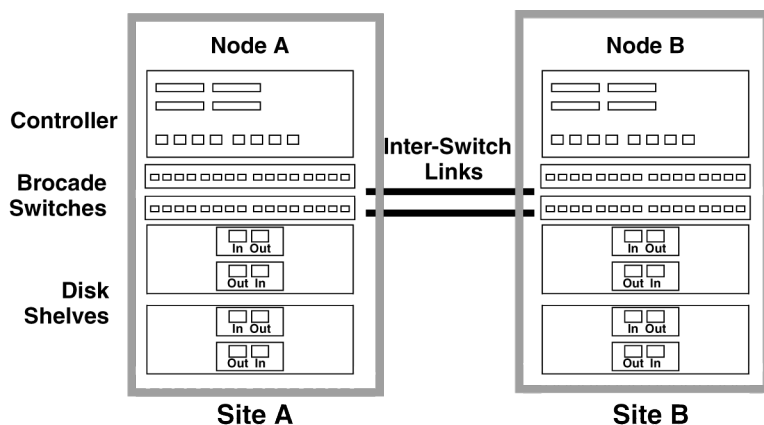
- Ethernet IP address for both the switches and nodes
  - Note:** The switches ship with a default IP address (10.77.77.77), which you can use if the switches are not attached to a network.
- Ethernet subnetmask
- Gateway address

### About this task

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

**Note:** If you are using SAS disk shelves, the disk shelves connect to FibreBridge 6500N bridges. To see an example of a cabled fabric-attached MetroCluster system with FibreBridges and SAS disk shelves, see the *Configuring a MetroCluster system with SAS disk shelves and FibreBridge 6500N bridges* at the NetApp Support Site at [support.netapp.com](http://support.netapp.com).

The following illustration shows a fabric-attached MetroCluster configuration:



### Related concepts

[Setup requirements and restrictions for fabric-attached MetroCluster configurations with disks](#) on page 133

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

[Configuring an HA pair](#) on page 50

Bringing up and configuring a standard or mirrored HA pair for the first time can require enabling HA mode capability and failover, setting options, configuring network connections, and testing the configuration.

[Recovering from a disaster by using MetroCluster configurations](#) on page 202

In situations such as prolonged power outages or natural disasters, you can use the 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.

### Related tasks

[Disabling the change\\_fsid option in MetroCluster configurations](#) on page 56

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.

### Steps

1. [Configuring the switches](#) on page 145
2. [Cabling Node A](#) on page 145
3. [Cabling Node B](#) on page 149



4. [Assigning disk pools](#) on page 152
5. [Verifying aggregate status](#) on page 153
6. [Setting up a shared-switches configuration](#) on page 153
7. [Fabric-attached MetroCluster configuration on single-enclosure HA pairs](#) on page 155
8. [Configuration differences for fabric-attached MetroCluster configurations on single-enclosure HA pairs](#) on page 156

## Configuring the switches

You need to configure the switches in a fabric-attached MetroCluster configuration.

### Step

To configure your Brocade or Cisco switches, see the *Fabric-attached MetroCluster Systems Brocade Switch Configuration Guide* or *Fabric-attached MetroCluster Systems Cisco Switch Configuration Guide* for your switch model at the NetApp Support 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 see the Interoperability Matrix at [support.netapp.com/matrix](http://support.netapp.com/matrix), when installing and configuring your MetroCluster switches. For information about ISL encryption on a Brocade switch, see the *Fabric-attached MetroCluster Systems Brocade Switch Configuration Guide* at the NetApp Support Site.

### After you finish

Proceed to configure Node A.

#### Related information

[NetApp Support Site: support.netapp.com](http://support.netapp.com)

## 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 HA interconnect to the switches, and ensure that the disk shelves in the configuration belong to the correct pools.

### About this task

If you are using SAS disk shelves, the SAS disk shelves connect to the FibreBridge 6500N bridges and the bridges connect to the switches.

### Steps


1. [Cabling the controller](#) on page 145
2. [Cabling the shelves](#) on page 146
3. [Cabling the FC-VI adapter and inter-switch link](#) on page 148

## Cabling the controller

You must cable the Fibre Channel ports on the controller to the Brocade or Cisco switches.

### About this task

The following illustration shows the cabling of the controller to the Brocade switches:



### 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 0a on the controller.
2. 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.
3. 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.

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

### After you finish

Proceed to cable disk shelves to the switches.

#### Related concepts

[Determining which Fibre Channel ports to use for Fibre Channel disk shelf connections](#) on page 36

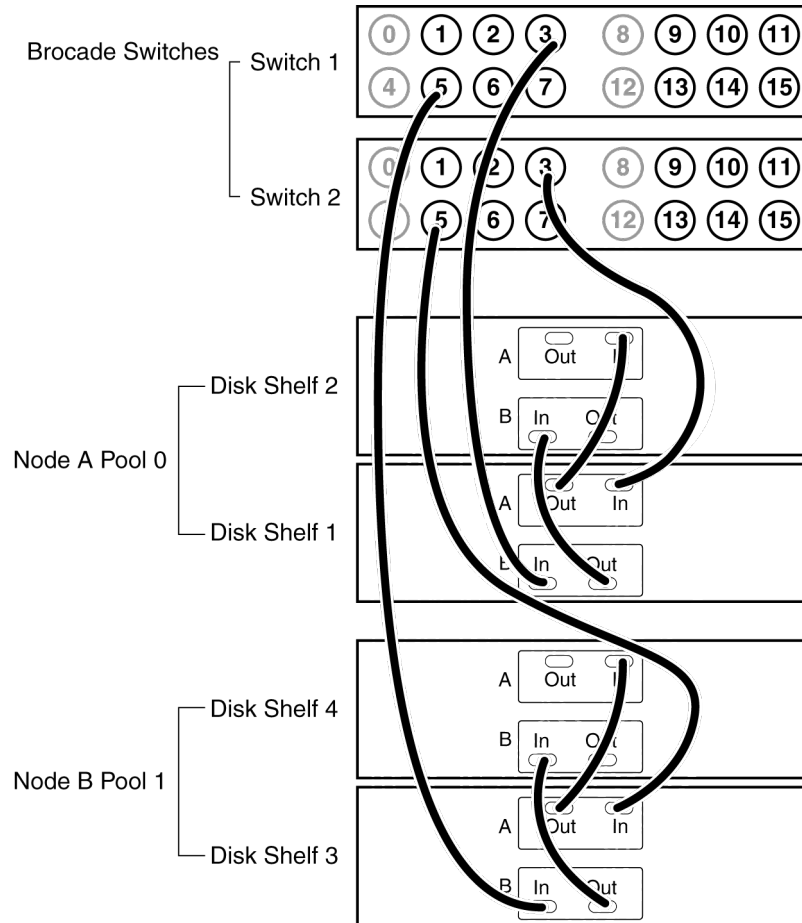
Before cabling your HA pair, you need to identify which Fibre Channel ports to use to connect your disk shelves to each storage system, and in what order to connect them.

### Cabling the shelves

You must cable the DS14mk4 FC disk shelf loops on Node A directly to the switches.

#### About this task

To cable SAS disk shelves and FibreBridge 6500N, see *Configuring a MetroCluster system with SAS disk shelves and FibreBridge 6500N bridges* on the NetApp Support Site.



**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.  
 In the example, switch port 3 is used.
  - 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, and then set the terminate switches on the last disk shelf to On.

**Note:** 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.

**After you finish**

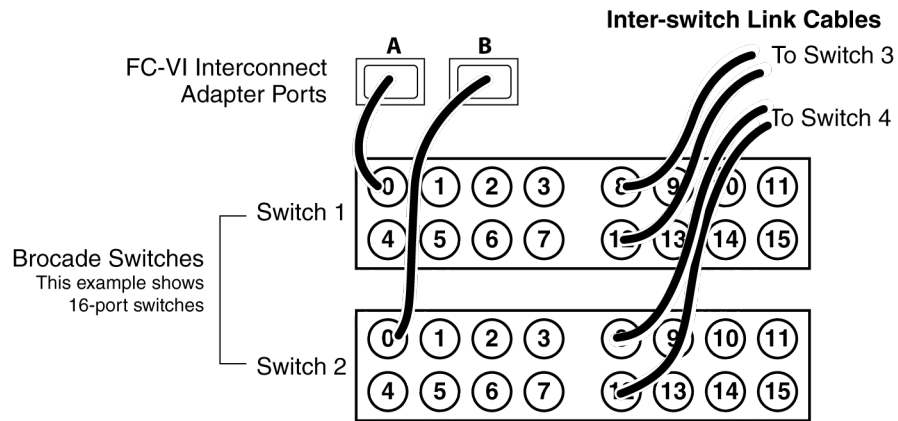
Proceed to cable the FC-VI adapter and interswitch connections.

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

You must cable the HA interconnect and Inter-Switch Link on Node A.

**About this task**

The following illustration shows the cabling of the FC-VI HA interconnect and Inter-Switch Links to the Brocade switches:



**Steps**

- 1. Connect one port of the FC-VI adapter to a port on switch 1 and the second port to the same port on switch 2.

In the example we are using 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. You must ensure that you have the FC-VI adapter in the correct slot for your system, as shown in the *Hardware Universe* at [hwu.netapp.com](http://hwu.netapp.com).

- 2. Connect an Inter-Switch Link cable to a port on each switch.

In the example, we are using ports 8 and 12 on switch 1 and switch 2 for the Inter-Switch Links.

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

**After you finish**

Cable Node B.

**Related information**

[NetApp Support Site: support.netapp.com](http://support.netapp.com)

## 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 HA interconnect to the switches, and ensure that the disk shelves in the configuration belong to the correct pools.

### About this task

If you are using SAS disk shelves, the SAS disk shelves connect to the FibreBridge 6500N bridge and the bridges connect to the switches.

### Steps


1. [Cabling the controller](#) on page 149
2. [Cabling the shelves](#) on page 150
3. [Cabling the FC-VI adapter and Inter-Switch Link](#) on page 151

## Cabling the controller

You must cable the Fibre Channel ports on the controller to the Brocade or Cisco switches.

### About this task

The following illustration shows the cabling of the controller to the Brocade switches:



### Steps

1. Determine the 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 0a 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.
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.

### After you finish

Proceed to cable disk shelves to the switches.

### Related concepts

[Determining which Fibre Channel ports to use for Fibre Channel disk shelf connections](#) on page


Before cabling your HA pair, 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.

### Cabling the shelves

You must cable the DS14mk2 or DS14mk4 disk shelf loops on Node B directly to the switches.

#### About this task

To cable SAS disk shelves and FibreBridge 6500N bridges, see *Configuring a MetroCluster system with SAS disk shelves and FibreBridge 6500N bridges* on the NetApp Support Site.



**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 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 7.
  - 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:** 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 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.

**After you finish**

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

**Related information**

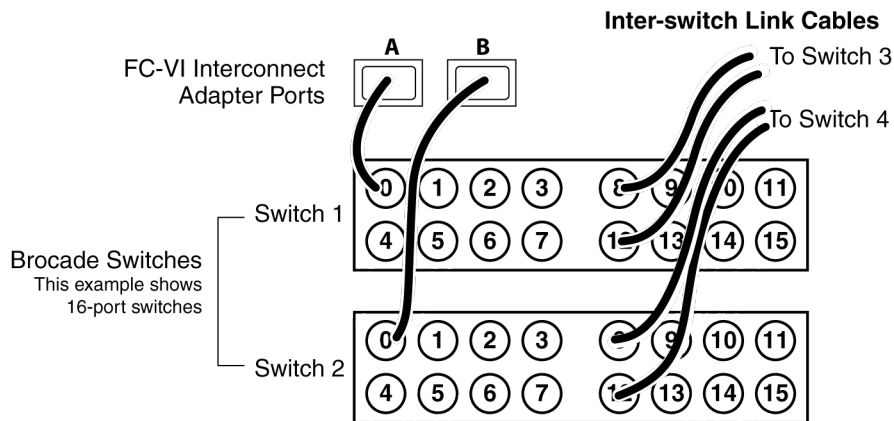
*NetApp Support Site: [support.netapp.com](http://support.netapp.com)*

**Cabling the FC-VI adapter and Inter-Switch Link**

You must cable the HA interconnect and Inter-Switch Link (ISL) on Node B.

**About this task**

The following illustration shows the cabling of the FC-VI HA interconnect and Inter-Switch Links to the Brocade switches:



**Steps**

1. Connect one port of the FC-VI adapter to another port that you have reserved for the FC-VI and ISL 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 *Hardware Universe* at the NetApp Support Site.

2. Connect an inter-switch link cable to a port on each switch.

The example uses port 8 and port 12 on switch 1 and switch 2 for the ISLs.

**Note:** If using dual ISLs, traffic isolation must be configured on the switches.

3. Verify that the status of the FC-VI links are UP and collect the node's FC Port Name (WWPN) for each FC-VI adapter port by entering the following command on both the nodes:

```
sysconfig -v adapter_slot
```

4. Verify that the FC-VI adapter connection is correct and the FC-VI adapter port is able to communicate to its peer FC-VI adapter port on the partner node by entering the following command on both the nodes:

```
fcvi ping adapter address
```

where *adapter* is the fcvi adapter port of the local node that is used to send the ping request.

*address* corresponds to FC Port Name (WWPN) of the peer FC-VI adapter port on the partner node.

### After you finish

Assign disks to disk pools.

#### Related information

[NetApp Support Site: support.netapp.com](http://support.netapp.com)

## Assigning disk pools

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.
2. Assign the local disks to Node A pool 0 by entering the following command at the console:

```
disk assign switch2:port3.* -p 0
```



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.* -p 1
```

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.

5. Assign the local disks to Node B pool 0 by entering the following command at the console:

```
disk assign switch4:port3.* -p 0
```

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:port5.* -p 1
```

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 aggregate status

You must use a specific procedure to verify your disk paths for configurations using DS14mk4 FC or SAS shelves.

### 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 7-Mode* for information about the `aggr status` command.

3. Repeat Steps 1 and 2 on Node B.

## Setting up a shared-switches configuration

In a shared-switches configuration, two fabric-attached MetroCluster configurations share the same four switches and the ISLs between them.

Shared-switches configuration are cost-effective because you can use the four Brocade switches between two MetroCluster configurations.

### Related concepts

[Requirements for a shared-switches MetroCluster configuration with array LUNs](#) on page 166

There are specific requirements for creating a shared-switches MetroCluster configuration with array LUNs.

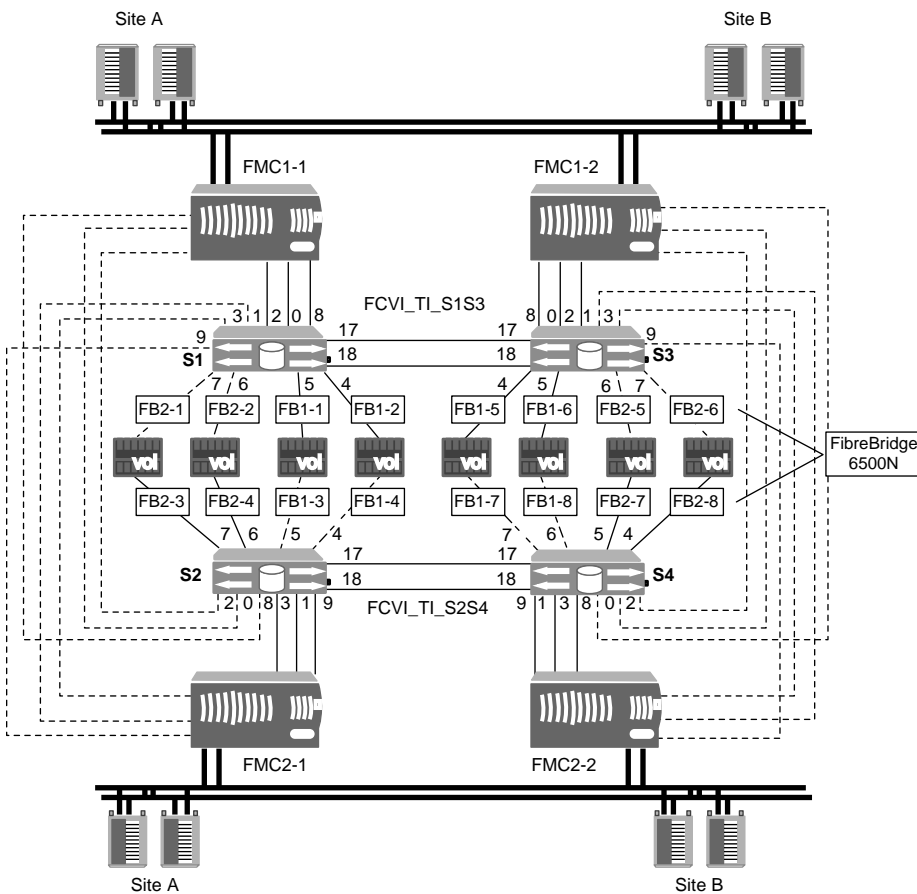
### Cabling the FC-VI adapter and ISL in a shared-switches configuration

Cabling in a shared-switches configuration involves cabling the FC-VI and HBA ports to the Brocade switches, and cabling the FibreBridge 6500N bridges to the Brocade switches and the SAS disk shelves.

#### About this task

In a shared-switches configuration, four nodes join together to form two fabric-attached MetroCluster configurations. In the illustration, FMC1-1 and FMC1-2 form one fabric-attached MetroCluster configuration, FMC1. FMC2-1 and FMC2-2 form another fabric-attached MetroCluster configuration, FMC2.

The following illustration shows the shared-switches configuration:



#### Steps

1. Connect the ports of the FC-VI adapter to primary and secondary switches for both the storage systems in the MetroCluster configuration.

In the illustration, the FC-VI adapter from storage system FMC1-1 connects to primary switch S1 and secondary switch S2 through port 0.

The FC-VI adapter from storage system FMC2-1 connects to primary switch S2 and secondary switch S1 through port 1.

2. Connect the ports of the HBA to the primary and secondary switches.

In the illustration, the HBA from storage system FMC1-1 connects to primary switch S1 and secondary switch S2 through ports 2 and 8.

The HBA from storage system FMC2-1 connects to primary switch S2 and secondary switch S1 through ports 3 and 9.

3. If you are using the FibreBridge 6500N bridge, complete the following steps:

For details about installing FibreBridge 6500N bridge as part of your MetroCluster configuration, see *Configuring a MetroCluster system with SAS disk shelves and FibreBridge 6500N bridges*, available at the NetApp Support Site.

- a. Connect the FibreBridge 6500N bridge to the switches.

In the illustration, the FibreBridge bridge FB1-1 of the FMC1-1 connects to switch S1 through port 5.

- b. Connect the FibreBridge 6500N bridge to the disk shelves.

4. Repeat steps 1 through 3 at another site.

5. When cabling at both the sites are complete, connect an ISL cable to a port on each switch.

In the illustration, the fabric containing switches S1 and S3 has two ISLs and the fabric containing switches S2 and S4 also has two ISLs.

The switch S1 connects to switch S3 and switch S2 connects to switch S4 through ports 17 and 18.

### After you finish

Configure the traffic isolation zone on the switches. In the illustration, FCVI\_TI\_S1S3 refers to the traffic isolation zone created using ports 0, 1, and 17 of switches S1 and S3. Similarly, FCVI\_TI\_S2S4 refers to the traffic isolation zone created using ports 0, 1, and 17 of switches S2 and S4. For more information, see the *Fabric-attached MetroCluster Systems Brocade Switch Configuration Guide*.

After cabling the shared-switches configuration, you must define the preferred primary port.

### Related information

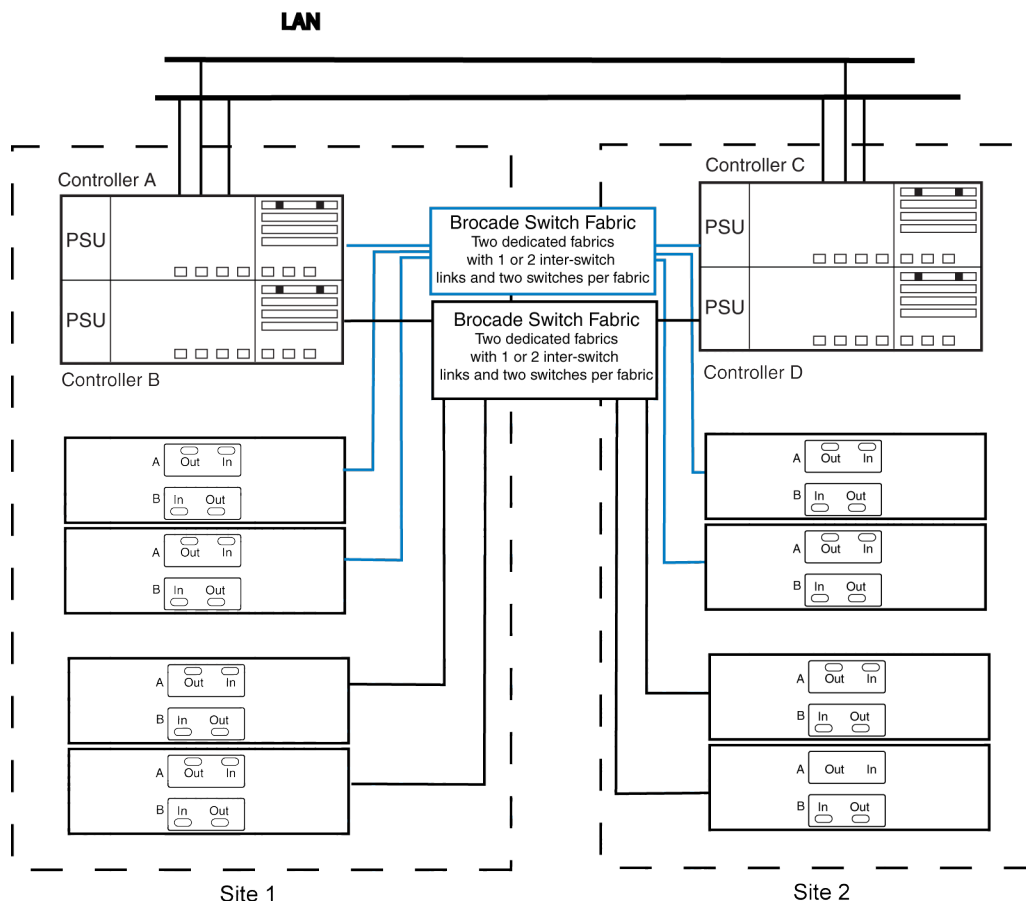
[NetApp Support Site: support.netapp.com](http://support.netapp.com)

## Fabric-attached MetroCluster configuration on single-enclosure HA pairs

You can configure a fabric-attached MetroCluster configuration between a pair of single-enclosure HA pair systems. In this configuration, the HA pair between the two controllers in each chassis is deactivated, and two separate, side-by-side MetroCluster configurations are formed between the four controllers.

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

The following figure shows a fabric-attached MetroCluster configuration on single-enclosure HA pair systems:



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

### Configuration differences for fabric-attached MetroCluster configurations on single-enclosure HA pairs

When configuring a fabric-attached MetroCluster configuration between single-enclosure HA pair (systems with two controllers in the same chassis), you get two separate MetroCluster configurations.

A single-enclosure HA pair can be connected to another HA pair to create two separate fabric-attached MetroCluster configurations. The internal InfiniBand connection in each system is automatically deactivated when the FC-VI card is installed in the controller.

You must cable each fabric-attached MetroCluster configuration separately by using the normal procedures for each and assign the storage appropriately.

#### Related concepts

[Fabric-attached MetroCluster configuration on single-enclosure HA pairs](#) on page 155

You can configure a fabric-attached MetroCluster configuration between a pair of single-enclosure HA pair systems. In this configuration, the HA pair between the two controllers in each chassis is

deactivated, and two separate, side-by-side MetroCluster configurations are formed between the four controllers.

## Setting a preferred primary port in a MetroCluster configuration

After cabling a MetroCluster configuration, you can define the primary port between the four switches to carry the traffic. You can set the preferred primary port in a shared-switches configuration by using the `ic primary set` command.

### About this task

By default, the primary port is 0.

### Steps

1. To set the primary port, enter the following command:

```
ic primary set 0|1 [-r]
```

You must use the `-r` option for the primary port to take effect immediately.

#### Note:

- If the port that you have set as the preferred primary port is down, the other port will be used. Once the port is up, it will be set as the preferred primary port.
- If the FC-VI link between the preferred primary ports fails, the fabric-attached MetroCluster configuration fails over to the secondary FC-VI link. When the primary link is restored, the fabric-attached MetroCluster configuration again uses the primary link. The failover temporarily causes the logs to be unsynchronized.

If you want to set the preferred primary port to 1, enter the following command:

```
ic primary set 1 -r
```

In a shared-switches configuration, if you have set the preferred primary port of FMC1 (FMC1-1 and FMC1-2) to 1, the preferred primary port for FMC2 (FMC2-1 and FMC2-2) is set to 0.

2. To view the current primary port, enter the following command:

```
ic primary show
```

The primary port is displayed.

## Removing the preferred primary port in a fabric-attached MetroCluster configuration

You can remove the port that you assigned as primary port in a fabric-attached MetroCluster configuration.

### Step

To remove the primary port, enter the following command:

```
ic primary unset
```

## Enabling unidirectional port mirroring in a MetroCluster configuration

The unidirectional port mirroring allows you to use the FC-VI link effectively and therefore increases the performance throughput. By enabling the unidirectional port mirroring, one link in

an FC-VI interconnect will be used for outgoing data traffic and the other for incoming traffic between the two nodes in a MetroCluster configuration.

### Before you begin

- The preferred primary port is disabled on both the nodes.
- The FC-VI links in a MetroCluster configuration is up and working.

### About this task

- The unidirectional port mirroring must be enabled or disabled at both the nodes in a MetroCluster configuration.

### Steps

1. Verify the status of the FC-VI links by entering the following command:

```
sysconfig -v
```

2. For enabling or disabling unidirectional port mirroring, do one of the following:

| If you want to...                          | Then enter the following command...                                                                                                                                                               |
|--------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enable the unidirectional port mirroring.  | <code>ic unidirectional_mirroring on</code><br><b>Note:</b> This is a default behavior. Therefore, you must first check the status of unidirectional port mirroring before enabling this feature. |
| Disable the unidirectional port mirroring. | <code>ic unidirectional_mirroring off</code>                                                                                                                                                      |

**Note:** When reverting to a Data ONTAP release prior to Data ONTAP 8.2.1, the unidirectional port mirroring gets disabled automatically.

3. Verify the status of the unidirectional port mirroring by entering the following command:

```
ic unidirectional_mirroring status
```

If the unidirectional port mirroring is enabled, it displays the following output:

```
Unidirectional mirroring is enabled
For outgoing data primary link is 0 and current link in use is 0.
For incoming data primary link is 1 and current link in use is 1.
```

If the unidirectional port mirroring is disabled, it displays the following output:

```
Unidirectional mirroring is disabled
For outgoing data primary link is 0 and current link in use is 0.
For incoming data primary link is 0 and current link in use is 0.
```

**Note:** When the unidirectional port mirroring is disabled, the primary link and current link in use for both outgoing and incoming traffic will be the same.

## What unidirectional port mirroring is

The unidirectional port mirroring feature enables you to increase the performance throughput by using the ISLs effectively for sending and receiving the data for NVRAM mirroring. By enabling this feature, one link in an FC-VI interconnect is used for outgoing and the other for incoming data traffic in a MetroCluster configuration.

The FC-VI interconnect link to be used for outgoing and incoming data traffic is identified based on System Serial Number of the nodes in a MetroCluster configuration. The node with higher System Serial Number uses FC-VI link 0 for outgoing data traffic and FC-VI link 1 for incoming data traffic. The partner node will, therefore, use FC-VI link 0 for incoming data traffic and FC-VI link 1 for outgoing data traffic.

**Note:** If any of the FC-VI link between the MetroCluster nodes fails, the FC-VI link fails over to the other available FC-VI link. When the failed FC-VI link is restored, the MetroCluster configuration resumes using earlier link. The failover temporarily causes logs to be unsynchronized.

## Verifying the MetroCluster configuration

After cabling and configuring the MetroCluster configuration, you must perform certain steps to ensure that the configuration is working properly.

### Before you begin

- You must have access to the following sessions:
  - Telnet or SSH access to the storage controllers
  - Telnet or SSH access to the Fibre Channel switches
  - Telnet access to the FibreBridge 6500N

### Steps

1. [Gathering the information for verifying the MetroCluster configuration](#) on page 159
2. [Testing takeover and giveback](#) on page 160
3. [Testing MetroCluster configuration during a switch failure](#) on page 160
4. [Testing a MetroCluster configuration during FibreBridge 6500N bridge failure](#) on page 161
5. [Testing a MetroCluster configuration during cluster failover on disaster](#) on page 161

## Gathering the information for verifying the MetroCluster configuration

To verify that the MetroCluster configuration is configured correctly, you must first collect information related to various components that are used in the configuration, such as switches, FibreBridge 6500N and so on. This information gathered is used for is then used for verifying a takeover and giveback operations.

### Steps

1. Collect and make a note of the IP address, user IDs, and passwords of the storage controllers at each node:

|            | Node A | Node B |
|------------|--------|--------|
| IP address |        |        |
| User ID    |        |        |
| Password   |        |        |

2. Collect and make a note of the IP addresses, User IDs and password of the four switches used at both the nodes:

|            | Node A_Switch 1 | Node A_Switch 2 | Node B_Switch 3 | Node B_Switch 4 |
|------------|-----------------|-----------------|-----------------|-----------------|
| IP address |                 |                 |                 |                 |
| User ID    |                 |                 |                 |                 |
| Password   |                 |                 |                 |                 |
| ISL1       |                 |                 |                 |                 |
| ISL2       |                 |                 |                 |                 |

3. Collect and make a note of the IP address, UserID and password of the FibreBridge 6500N at both the nodes:

|            | <b>Node A_Bridge 1</b> | <b>Node B_Bridge2</b> |
|------------|------------------------|-----------------------|
| IP address |                        |                       |
| User ID    |                        |                       |
| Password   |                        |                       |

**After you finish**

Perform the takeover and giveback operation.

**Testing takeover and giveback**

After you configure all aspects of your HA pair, you need to verify that it is operating as expected in maintaining uninterrupted access to both nodes' storage during takeover and giveback operations. Throughout the takeover process, the local (or takeover) node should continue serving the data normally provided by the partner node. During giveback, control and delivery of the partner's storage should return transparently to the partner node.

**Steps**

1. Check the cabling on the HA 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**

See the man page for command details.

The local node takes over the partner node and gives the following output:

```
Failover monitor: takeover completed
```

4. Use the `sysconfig -r` command to ensure that the local (takeover) node can access its partner's disks.
5. Give back the partner node's data service after it displays the `Waiting for giveback` message 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 saw the message that giveback was completed successfully:

| <b>If takeover and giveback...</b> | <b>Then...</b>                                                           |
|------------------------------------|--------------------------------------------------------------------------|
| Is completed successfully          | Repeat Step 2 through Step 5 on the partner node.                        |
| Fails                              | Correct the takeover or giveback failure and then repeat this procedure. |

**Testing MetroCluster configuration during a switch failure**

To ensure that an FC switch is configured correctly and the failure of the switch does not cause a node failover, you must manually test the FC switch used in a MetroCluster configuration.

**Before you begin**

You must have collected all the required information for the components used in a MetroCluster configuration.



### About this task

A switch failure must not initiate a disaster recovery and therefore during a switch failure, you must be able to access the data in both the nodes in a MetroCluster configuration.

### Steps

1. Fail one FC switch by powering it off.
2. Verify that the data on both nodes is still accessible.
3. Verify that takeover capability is enabled by entering the following command:

```
cf status
```

**Note:** After powering off the switch, the takeover capability might be disabled briefly. It is subsequently reenabled automatically.

4. Power on the switch.
5. Verify that the system returns to healthy state and all the disks are available.

## Testing a MetroCluster configuration during FibreBridge 6500N bridge failure

To ensure that the MetroCluster configuration and the FibreBridge 6500N bridge are configured correctly, you must manually fail the bridge. Failure of the FibreBridge 6500N bridge must not initiate a disaster recovery and you must be able to access the data from both the nodes in the MetroCluster configuration.

### Before you begin

You must have collected all the required information related to the components used in the MetroCluster configuration.

### Steps

1. Fail one FibreBridge 6500N bridge by powering it off.
2. Verify that the data on both the nodes is accessible.
3. Power on the bridge.
4. Verify that the system returns to an healthy state and data on both the nodes is accessible.

## Testing a MetroCluster configuration during cluster failover on disaster

You should ensure that the MetroCluster installation is correct by testing whether you are able to perform a disaster recovery operations.

### About this task

Before you perform this task, you must keep in mind that this task is disruptive.

### Steps

1. Halt the MetroCluster configuration by powering off the switches at one site.
2. Perform the disaster recovery operations by using the `cf forcetakeover` command.

*[Recovering from a disaster](#)*

## Implementing a stretch or fabric-attached MetroCluster configuration with array LUNs

---

Careful planning, a systematic configuration process, and testing is key to a successful implementation of a MetroCluster configuration. There are requirements and configuration and testing procedures that are unique to setting up a MetroCluster configuration to use array LUNs.

If your MetroCluster configuration includes both array LUNs and native disks, complete the setup and testing of one type of storage before starting configuration of the other type of storage.

**Note:** MetroCluster configurations are supported only with Data ONTAP operating in 7-Mode.

### Related concepts

[Types of systems that can use array LUNs in a MetroCluster configuration](#) on page 163  
Support for which systems can use array LUNs in a MetroCluster configuration differs according to Data ONTAP release.

[Planning for a MetroCluster configuration with array LUNs](#) on page 162  
Creating a detailed plan for your MetroCluster configuration helps you understand the unique requirements for a MetroCluster configuration that uses LUNs on storage arrays. A plan also helps you communicate with other people involved in the installation. Installing a MetroCluster configuration involves connecting and configuring a number of devices, which might be done by different people.

[Connecting devices in a MetroCluster configuration with array LUNs](#) on page 175  
You should plan the cabling required for your MetroCluster configuration before you start connecting the devices. It is helpful to have a port-to-port connectivity diagram to use for reference while you are connecting the devices in a MetroCluster configuration with array LUNs.

[Setting up Data ONTAP after connecting devices in a MetroCluster configuration with array LUNs](#) on page 191  
After connecting the devices in the MetroCluster configuration, you need to set up the Data ONTAP systems to use the storage on the storage array. You also need to set up any desired Data ONTAP features.

[Testing a MetroCluster configuration with array LUNs](#) on page 191  
It is important to test that your MetroCluster configuration is properly set up before putting it into a production environment.

## Planning for a MetroCluster configuration with array LUNs

Creating a detailed plan for your MetroCluster configuration helps you understand the unique requirements for a MetroCluster configuration that uses LUNs on storage arrays. A plan also helps you communicate with other people involved in the installation. Installing a MetroCluster configuration involves connecting and configuring a number of devices, which might be done by different people.

### Related concepts

[Types of systems that can use array LUNs in a MetroCluster configuration](#) on page 163  
Support for which systems can use array LUNs in a MetroCluster configuration differs according to Data ONTAP release.

[Requirements for a MetroCluster configuration with array LUNs](#) on page 164  
There are some unique requirements for setting up a MetroCluster configuration with array LUNs.

[Recommended fabric-attached MetroCluster configuration with array LUNs](#) on page 166  
When configuring a fabric-attached MetroCluster configuration with Data ONTAP systems and storage arrays, you should follow the best practice recommendations for how to connect the

Implementing a stretch or fabric-attached MetroCluster configuration with array LUNs

hardware components. The Interoperability Matrix at [support.netapp.com](http://support.netapp.com) contains information about the hardware components that are supported for a MetroCluster configuration using array LUNs without a PVR.

*Recommended stretch MetroCluster configuration with array LUNs* on page 169

When configuring a stretch MetroCluster configuration with Data ONTAP systems and storage arrays, you should follow the best practice recommendations for how to connect the hardware components. The Interoperability Matrix at [support.netapp.com](http://support.netapp.com) contains information about the hardware components that are supported for a MetroCluster configuration using array LUNs.

*Cabling guidelines for a MetroCluster configuration with array LUNs* on page 172

There are a number of cabling guidelines that you need to review before connecting the devices in a MetroCluster configuration with array LUNs.

*Planning zoning for a MetroCluster configuration with array LUNs* on page 173

Switch zoning defines paths between connected nodes based on the node's unique WWN. Sketching out the zoning plan enables you to correct errors before zones are configured, and makes it easier to communicate zoning information to the person who configures the switches.

### Types of systems that can use array LUNs in a MetroCluster configuration

Support for which systems can use array LUNs in a MetroCluster configuration differs according to Data ONTAP release.

| Data ONTAP release | Systems that can use array LUNs in a MetroCluster configuration                                                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Prior to 8.2       | V-Series systems only, and those systems can use only array LUNs                                                                                                                                                                       |
| 8.2                | V-Series systems only, and those systems can use array LUNs, disks, or both                                                                                                                                                            |
| 8.2.1 and later    | <ul style="list-style-type: none"> <li>• V-Series systems, which can use array LUNs, disks, or both</li> <li>• FAS systems that can be licensed to use array LUNs</li> </ul> These systems can use disks also (or instead) if desired. |

**Note:** V-Series ("V") systems and new FAS platforms released in Data ONTAP 8.2.1 and later can use array LUNs if the proper license is installed. In the documentation, these systems are collectively referred to as Data ONTAP systems.

### Implementation overview for a MetroCluster configuration with array LUNs

Implementing a MetroCluster configuration to use LUNs from storage arrays requires planning the implementation, installing hardware, connecting multiple devices, configuring Data ONTAP, and testing the MetroCluster configuration to ensure that it is operating correctly.

The following tasks must be completed, in the order shown, to set up a MetroCluster configuration to work with storage arrays. Storage array configuration is performed by the storage array administrator or the storage array vendor. Zoning is often performed by a switch administrator.

1. Planning your MetroCluster implementation.
2. Setting up the storage array to present array LUNs to Data ONTAP and configuring the parameters that Data ONTAP requires for a storage array to work with Data ONTAP.
3. Installing the FC-VI adapter on each Data ONTAP system, if it is not installed already.

**Note:** New Data ONTAP systems that are ordered for a MetroCluster configuration are shipped with the correct FC-VI adapter. If you are configuring a MetroCluster configuration with existing Data ONTAP systems, you might need to install a new adapter or move an existing adapter to another slot. The *Hardware Universe* at [hwu.netapp.com](http://hwu.netapp.com) contains information about the required FC-VI adapter for your platform and the slot to install it in.

4. Connecting the local Data ONTAP system to the fabric.
5. Connecting the remote Data ONTAP system to the fabric.
6. Connecting the switched fabric.
7. Connecting the storage array to the fabric.
8. Configuring zoning.
9. Assigning array LUNs to specific Data ONTAP systems.
10. Configuring Data ONTAP features.
11. Testing the MetroCluster configuration.

### Related concepts

[Connecting devices in a MetroCluster configuration with array LUNs](#) on page 175

You should plan the cabling required for your MetroCluster configuration before you start connecting the devices. It is helpful to have a port-to-port connectivity diagram to use for reference while you are connecting the devices in a MetroCluster configuration with array LUNs.

[Setting up Data ONTAP after connecting devices in a MetroCluster configuration with array LUNs](#) on page 191

After connecting the devices in the MetroCluster configuration, you need to set up the Data ONTAP systems to use the storage on the storage array. You also need to set up any desired Data ONTAP features.

[Testing a MetroCluster configuration with array LUNs](#) on page 191

It is important to test that your MetroCluster configuration is properly set up before putting it into a production environment.

### Related tasks

[Connecting the local systems in a MetroCluster configuration with array LUNs](#) on page 176

The first stage in connecting the devices in a MetroCluster configuration that uses arrays LUNs is to connect a local Data ONTAP system to the fabric.

[Connecting the remote systems in a MetroCluster configuration with array LUNs](#) on page 179

After connecting the local Data ONTAP systems to the fabric, you need to connect the remote systems to the fabric.

[Connecting the switch fabric in a MetroCluster configuration with array LUNs](#) on page 182

Connecting the switch fabric involves connecting the ISL cables and applying power to the switches.

[Connecting the fabric and storage array in a MetroCluster configuration with array LUNs](#) on page 183

All storage arrays, regardless of model, must be configured to allow Data ONTAP to access a specific LUN on two (primary and secondary) storage array ports.

[Configuring zoning in a MetroCluster configuration with array LUNs](#) on page 185

Switch zoning defines paths between connected nodes based on the node's unique WWN. Single-initiator zoning is recommended for a configuration with array LUNs.

## Requirements for a MetroCluster configuration with array LUNs

There are some unique requirements for setting up a MetroCluster configuration with array LUNs.

The Interoperability Matrix at [support.netapp.com](http://support.netapp.com) contains the latest policies for MetroCluster configurations that use array LUNs. It includes information about supported storage arrays, switches, and Data ONTAP systems (V-Series systems and FAS systems) that are supported for use with array LUNs. The *Interoperability Matrix* is the final authority for information about requirements and restrictions for MetroCluster configurations that use array LUNs.

### Requirements for Data ONTAP systems using array LUNs in a MetroCluster configuration

- The platform must be identified in the *Interoperability Matrix* as supported for MetroCluster configurations.
  - Prior to Data ONTAP 8.2.1, only V-Series systems (not FAS systems) can be deployed in a MetroCluster configuration that uses array LUNs.
  - Starting in Data ONTAP 8.2.1, V-Series systems and FAS systems that can be licensed to use array LUNs can be deployed in a MetroCluster.
- The Data ONTAP systems in the MetroCluster configuration must be all V-Series systems or all FAS systems, and both nodes must be the same model.
- A system with two controllers in the same enclosure, such as a 32xx system, requires an FC-VI adapter.

The *Hardware Universe* contains information about adapter requirements for different models.

In releases prior to Data ONTAP 8.2, MetroCluster configurations with V-Series systems could not use native disks. Starting in Data ONTAP 8.2, a MetroCluster configuration with V-Series systems can include native disks, storage arrays, or both.

**Note:** Connecting a single FC initiator port to multiple storage array target ports is not supported in a MetroCluster configuration. This functionality is supported for non MetroCluster configurations starting in Data ONTAP 8.2, as described in the *FlexArray Virtualization Installation Requirements and Reference Guide*.

### Requirements for storage arrays

- The storage arrays must be identified in the *Interoperability Matrix* as supported for MetroCluster configurations.
- The storage arrays in the MetroCluster configuration must be symmetric, which means the following:
  - The two storage arrays must be in the same vendor family.  
The *FlexArray Virtualization Implementation Guide for Third-Party Storage* contains details about storage array families.
  - A pool of array LUNs for the mirrored storage must be created on each of the two storage arrays.  
The array LUNs must be the same size.
  - Disk types (SATA, SSD, or SAS) used for mirrored storage must be the same on both storage arrays.
  - Storage arrays that provide tiered configurations (for example, Hitachi) must use the same tiered topology on each side of the MetroCluster configuration.
  - The root volume must be mirrored for successful takeover to occur.

### Requirements for FC switches

- The switches and switch firmware must be identified in the *Interoperability Matrix* as supported for MetroCluster configurations.
- Each fabric must have two switches.
- Each Data ONTAP system must be connected to storage using redundant components so that there is redundancy in case of device failure.

The *Interoperability Matrix* contains the latest information about switch requirements and supported switches. The configuration and firmware requirements for switches in a MetroCluster environment might differ from those in other configurations.

### Zoning requirements

- Single-initiator to single-target zoning is recommended.

## Implementing a stretch or fabric-attached MetroCluster configuration with array LUNs

Single-initiator to single-target zoning limits each zone to a single FC initiator port. It also improves discovery and boot time because the FC initiators on the Data ONTAP systems do not attempt to discover each other.

- FC-VI ports on the FC adapters must be zoned end-to-end across the fabric.  
The *Interoperability Matrix* contains specific guidelines for FC-VI.

**Requirements for ISLs**

Data ONTAP supports using one or two ISLs, depending on the configuration. The *Interoperability Matrix* contains information about the number of ISLs for specific configurations.

For specific information about Traffic Isolation zoning for FC-VI traffic and for information about ISL encryption, see the "Brocade's Zoning Feature" section in the *Fabric-attached MetroCluster Systems Brocade Switch Configuration Guide* on the NetApp Support Site at [support.netapp.com](http://support.netapp.com). Regular zoning guidelines apply for Data ONTAP systems and storage arrays.

**SyncMirror requirements**

- SyncMirror is required for a MetroCluster configuration.  
Licenses are no longer required for SyncMirror. You need to enable SyncMirror with the `cf.mode` and `cf.remote_syncmirror.enable` options.
- Two separate storage arrays are required for the mirrored storage.  
**Note:** Using two different storage arrays for mirroring is optional in configurations that are not MetroCluster configurations.
- Two sets of LUNs are required—one set for the aggregate on the local storage array (pool0) and another set at the remote storage array for the mirror of the aggregate (the other plex of the aggregate, pool1).

The *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode* provides more details about setting up SyncMirror for a MetroCluster configuration that uses array LUNs.

**Requirements for a shared-switches MetroCluster configuration with array LUNs**

There are specific requirements for creating a shared-switches MetroCluster configuration with array LUNs.

You need the following to create a shared-switches MetroCluster configuration with array LUNs:

- In releases prior to Data ONTAP 8.2, four Brocade 5100 switches. In Data ONTAP 8.2 and later, four Brocade 5100 switches or 6510 switches.
- Two storage arrays at each site.  
One storage array at one site is used for the first MetroCluster configuration and the other storage array at the site is used for the second MetroCluster configuration.  
**Note:** A single storage array at each site is supported if the target ports on the storage array and the array LUNs are not shared between the two MetroCluster configurations.
- A minimum of two ISLs per fabric with TI enabled.

For more information about configuring Brocade switches, see the *Fabric-attached MetroCluster Systems Brocade Switch Configuration Guide*.

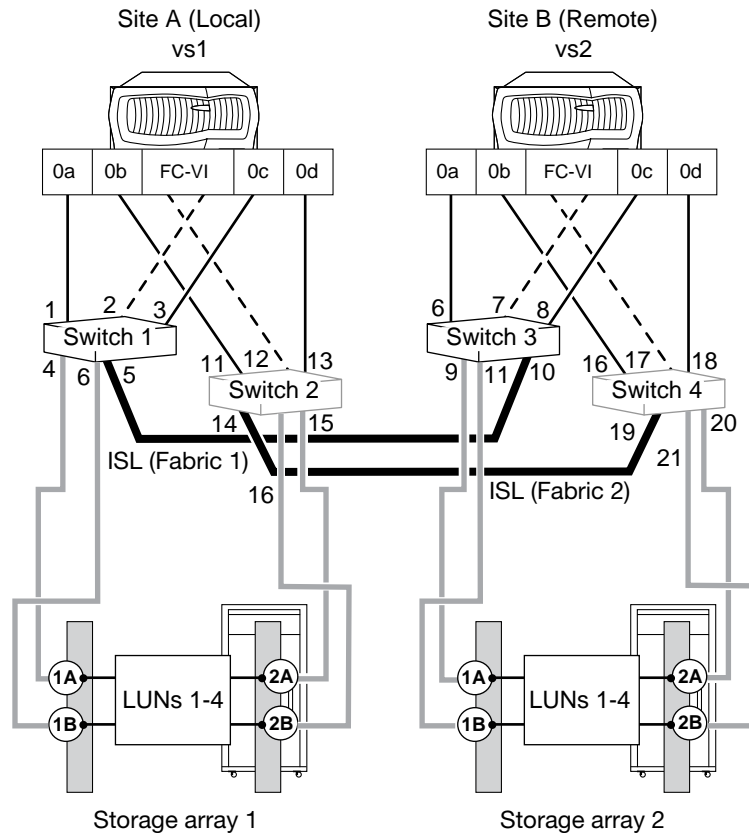
**Recommended fabric-attached MetroCluster configuration with array LUNs**

When configuring a fabric-attached MetroCluster configuration with Data ONTAP systems and storage arrays, you should follow the best practice recommendations for how to connect the hardware components. The Interoperability Matrix at [support.netapp.com](http://support.netapp.com) contains information about the hardware components that are supported for a MetroCluster configuration using array LUNs without a PVR.

Implementing a stretch or fabric-attached MetroCluster configuration with array LUNs

The following illustration shows the components and best practice configuration of a fabric-attached MetroCluster configuration that uses array LUNs from storage arrays. The fabric-attached MetroCluster configuration shown provides the same single-point-of-failure protections that are available for all mirrored HA pairs.

**Note:** FC initiator port names on Data ONTAP systems differ among some platforms. Your platform's port names might not match the port names shown in the illustration.



The following sections describe the connections for the fabric-attached MetroCluster configuration in the previous illustration.

**Interconnections on Data ONTAP systems**

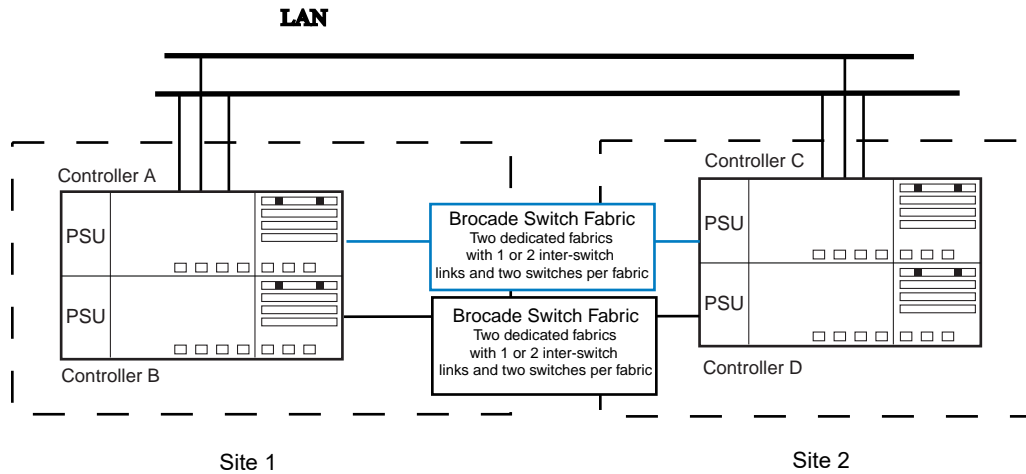
The Data ONTAP systems in the sample MetroCluster illustration are configured to be an HA pair. They are interconnected by connecting ports A and B on each system's FC-VI adapter to alternate switch fabrics.

| Data ONTAP system | Switch            | Fabric |
|-------------------|-------------------|--------|
| vs1: Port A       | Switch 2, Port 12 | 2      |
| vs1: Port B       | Switch 1, port 2  | 1      |
| vs2: Port A       | Switch 4, port 17 | 2      |
| vs2: Port B       | Switch 3, port 7  | 1      |

If your Data ONTAP systems are dual-enclosure systems with two controllers in the same chassis (such as a 32xx system), you connect the top controller in one system to the top controller in the other system through the fabric. You connect the bottom controllers between the fabric in the same way. These two separate connections result in two separate fabric-attached MetroCluster configurations.

**Note:** The internal InfiniBand connection in each system is automatically deactivated when the FC-VI adapter is installed in the controller.

The following illustration shows the connections of the controllers to the switch fabric for dual-enclosure systems with two controllers in the same chassis. Controller A and Controller C form one MetroCluster configuration. Controller B and Controller D form the second MetroCluster configuration:



### Inter-Switch Link connections (ISLs)

Fabric-attached MetroCluster configurations use a switched fabric to connect the local half of the configuration to the remote half of the configuration. The sample MetroCluster illustration shows the following:

- Switches 1 and 3 are connected to each other (Fabric 1).  
 The first fabric in the MetroCluster configuration begins from Switch 1 on Site A (local) and is completed by connecting the ISL cable to the first switch on Site B (remote)—Switch 3.
- Switches 2 and 4 are also connected (Fabric 2).  
 The second fabric is created using Switch 2 on Site A (local), connected through a second ISL cable, to the second switch on Site B (remote)—Switch 4.

The following table lists the ISLs in this configuration:

| ISL connection      | Switch              | Fabric   |
|---------------------|---------------------|----------|
| Port 5 on switch 1  | Port 10 on switch 3 | Fabric 1 |
| Port 14 on switch 2 | Port 19 on switch 4 | Fabric 2 |

### Data ONTAP system-to-switch connections

The best practice connections in the recommended MetroCluster illustration eliminate a single point of failure in the following ways:

- FC initiator ports on the same FC controller chip (for example, port 0a and 0b) connect to alternate fabrics.
- Multiple paths and zones ensure that FC initiator ports on the same controller chips access the array LUN from different Data ONTAP systems and switch fabrics.

The following table lists the connections from the Data ONTAP system to the switch:

| Data ONTAP system and FC initiator port | Switch           |
|-----------------------------------------|------------------|
| vs1: FC port 0a                         | Switch 1: Port 1 |



| Data ONTAP system and FC initiator port | Switch            |
|-----------------------------------------|-------------------|
| vs1: FC port 0b                         | Switch 2: Port 11 |
| vs1: FC port 0c                         | Switch 1: Port 3  |
| vs1: FC port 0d                         | Switch 2: Port 13 |
| vs2: FC port 0a                         | Switch 3: Port 6  |
| vs2: FC port 0b                         | Switch 4: Port 16 |
| vs2: FC port 0c                         | Switch 3: Port 8  |
| vs2: FC port 0d                         | Switch 4: Port 18 |

**Storage array-to-switch connections**

Best practice connections from the storage array to the switch are as follows:

- Ports 1A and 2A on each storage array connect to alternate fabrics.
- Ports 1B and 2B on each storage array connect to alternate fabrics.
- Data ONTAP systems are configured to access any array LUN on two storage array paths (1A and 2A or 1B and 2B).

| Storage array port | Switch            | Fabric |
|--------------------|-------------------|--------|
| Array 1: Port 1A   | Switch 1: Port 4  | 1      |
| Array 1: Port 2A   | Switch 2: Port 15 | 2      |
| Array 1: Port 1B   | Switch 1: Port 6  | 1      |
| Array 1: Port 2B   | Switch 2: Port 16 | 2      |
| Array 2: Port 1A   | Switch 3: Port 9  | 1      |
| Array 2: Port 2A   | Switch 4: Port 20 | 2      |
| Array 2:Port 1B    | Switch 3: Port 11 | 1      |
| Array 2:Port 2B    | Switch 4: Port 21 | 2      |

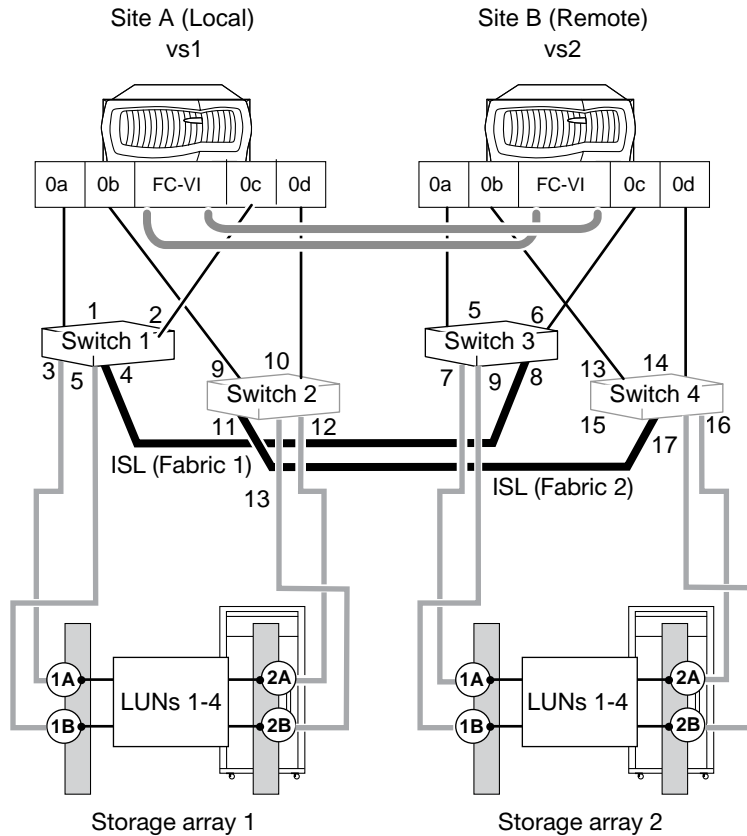
**Recommended stretch MetroCluster configuration with array LUNs**

When configuring a stretch MetroCluster configuration with Data ONTAP systems and storage arrays, you should follow the best practice recommendations for how to connect the hardware components. The Interoperability Matrix at [support.netapp.com](http://support.netapp.com) contains information about the hardware components that are supported for a MetroCluster configuration using array LUNs.

The following illustration shows the components and best practice for a stretch MetroCluster configuration with storage arrays.

**Note:** FC initiator port names on Data ONTAP systems differ between some platforms. Your platform's port names might not match the port names shown in the illustration. The *Hardware Universe* contains information that helps you determine in which slot to install the FC-VI adapter for your platform.

Implementing a stretch or fabric-attached MetroCluster configuration with array LUNs



**Interconnections on Data ONTAP systems**


In a stretch MetroCluster configuration, the Data ONTAP systems are connected directly.

The sample illustration shows the connections between a dual-enclosure HA pair for systems that can contain only one controller in the enclosure.

For Data ONTAP systems that are dual-enclosure systems with two controllers in the same chassis, such as 31xx systems, you install a Data ONTAP system at each site. To implement the stretch MetroCluster configuration, you must install an FC-VI adapter in each controller to provide the HA interconnect between the systems.

**Note:** When the FC-VI adapter is installed in the system, the internal InfiniBand interconnect is automatically disabled.

The following illustration shows the FC-VI port connections between HA pair systems such as the 31xx. Each enclosure contains two controllers. You configure two separate MetroCluster configurations among the four controllers.



In the stretch MetroCluster illustration, one MetroCluster configuration is between Controller A and Controller C.

| From Controller A | To Controller C |
|-------------------|-----------------|
| FC-VI port A      | FC-VI port A    |
| FC-VI port B      | FC-VI port B    |

In the stretch MetroCluster illustration, the other MetroCluster configuration is between Controller B and Controller D.

| From Controller B | To Controller D |
|-------------------|-----------------|
| FC-VI port A      | FC-VI port A    |
| FC-VI port B      | FC-VI port B    |

Some older V-Series models use NVRAM to connect the systems in a MetroCluster configuration. The following table lists the NVRAM connections on V-Series systems that connect to each other through NVRAM.

| vs1                | vs2                |
|--------------------|--------------------|
| NVRAM port L02 Ph2 | NVRAM port L02 Ph2 |
| NVRAM port L01 Ph1 | NVRAM port L01 Ph1 |

### Inter-Switch Link connections (ISLs)

Stretch MetroCluster configurations use a switched fabric to connect the local half of the configuration to the remote half of the configuration. In the stretch MetroCluster illustration, Switches 1 and 3 are connected to each other (Fabric 1). Switches 2 and 4 are also connected (Fabric 2).

- The first fabric in the MetroCluster configuration begins from Switch 1 on Site A (local) and is completed by connecting the ISL cable to the first switch on Site B (remote)—Switch 3.
- The second fabric is created using Switch 2 on Site A (local), connected through a second ISL cable to the second switch on Site B (remote)—Switch 4.

| ISL connection      | Switch              | Fabric   |
|---------------------|---------------------|----------|
| Port 4 on switch 1  | Port 8 on switch 3  | Fabric 1 |
| Port 11 on switch 2 | Port 15 on switch 4 | Fabric 2 |

### Data ONTAP system-to-switch connections

The best practice connections in the stretch MetroCluster with storage arrays illustration eliminate a single-point-of-failure in the following ways:

- FC initiator ports on the same FC controller chip (for example, port 0a and 0b) connect to alternate fabrics.
- Multiple paths and zones ensure that FC initiator ports on the same controller chips access the array LUN from different Data ONTAP systems and switch fabrics.

The following table lists the connections from the Data ONTAP system to the switch:

| Data ONTAP system and FC initiator port | Switch           |
|-----------------------------------------|------------------|
| vs1: FC port 0a                         | Switch 1: Port 1 |
| vs1: FC port 0b                         | Switch 2: Port 9 |
| vs1: FC port 0c                         | Switch 1: Port 2 |

| Data ONTAP system and FC initiator port | Switch            |
|-----------------------------------------|-------------------|
| vs1: FC port 0d                         | Switch 2: Port 10 |
| vs2: FC port 0a                         | Switch 3: Port 5  |
| vs2: FC port 0b                         | Switch 4: Port 13 |
| vs2: FC port 0c                         | Switch 3: Port 6  |
| vs2: FC port 0d                         | Switch 4: Port 14 |

**Storage array-to-switch connections**

The best practice connections from the storage array to the switch are as follows:

- Ports 1A and 2A on each storage array connect to alternate fabrics.
- Ports 1B and 2B on each storage array connect to alternate fabrics.
- Data ONTAP systems are configured to access any array LUN on two storage array paths (1A and 2A or 1B and 2B).

| Storage array port | Switch            | Fabric |
|--------------------|-------------------|--------|
| Array 1: Port 1A   | Switch 1: Port 3  | 1      |
| Array 1: Port 2A   | Switch 2: Port 12 | 2      |
| Array 1: Port 1B   | Switch 1: Port 5  | 1      |
| Array 1: Port 2B   | Switch 2: Port 13 | 2      |
| Array 2: Port 1A   | Switch 3: Port 7  | 1      |
| Array 2: Port 2A   | Switch 4: Port 16 | 2      |
| Array 2: Port 1B   | Switch 3: Port 9  | 1      |
| Array 2: Port 2B   | Switch 4: Port 17 | 2      |

**Related concepts**

[Connecting devices in a MetroCluster configuration with array LUNs](#) on page 175

You should plan the cabling required for your MetroCluster configuration before you start connecting the devices. It is helpful to have a port-to-port connectivity diagram to use for reference while you are connecting the devices in a MetroCluster configuration with array LUNs.

**Cabling guidelines for a MetroCluster configuration with array LUNs**

There are a number of cabling guidelines that you need to review before connecting the devices in a MetroCluster configuration with array LUNs.

**Cabling guidelines for a dual-enclosure HA pair with multiple controllers in the same enclosure**

Some platforms, for example, 31xx systems, support two controllers in the same enclosure. For a MetroCluster configuration, you must have two such systems. You configure the two systems into a pair of MetroCluster configurations by connecting the FC-VI adapter between the two top controllers and then connecting the FC-VI adapter between the bottom controllers.

In such a configuration, the internal InfiniBand connections between the controllers in the same enclosure are automatically deactivated. Therefore, the two controllers in the enclosure are no longer in an HA pair with each other. Each controller is connected through FC-VI connections to another controller of the same type, so the four controllers form two independent MetroCluster configurations.

### Cabling guidelines for partner-to-partner connections

The cabling for the MetroCluster partners is different on each node of the HA pair. Use the following guidelines to plan your cabling for partner-to-partner connections.

- Each node in the HA pair must have a different name so that you can distinguish them.
- One node needs to be identified as being in Site A and the other as being in Site B.  
For example, the local node could be System 1, Site A, and the remote node could be System 2, Site B.
- Each port on the node must be connected to the same fabric.  
For example, if Port A of the FC-VI adapter on the local node is connected to Switch 2 and Port A of the FC-VI adapter on the remote node is connected to Switch 4, then Switch 2 and Switch 4 must be connected by the ISL, thereby connecting the nodes to the same fabric.

### Cabling guidelines for Data ONTAP system-to-switch connections

- FC ports on the same channel controller chip cannot be connected to the same switch.  
One port must be connected to one switch and the other port must be connected to the other switch. For example, if onboard port 0a is connected to Switch 3, you cannot connect onboard port 0b to Switch 3; port 0b must be connected to Switch 4.  
  
**Note:** Connecting both onboard ports of the pair to the same switch port number can simplify cabling and management of the MetroCluster configuration. For example, if port 0a is connected to Port 1, Switch 1, connect port 0b to Port 1, Switch 2.
- All switches within a fabric must be the same switch model and have the same number of licensed ports.

### Cabling guidelines for Inter-Switch connections

You can connect an Inter-Switch Link (ISL) to any available switch port.

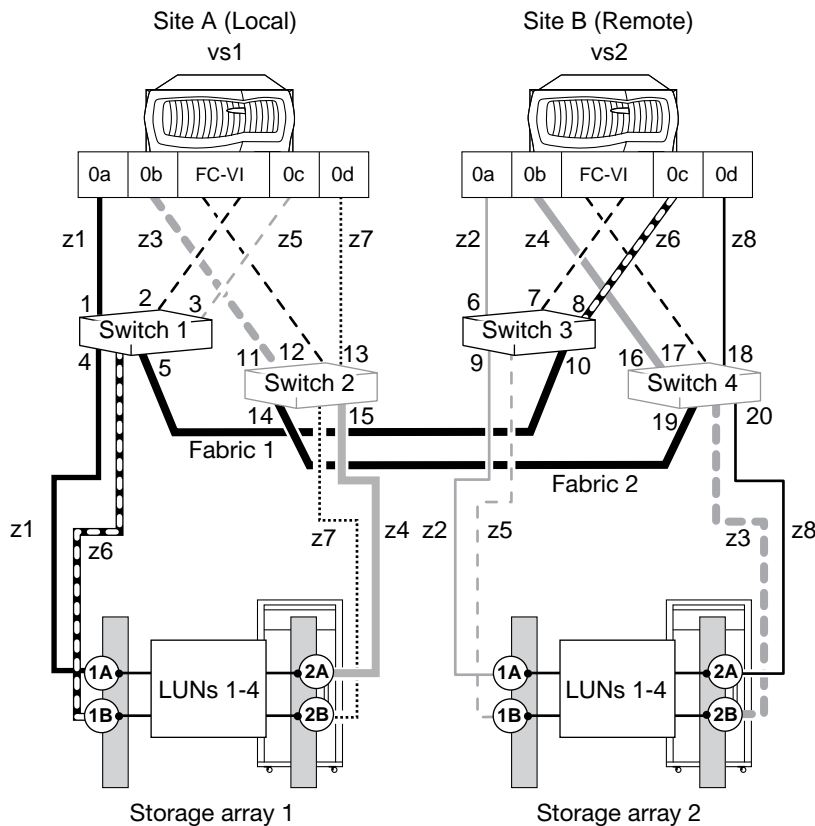
### Cabling guidelines for Data ONTAP system-to-storage array connections

Cabling between the Data ONTAP systems and the storage arrays must be redundant and conform to the required number of paths (two).

## Planning zoning for a MetroCluster configuration with array LUNs

Switch zoning defines paths between connected nodes based on the node's unique WWN. Sketching out the zoning plan enables you to correct errors before zones are configured, and makes it easier to communicate zoning information to the person who configures the switches.

You can use the following example as a reference when determining your zoning plan. The example shows single-initiator zoning for a fabric-attached MetroCluster configuration. The lines in the following example represent zones rather than connections; each line is labeled with its zone number:



In the sample illustration, four array LUNs are allocated on each storage array for the MetroCluster configuration. LUNs of equal size are provisioned on the storage arrays at both sites, which is a SyncMirror requirement. Each FC initiator port on each Data ONTAP system has a path to each of its LUNs on the storage arrays. The ports on the storage array are redundant, and are configured as follows:

- Storage array 1
  - Ports 1A and 2A are a redundant port pair.
  - Ports 1B and 2B are a redundant port pair.
  - In each port pair, both ports can access LUNs 1 through 4 because they are alternate paths.
- Storage array 2
  - Ports 1A and 2A are a redundant port pair.
  - Ports 1B and 2B are a redundant port pair.
  - In each port pair, both ports can access LUNs 1 through 4 because they are alternate paths.

Switches are zoned so that there are only two paths to each array LUN, one unique path from each FC initiator port on a Data ONTAP system through each switch. If there are multiple connections between a Data ONTAP system and the switch, the best practice recommendation is to put each connection into a separate zone.

The following table shows the zones for this example:

| Zone            | Data ONTAP system and FC initiator port | Storage array port      |
|-----------------|-----------------------------------------|-------------------------|
| <b>Switch 1</b> |                                         |                         |
| z1              | vs1:Port 0a                             | Storage array 1:Port 1A |

| Zone            | Data ONTAP system and FC initiator port | Storage array port      |
|-----------------|-----------------------------------------|-------------------------|
| z5              | vs1:Port 0c                             | Storage array 2:Port 1B |
| <b>Switch 2</b> |                                         |                         |
| z3              | vs1:Port 0b                             | Storage array 2:Port 2B |
| z7              | vs1:Port 0d                             | Storage array 1:Port 2A |
| <b>Switch 3</b> |                                         |                         |
| z2              | vs2:Port 0a                             | Storage array 2:Port 1A |
| z6              | vs2:Port 0c                             | Storage array 1:Port 1B |
| <b>Switch 4</b> |                                         |                         |
| z8              | vs2:Port 0d                             | Storage array 2:Port 2A |
| z4              | vs2:Port 0b                             | Storage array 1:Port 2B |

**Related tasks**

*Configuring zoning in a MetroCluster configuration with array LUNs* on page 185

Switch zoning defines paths between connected nodes based on the node's unique WWN. Single-initiator zoning is recommended for a configuration with array LUNs.

*Testing zoning of FC-VI ports in a MetroCluster configuration with array LUNs* on page 191

To test zoning of the FC-VI ports, you need to check that you have not crossed the FC-VI "a" ports and "b" ports in your zones.

## Connecting devices in a MetroCluster configuration with array LUNs

You should plan the cabling required for your MetroCluster configuration before you start connecting the devices. It is helpful to have a port-to-port connectivity diagram to use for reference while you are connecting the devices in a MetroCluster configuration with array LUNs.

**Note:** The instructions for connecting devices in a MetroCluster configuration use the term FC-VI adapter. Alternative terms for FC-VI adapter are VI-MC adapter and VI-MetroCluster adapter.

**Related concepts**

*Requirements for a MetroCluster configuration with array LUNs* on page 164

There are some unique requirements for setting up a MetroCluster configuration with array LUNs.

*Recommended fabric-attached MetroCluster configuration with array LUNs* on page 166

When configuring a fabric-attached MetroCluster configuration with Data ONTAP systems and storage arrays, you should follow the best practice recommendations for how to connect the hardware components. The Interoperability Matrix at [support.netapp.com](http://support.netapp.com) contains information about the hardware components that are supported for a MetroCluster configuration using array LUNs without a PVR.

*Recommended stretch MetroCluster configuration with array LUNs* on page 169

When configuring a stretch MetroCluster configuration with Data ONTAP systems and storage arrays, you should follow the best practice recommendations for how to connect the hardware components. The Interoperability Matrix at [support.netapp.com](http://support.netapp.com) contains information about the hardware components that are supported for a MetroCluster configuration using array LUNs.

*Cabling guidelines for a MetroCluster configuration with array LUNs* on page 172

There are a number of cabling guidelines that you need to review before connecting the devices in a MetroCluster configuration with array LUNs.

*Setting up Data ONTAP after connecting devices in a MetroCluster configuration with array LUNs* on page 191

After connecting the devices in the MetroCluster configuration, you need to set up the Data ONTAP systems to use the storage on the storage array. You also need to set up any desired Data ONTAP features.

*Testing a MetroCluster configuration with array LUNs* on page 191

It is important to test that your MetroCluster configuration is properly set up before putting it into a production environment.

### Steps

1. *Connecting the local systems in a MetroCluster configuration with array LUNs* on page 176
2. *Connecting the remote systems in a MetroCluster configuration with array LUNs* on page 179
3. *Connecting the switch fabric in a MetroCluster configuration with array LUNs* on page 182
4. *Connecting the fabric and storage array in a MetroCluster configuration with array LUNs* on page 183
5. *Configuring zoning in a MetroCluster configuration with array LUNs* on page 185
6. *Setting up a shared-switches configuration* on page 186
7. *Setting a preferred primary port in a MetroCluster configuration* on page 189
8. *Removing the preferred primary port in a fabric-attached MetroCluster configuration* on page 189
9. *Enabling unidirectional port mirroring in a MetroCluster configuration* on page 189
10. *Changing the configuration speed of a stretch MetroCluster configuration* on page 190

## Connecting the local systems in a MetroCluster configuration with array LUNs

The first stage in connecting the devices in a MetroCluster configuration that uses arrays LUNs is to connect a local Data ONTAP system to the fabric.

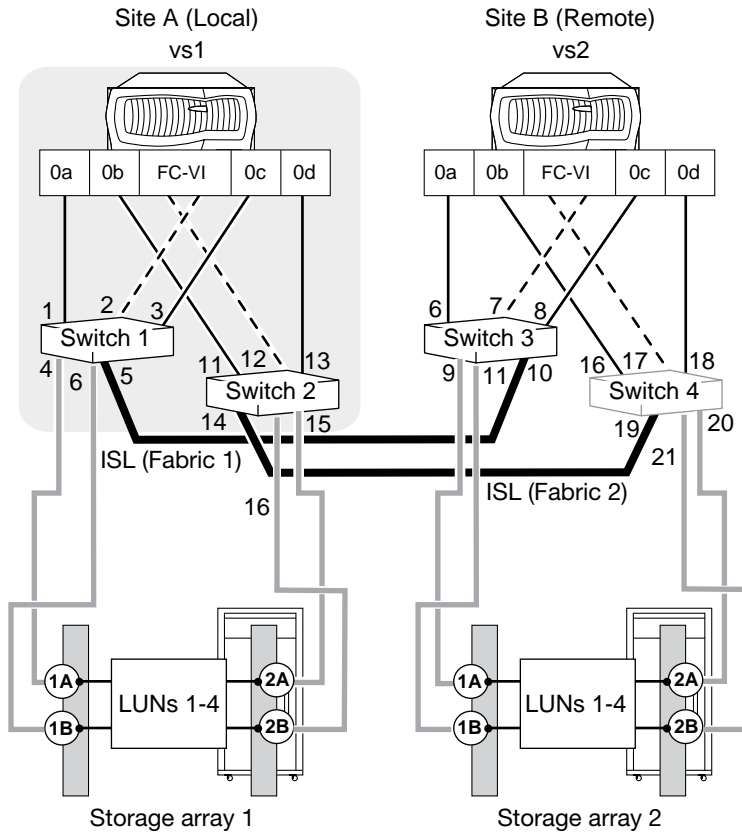
### Before you begin

- MetroCluster configuration planning must be complete.
- Storage array LUNs must already be created and presented to Data ONTAP.

### About this task

Use the following fabric-attached MetroCluster illustration as a reference for connecting the local systems in a MetroCluster configuration that uses array LUNs. Differences for stretch MetroCluster configurations are noted in the procedure.





**Note:** A port pair consists of two FC initiator ports that are used to access the same set of array LUNs. For example, FC initiators 0a and 0d are accessing the same LUNs. See the *FlexArray Virtualization Installation Requirements and Reference Guide* for specific information about port pairs for your platform.

### Steps

1. On the local Data ONTAP system, locate the NVRAM or FC-VI interconnect module.
2. Take one of the following actions to connect the local Data ONTAP system to the remote Data ONTAP system.

| For a...                                    | Do this...                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Fabric-attached configuration               | <ol style="list-style-type: none"> <li>a. Cable FC-VI Port A to one fabric (Port A to Switch 2 in the previous illustration).</li> <li>b. Cable FC-VI Port B to the alternate fabric (Port B to Switch 1 in the previous illustration).</li> <li>c. For dual-enclosure systems, repeat substeps a and b for the bottom controllers.</li> </ol> |
| Stretch configuration with an FC-VI adapter | Connect the cables between the FC-VI ports directly (point-to-point).                                                                                                                                                                                                                                                                          |
| Stretch configuration with an NVRAM adapter | <ol style="list-style-type: none"> <li>a. Install a copper-to-fiber converter in L01 Ph1, and cable L01 Ph1 on the local node to L01 Ph1 on the remote node.</li> <li>b. Install a copper-to-fiber converter in L02 Ph2, and cable L02 Ph2 on the local node to L02 Ph2 on the remote node.</li> </ol>                                         |

- 
- | For a...                                             | Do this...                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Dual-enclosure HA pair systems with an FC-VI adapter | <ul style="list-style-type: none"><li>a. Connect port A of the FC-VI adapter on the top controller of the local site (vs1) to port A of the corresponding FC-VI adapter at the remote site (vs2).</li><li>b. Connect port B of the FC-VI adapter on the top controller of the local site (vs1) to port B of the corresponding FC-VI adapter at the remote site (vs2).</li><li>c. Repeat substeps a and b for connecting the FC-VI adapter on the bottom controller.</li></ul> |
- 
3. Connect the FC initiator ports on the Data ONTAP system to the switched fabric:
    - a. Identify one of the two FC initiator port pairs on the Data ONTAP system.
    - b. Cable one port of the FC initiator port pair to one fabric.
    - c. Cable the other port of the pair to the alternate fabric.  
In the sample illustration, these are ports 0a and 0b.
    - d. Identify the other onboard FC initiator port pair on the Data ONTAP system.
    - e. Cable one port of the pair to one fabric.
    - f. Cable another port of the pair to the alternate fabric.  
In the sample illustration, these are ports 0c and 0d.
  4. Connect the Data ONTAP system to a tape backup device through a separate FC initiator port or SCSI tape adapter.
  5. Connect a console cable to the console port on the Data ONTAP system, then connect the console cable to the adapter.  
Use the RJ-45 to DB-9 adapter that is included with your system.
  6. Install the cable management tray:
    - a. Pinch the arms of the tray and fit the holes in the arms through the motherboard tray pins.
    - b. Push the cables into the cable holders, thread the adapter cables through the top rows of the cable holders, and then thread the port cables through the lower cable holders.
  7. Connect the Data ONTAP system to the Ethernet network by plugging the network cable into the networking port.  
If you are connecting more than one network cable to the network, connect to the ports sequentially. Use the cable management tray to direct all the cabling from your system.
  8. Connect the remote management device from the back of the Data ONTAP system to the network using an Ethernet cable.  
The network switch port for the remote management device connection must negotiate down to 10/100 or autonegotiate.
  9. If applicable, turn on any tape backup devices.
  10. For each power supply on the Data ONTAP system, take the following steps:
    - a. Ensure that the power switch is in the Off (0) position.
    - b. Connect the socket end of the power cord to the power plug on the power supply.
    - c. Secure the power cord with the retaining adjustable clip on the power supply.
    - d. Plug the other end of the power cord into a grounded electrical outlet.  
**Note:** To obtain power supply redundancy, you must connect the second power supply to a separate AC circuit.
  11. Start a communications program.  
You must use some form of communications program to perform initial network setup and configuration of the Data ONTAP system. You can start a communications program through the remote management device or through the console after connecting to the serial port.

### After you finish

Connect the remote Data ONTAP system to the fabric.

#### Related concepts

[Recommended fabric-attached MetroCluster configuration with array LUNs](#) on page 166  
When configuring a fabric-attached MetroCluster configuration with Data ONTAP systems and storage arrays, you should follow the best practice recommendations for how to connect the hardware components. The Interoperability Matrix at [support.netapp.com](http://support.netapp.com) contains information about the hardware components that are supported for a MetroCluster configuration using array LUNs without a PVR.

[Recommended stretch MetroCluster configuration with array LUNs](#) on page 169  
When configuring a stretch MetroCluster configuration with Data ONTAP systems and storage arrays, you should follow the best practice recommendations for how to connect the hardware components. The Interoperability Matrix at [support.netapp.com](http://support.netapp.com) contains information about the hardware components that are supported for a MetroCluster configuration using array LUNs.

[Cabling guidelines for a MetroCluster configuration with array LUNs](#) on page 172  
There are a number of cabling guidelines that you need to review before connecting the devices in a MetroCluster configuration with array LUNs.

#### Related tasks

[Connecting the remote systems in a MetroCluster configuration with array LUNs](#) on page 179  
After connecting the local Data ONTAP systems to the fabric, you need to connect the remote systems to the fabric.

## Connecting the remote systems in a MetroCluster configuration with array LUNs

After connecting the local Data ONTAP systems to the fabric, you need to connect the remote systems to the fabric.

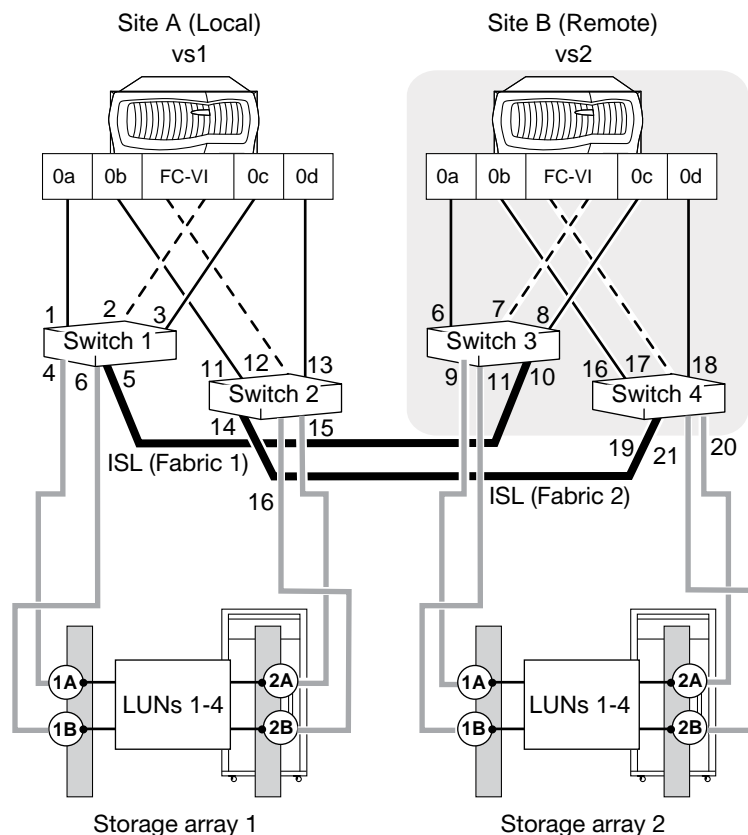
### Before you begin

The local Data ONTAP systems must be connected to the fabric.

### About this task

Use the following illustration of a fabric-attached MetroCluster configuration as a reference for connecting the remote Data ONTAP systems. Differences for stretch MetroCluster configurations are noted in the procedure.

**Note:** A port pair consists of two FC initiator ports that are used to access the same set of array LUNs. For example, FC initiators 0a and 0d are accessing the same LUNs. See the *FlexArray Virtualization Installation Requirements and Reference Guide* for specific information about port pairs.



**Steps**

1. Connect the local Data ONTAP system to the remote Data ONTAP system.

| For a...                                 | Do this...                                                                                                                                                                                                                                     |
|------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Fabric-attached configuration            | <ol style="list-style-type: none"> <li>a. Cable FC-VI Port A to one fabric (Port A to Switch 4 in the sample illustration).</li> <li>b. Cable FC-VI Port B to the alternate fabric (Port B to Switch 3 in the sample illustration).</li> </ol> |
| Stretch configuration with an FC-VI card | Use a point-to-point connection.                                                                                                                                                                                                               |

2. Connect the FC initiator ports to the switched fabric:
  - a. Identify one of the two FC initiator port pairs on the Data ONTAP system.
  - b. Cable one port of the FC initiator port pair to one fabric.
  - c. Cable the other port of the pair to the alternate fabric.  
 In the sample illustration, these are ports 0a and 0b.
  - d. Identify the other FC initiator port pair on the Data ONTAP system.
  - e. Cable one port of the pair to one fabric.
  - f. Cable another port of the pair to the alternate fabric.  
 In the sample illustration, these are ports 0c and 0d.
3. Connect the Data ONTAP system to a tape backup device through a separate FC initiator port or SCSI tape adapter.
4. Connect a console cable to the console port on the Data ONTAP system.

## Implementing a stretch or fabric-attached MetroCluster configuration with array LUNs

Use the RJ-45 to DB-9 adapter that is included with your system. Connect the console cable to the adapter.

5. Install the cable management tray.
  - a. Pinch the arms of the tray and fit the holes in the arms through the motherboard tray pins.
  - b. Push the cables into the cable holders, thread the adapter cables through the top rows of the cable holders, and then thread the port cables through the lower cable holders.
6. Connect the Data ONTAP system to the Ethernet network by plugging the network cable into the networking port.

If you are connecting more than one network cable to the network, connect to the ports sequentially. Use the cable management tray to direct all the cabling from your system.

7. Connect the remote management device from the back of the Data ONTAP system to the network using an Ethernet cable.

The network switch port for the remote management device connection must negotiate down to 10/100 or autonegotiate.

8. If applicable, turn on any tape backup devices.
9. For each power supply on the Data ONTAP system, take the following steps:
  - a. Ensure that the power switch is in the Off (0) position.
  - b. Connect the socket end of the power cord to the power plug on the power supply.
  - c. Secure the power cord with the retaining adjustable clip on the power supply.
  - d. Plug the other end of the power cord into a grounded electrical outlet.

**Note:** To obtain power supply redundancy, you must connect the second power supply to a separate AC circuit.

10. Start a communications program.

You must use some form of communications program to perform initial network setup and configuration of the Data ONTAP system. You can start a communications program through the remote management device or through the console after connecting to the serial port.

### After you finish

Connect the switch fabric in the MetroCluster configuration.

#### Related concepts

*Recommended fabric-attached MetroCluster configuration with array LUNs* on page 166  
When configuring a fabric-attached MetroCluster configuration with Data ONTAP systems and storage arrays, you should follow the best practice recommendations for how to connect the hardware components. The Interoperability Matrix at [support.netapp.com](http://support.netapp.com) contains information about the hardware components that are supported for a MetroCluster configuration using array LUNs without a PVR.

*Recommended stretch MetroCluster configuration with array LUNs* on page 169  
When configuring a stretch MetroCluster configuration with Data ONTAP systems and storage arrays, you should follow the best practice recommendations for how to connect the hardware components. The Interoperability Matrix at [support.netapp.com](http://support.netapp.com) contains information about the hardware components that are supported for a MetroCluster configuration using array LUNs.

*Cabling guidelines for a MetroCluster configuration with array LUNs* on page 172  
There are a number of cabling guidelines that you need to review before connecting the devices in a MetroCluster configuration with array LUNs.

#### Related tasks

*Connecting the switch fabric in a MetroCluster configuration with array LUNs* on page 182

Connecting the switch fabric involves connecting the ISL cables and applying power to the switches.

## Connecting the switch fabric in a MetroCluster configuration with array LUNs

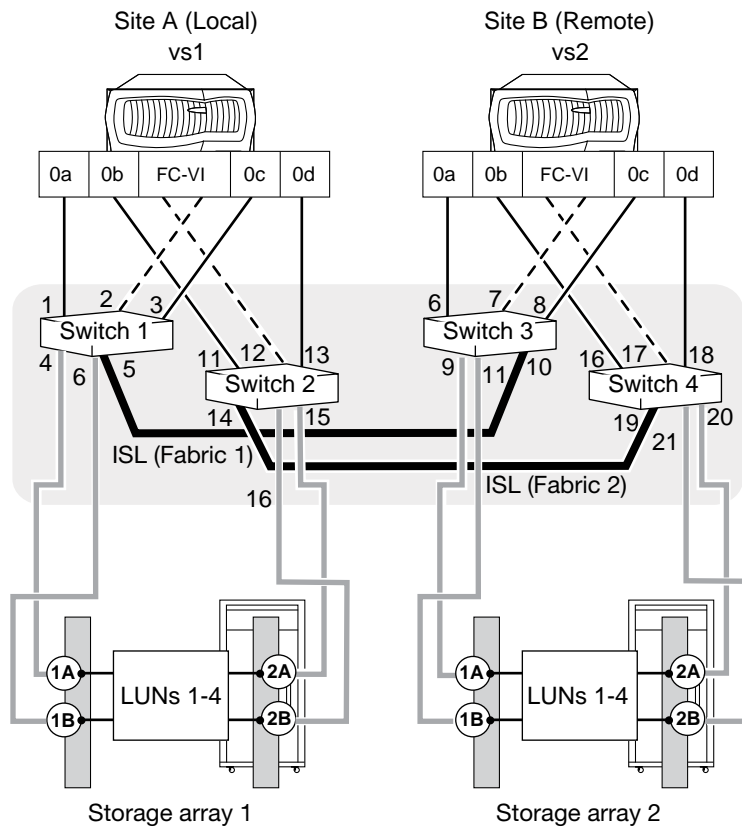
Connecting the switch fabric involves connecting the ISL cables and applying power to the switches.

### Before you begin

The local and remote Data ONTAP systems must be connected to the fabric.

### About this task

The following illustration provides a reference for connecting the switch fabric:



### Steps

1. Connect the switched fabric:
  - a. If you are configuring a new switch fabric, set the domain ID to a unique domain ID for each switch in the fabric.  
See your switch documentation for details.
  - b. Connect an ISL cable to a switch on one fabric and to another switch on the same fabric.  
In the sample illustration, Fabric 1, Switch 1, Port 5 connects to Fabric 1, Switch 3, Port 10.
  - c. Connect an ISL cable on a switch on the alternate fabric to another switch on the alternate fabric.  
In the sample illustration, Fabric 2, Switch 2, Port 14 connects to Fabric 2, Switch 4, Port 19.

**Note:** You must install a long distance SFP adapter in each port that you use to connect an ISL cable. You might also need to install an additional switch license to provide ISL support.

2. Make sure that all switch IDs are set, and then turn on each switch 10 minutes apart from one another.

### After you finish

Connect the fabric and the storage arrays.

#### Related concepts

*Recommended fabric-attached MetroCluster configuration with array LUNs* on page 166  
When configuring a fabric-attached MetroCluster configuration with Data ONTAP systems and storage arrays, you should follow the best practice recommendations for how to connect the hardware components. The Interoperability Matrix at [support.netapp.com](http://support.netapp.com) contains information about the hardware components that are supported for a MetroCluster configuration using array LUNs without a PVR.

*Recommended stretch MetroCluster configuration with array LUNs* on page 169  
When configuring a stretch MetroCluster configuration with Data ONTAP systems and storage arrays, you should follow the best practice recommendations for how to connect the hardware components. The Interoperability Matrix at [support.netapp.com](http://support.netapp.com) contains information about the hardware components that are supported for a MetroCluster configuration using array LUNs.

*Cabling guidelines for a MetroCluster configuration with array LUNs* on page 172  
There are a number of cabling guidelines that you need to review before connecting the devices in a MetroCluster configuration with array LUNs.

#### Related tasks

*Connecting the fabric and storage array in a MetroCluster configuration with array LUNs* on page 183  
All storage arrays, regardless of model, must be configured to allow Data ONTAP to access a specific LUN on two (primary and secondary) storage array ports.

## Connecting the fabric and storage array in a MetroCluster configuration with array LUNs

All storage arrays, regardless of model, must be configured to allow Data ONTAP to access a specific LUN on two (primary and secondary) storage array ports.

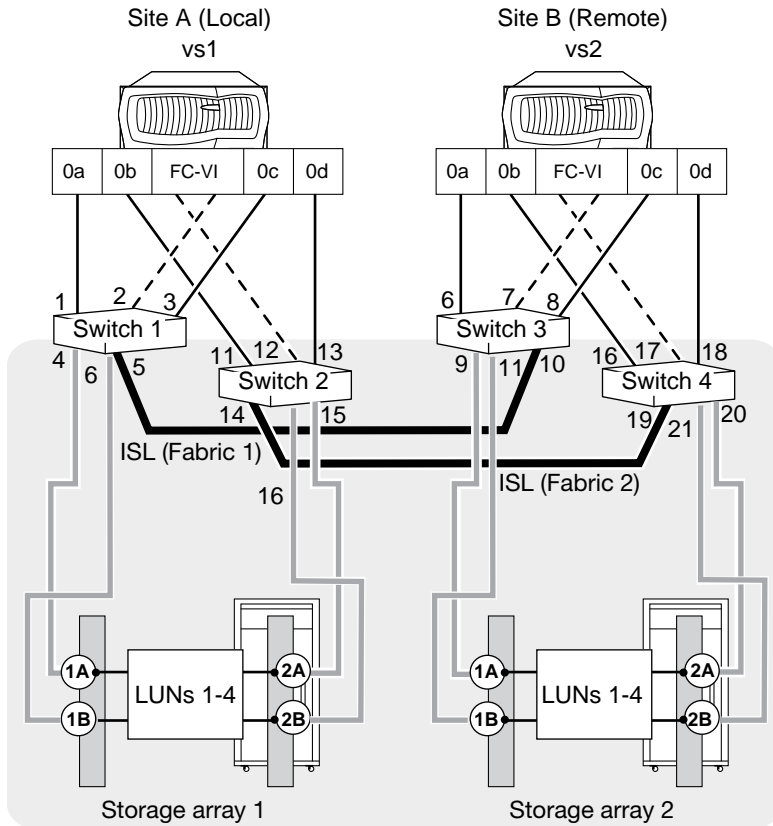
### Before you begin

- The local and remote Data ONTAP systems must be connected to the fabric.
- The fabric must be connected and the switches must be powered up.
- Your storage array documentation verified that the ports you plan to use do not access more than the number of array LUNs and host groups that are supported for that storage array model.

### About this task

The following illustration provides a reference for connecting the fabric and the storage arrays:

## Implementing a stretch or fabric-attached MetroCluster configuration with array LUNs

**Steps**

1. Connect the ports on the storage array at Site A:
  - a. Connect storage array controller 1A to any port on one fabric.  
In the sample illustration, this is Switch 1, Port 4, Fabric 1.
  - b. Connect controller 2A to any port on the alternate fabric.  
In the sample illustration this is Switch 2, Port 15, Fabric 2.
  - c. Connect storage array controller 1B to any port on one fabric.  
In the sample illustration, this is Switch 1, Port 6, Fabric 1.
  - d. Connect controller 2B to any port on the alternate fabric.  
In the sample illustration this is Switch 2, Port 16, Fabric 2.
  - e. Connect additional controller ports and fabrics, as required by your MetroCluster configuration.
2. Connect the ports on the storage array at Site B:
  - a. Connect storage array controller 1A to any port on one fabric.  
In the sample illustration, this is Switch 3, Port 9, Fabric 1.
  - b. Connect controller 2A to any port on the alternate fabric.  
In the sample illustration, this is Switch 4, Port 20, Fabric 2.
  - c. Connect additional controller ports and fabric, as required by your MetroCluster configuration.
  - d. Connect storage array controller 1B to any port on one fabric.  
In the sample illustration, this is Switch 3, Port 11, Fabric 1.
  - e. Connect controller 2B to any port on the alternate fabric.  
In the sample illustration, this is Switch 4, Port 21, Fabric 2.



## After you finish

Configure the switch zoning.

### Related concepts

*Recommended fabric-attached MetroCluster configuration with array LUNs* on page 166  
When configuring a fabric-attached MetroCluster configuration with Data ONTAP systems and storage arrays, you should follow the best practice recommendations for how to connect the hardware components. The Interoperability Matrix at [support.netapp.com](http://support.netapp.com) contains information about the hardware components that are supported for a MetroCluster configuration using array LUNs without a PVR.

*Recommended stretch MetroCluster configuration with array LUNs* on page 169  
When configuring a stretch MetroCluster configuration with Data ONTAP systems and storage arrays, you should follow the best practice recommendations for how to connect the hardware components. The Interoperability Matrix at [support.netapp.com](http://support.netapp.com) contains information about the hardware components that are supported for a MetroCluster configuration using array LUNs.

*Cabling guidelines for a MetroCluster configuration with array LUNs* on page 172  
There are a number of cabling guidelines that you need to review before connecting the devices in a MetroCluster configuration with array LUNs.

### Related tasks

*Configuring zoning in a MetroCluster configuration with array LUNs* on page 185  
Switch zoning defines paths between connected nodes based on the node's unique WWN. Single-initiator zoning is recommended for a configuration with array LUNs.

## Configuring zoning in a MetroCluster configuration with array LUNs

Switch zoning defines paths between connected nodes based on the node's unique WWN. Single-initiator zoning is recommended for a configuration with array LUNs.

### Before you begin

- Zoning must be determined for each switch.
- The local and remote Data ONTAP systems must be connected to the fabric.
- The switch fabric must be connected and the switches must be powered up.
- The fabric and storage array must be connected.

### About this task

For a MetroCluster configuration with array LUNs, the FC-VI ports on the FC adapters must be zoned end-to-end across the fabric.

### Steps

1. Gather the WWPNs for the FC-VI ports on the FC-VI adapter of the Data ONTAP system.  
The WWPNs for the FC-VI ports are needed when you set up the zones for the FC-VI ports. The WWPNs for the FC-VI ports are available from the following sources:
  - The switch
  - Data ONTAP `sysconfig -M output`  
Gather WWPNs from entries such as the following: `!Qlogic 2352 FCVI Cluster Interconnect Adapter<adapter_wwpn>`
2. Zone FC-VI port "a" on the local Data ONTAP system to the FC-VI port "a" on the remote Data ONTAP system.
3. Zone FC-VI port "b" on the local Data ONTAP system to the FC-VI port "b" on the remote Data ONTAP system.

### After you finish

- If necessary, change the configuration speed of your stretch MetroCluster configuration.
- Set up Data ONTAP on the Data ONTAP systems so that they can access storage on the storage arrays, and so that you can take advantage of Data ONTAP features.

### Related concepts

*Planning zoning for a MetroCluster configuration with array LUNs* on page 173

Switch zoning defines paths between connected nodes based on the node's unique WWN. Sketching out the zoning plan enables you to correct errors before zones are configured, and makes it easier to communicate zoning information to the person who configures the switches.

### Related tasks

*Testing zoning of FC-VI ports in a MetroCluster configuration with array LUNs* on page 191

To test zoning of the FC-VI ports, you need to check that you have not crossed the FC-VI "a" ports and "b" ports in your zones.

## Setting up a shared-switches configuration

In a shared-switches configuration, two fabric-attached MetroCluster configurations share the same four switches and the ISLs between them.

Shared-switches configuration are cost-effective because you can use the four Brocade switches between two MetroCluster configurations.

## Cabling the FC-VI adapter and ISL in a shared-switches MetroCluster configuration with array LUNs

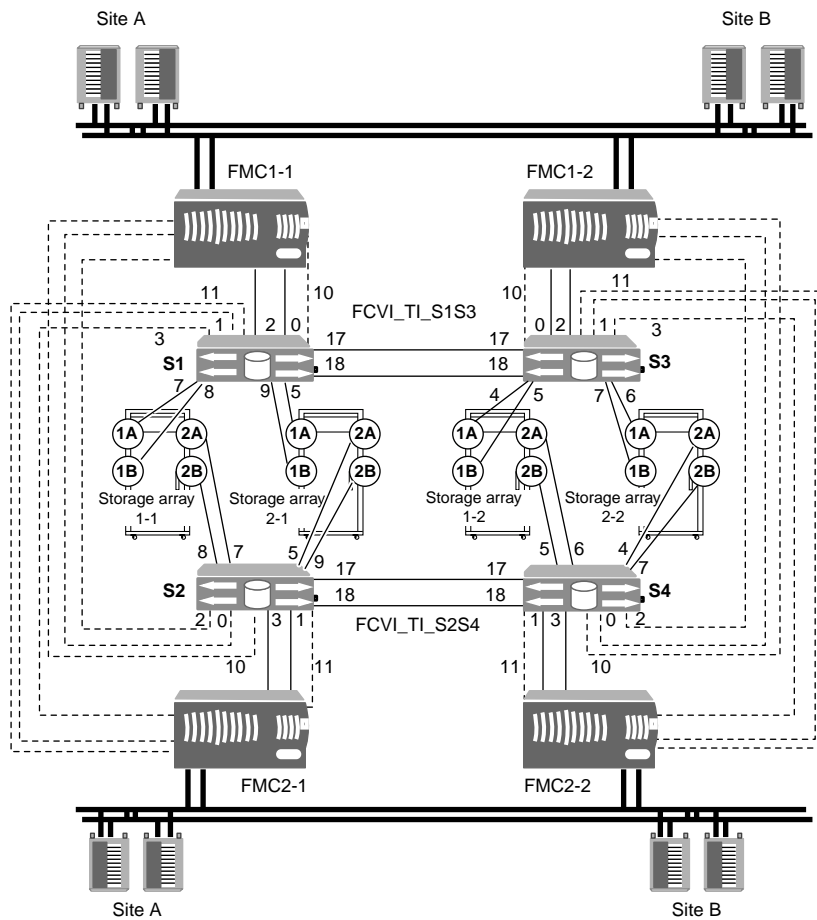
Cabling in a shared-switches MetroCluster configuration with array LUNs involves cabling the FC-VI and HBA ports to the switches, and cabling the switches to the storage arrays.

### About this task

In a shared-switches configuration, four nodes join together to form two fabric-attached MetroCluster configurations. In the illustration, the MetroCluster configurations are as follows:

- FMC1-1 and FMC1-2 form one fabric-attached MetroCluster configuration named FMC1.
- FMC2-1 and FMC2-2 form another fabric-attached MetroCluster configuration named FMC2.

Implementing a stretch or fabric-attached MetroCluster configuration with array LUNs



**Steps**

1. Connect the ports of the FC-VI adapter to primary and secondary switches for both Data ONTAP systems in the MetroCluster configuration.

The following table shows the FC-VI adapter connections for the sample configuration:

| FC-VI adapter on the Data ONTAP system... | Connects to these switches...                            |
|-------------------------------------------|----------------------------------------------------------|
| FMC1-1                                    | Primary switch S1 and secondary switch S2 through port 0 |
| FMC1-2                                    | Primary switch S3 and secondary switch S4 through port 0 |
| FMC2-1                                    | Primary switch S2 and secondary switch S1 through port 1 |
| FMC2-2                                    | Primary switch S3 and secondary switch S4 through port 1 |

2. Connect the FC initiator ports to the primary and secondary switches.

The FC initiator port connections for the sample configuration are shown in the following table:

| FC initiator ports on the Data ONTAP system | Connect to these switches...                                         |
|---------------------------------------------|----------------------------------------------------------------------|
| FMC 1-1                                     | Primary switch S1 and secondary switch S2 through port 2 and port 10 |

| FC initiator ports on the Data ONTAP system | Connect to these switches...                                         |
|---------------------------------------------|----------------------------------------------------------------------|
| FMC 1-2                                     | Primary switch S3 and secondary switch S4 through port 2 and port 10 |
| FMC 2-1                                     | Primary switch S2 and secondary switch S1 through port 3 and port 11 |
| FMC 2-2                                     | Primary switch S3 and secondary switch S4 through port 3 and port 11 |

3. Connect the storage arrays to the switches.

The following table shows the switch-to-storage-array connections for FMC1:

| Storage array and port | Switch and port  |
|------------------------|------------------|
| 1-1: Port 1A           | Switch 1: Port 7 |
| 1-1: Port 1B           | Switch 1: Port 8 |
| 1-1: Port 2A           | Switch 2: Port 7 |
| 1-1: Port 2B           | Switch 2: Port 8 |
| 2-1: Port 1A           | Switch 1: Port 5 |
| 2-1: Port 1B           | Switch 1: Port 9 |
| 2-1: Port 2A           | Switch 2: Port 5 |
| 2-1: Port 2B           | Switch 2: Port 9 |

4. Repeat steps 1 through 3 at another site.

The following table shows the switch-to-storage-array connections for FMC2:

| Storage array and port | Switch and port  |
|------------------------|------------------|
| 1-2: Port 1A           | Switch 3: Port 4 |
| 1-2: Port 1B           | Switch 3: Port 5 |
| 1-2: Port 2A           | Switch 4: Port 6 |
| 1-2: Port 2B           | Switch 4: Port 5 |
| 2-2: Port 1A           | Switch 3: Port 6 |
| 2-2: Port 1B           | Switch 3: Port 7 |
| 2-2: Port 2A           | Switch 4: Port 4 |
| 2-2: Port 2B           | Switch 4: Port 7 |

5. When cabling at both the sites is complete, connect an ISL (inter-switch link) cable to a port on each switch.

In the illustration, the fabric containing switches S1 and S3 has two ISLs and the fabric containing switches S2 and S4 also has two ISLs.

Switch S1 connects to switch S3 and switch S2 connects to switch S4 through ports 17 and 18.

**After you finish**

After you finish cabling the devices, you need to configure traffic isolation on the switches. This configuration is required when you are using dual ISLs. See the *Fabric-attached MetroCluster Systems Brocade Switch Configuration Guide* or the *Fabric-attached MetroCluster Systems Cisco Switch Configuration Guide* for more information.

### Related tasks

[Setting a preferred primary port in a MetroCluster configuration](#) on page 157

After cabling a MetroCluster configuration, you can define the primary port between the four switches to carry the traffic. You can set the preferred primary port in a shared-switches configuration by using the `ic primary set` command.

## Setting a preferred primary port in a MetroCluster configuration

After cabling a MetroCluster configuration, you can define the primary port between the four switches to carry the traffic. You can set the preferred primary port in a shared-switches configuration by using the `ic primary set` command.

### About this task

By default, the primary port is 0.

### Steps

1. To set the primary port, enter the following command:

```
ic primary set 0|1 [-r]
```

You must use the `-r` option for the primary port to take effect immediately.

#### Note:

- If the port that you have set as the preferred primary port is down, the other port will be used. Once the port is up, it will be set as the preferred primary port.
- If the FC-VI link between the preferred primary ports fails, the fabric-attached MetroCluster configuration fails over to the secondary FC-VI link. When the primary link is restored, the fabric-attached MetroCluster configuration again uses the primary link. The failover temporarily causes the logs to be unsynchronized.

If you want to set the preferred primary port to 1, enter the following command:

```
ic primary set 1 -r
```

In a shared-switches configuration, if you have set the preferred primary port of FMC1 (FMC1-1 and FMC1-2) to 1, the preferred primary port for FMC2 (FMC2-1 and FMC2-2) is set to 0.

2. To view the current primary port, enter the following command:

```
ic primary show
```

The primary port is displayed.

## Removing the preferred primary port in a fabric-attached MetroCluster configuration

You can remove the port that you assigned as primary port in a fabric-attached MetroCluster configuration.

### Step

To remove the primary port, enter the following command:

```
ic primary unset
```

## Enabling unidirectional port mirroring in a MetroCluster configuration

The unidirectional port mirroring allows you to use the FC-VI link effectively and therefore increases the performance throughput. By enabling the unidirectional port mirroring, one link in

an FC-VI interconnect will be used for outgoing data traffic and the other for incoming traffic between the two nodes in a MetroCluster configuration.

### Before you begin

- The preferred primary port is disabled on both the nodes.
- The FC-VI links in a MetroCluster configuration is up and working.

### About this task

- The unidirectional port mirroring must be enabled or disabled at both the nodes in a MetroCluster configuration.

### Steps

1. Verify the status of the FC-VI links by entering the following command:

```
sysconfig -v
```

2. For enabling or disabling unidirectional port mirroring, do one of the following:

| If you want to...                          | Then enter the following command...                                                                                                                                                               |
|--------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enable the unidirectional port mirroring.  | <code>ic unidirectional_mirroring on</code><br><b>Note:</b> This is a default behavior. Therefore, you must first check the status of unidirectional port mirroring before enabling this feature. |
| Disable the unidirectional port mirroring. | <code>ic unidirectional_mirroring off</code>                                                                                                                                                      |

**Note:** When reverting to a Data ONTAP release prior to Data ONTAP 8.2.1, the unidirectional port mirroring gets disabled automatically.

3. Verify the status of the unidirectional port mirroring by entering the following command:

```
ic unidirectional_mirroring status
```

If the unidirectional port mirroring is enabled, it displays the following output:

```
Unidirectional mirroring is enabled
For outgoing data primary link is 0 and current link in use is 0.
For incoming data primary link is 1 and current link in use is 1.
```

If the unidirectional port mirroring is disabled, it displays the following output:

```
Unidirectional mirroring is disabled
For outgoing data primary link is 0 and current link in use is 0.
For incoming data primary link is 0 and current link in use is 0.
```

**Note:** When the unidirectional port mirroring is disabled, the primary link and current link in use for both outgoing and incoming traffic will be the same.

## Changing the configuration speed of a stretch MetroCluster configuration

If the distance between nodes in a stretch MetroCluster configuration is greater than the supported default configuration speed, you must change the default configuration speed. If you modified the default configuration speed in a stretch MetroCluster configuration that uses an FC-VI adapter, you can reset the speed to the default configuration speed.

The default maximum speed between nodes in a stretch MetroCluster configuration depends on the speed at which the FC-VI adapter operates (for example, 4 GB or 8 GB). The maximum supported distance is included in the Interoperability Matrix at [support.netapp.com](http://support.netapp.com).

## Implementing a stretch or fabric-attached MetroCluster configuration with array LUNs

The procedures for changing the configuration speed of a stretch MetroCluster configuration are the same for MetroCluster configurations with FAS systems and MetroCluster configurations with V-Series systems.

**Related tasks**

[Changing the default configuration speed of a stretch MetroCluster configuration](#) on page 128  
The distance between your nodes and the FC-VI adapter speed dictates the default configuration speed of your stretch MetroCluster configuration. If the distance between the nodes is greater than what is supported with the default configuration speed, you must change the default configuration speed of the FC-VI adapter.

[Resetting a stretch MetroCluster configuration to the default speed](#) on page 131  
If you modified the default configuration speed in a stretch MetroCluster configuration by using an FC-VI adapter, you can reset the speed to the default configuration speed by using the `unset env` command at the boot environment prompt.

## Setting up Data ONTAP after connecting devices in a MetroCluster configuration with array LUNs

After connecting the devices in the MetroCluster configuration, you need to set up the Data ONTAP systems to use the storage on the storage array. You also need to set up any desired Data ONTAP features.

Perform the following tasks after connecting the devices in your MetroCluster configuration with array LUNs.

1. Set up the Data ONTAP systems as described in the *Data ONTAP Software Setup Guide for 7-Mode* for the systems using only array LUNs.
2. Create one or more mirrored aggregates.  
See the *Data ONTAP Data Protection Online Backup and Recovery Guide for 7-Mode*.
3. Test the MetroCluster configuration, as described in the *Data ONTAP High Availability and MetroCluster Configuration Guide for 7-Mode*.
4. Verify network and protocol setup, as described in the *Data ONTAP File Access and Protocols Management Guide for 7-Mode*.
5. Configure other Data ONTAP features as desired, for example, data protection features.  
See the relevant Data ONTAP guide for the feature you want to set up.

## Testing a MetroCluster configuration with array LUNs

It is important to test that your MetroCluster configuration is properly set up before putting it into a production environment.

**Steps**

1. [Testing zoning of FC-VI ports in a MetroCluster configuration with array LUNs](#) on page 191
2. [Verifying proper setup at MetroCluster sites with storage arrays](#) on page 192
3. [Simulating a disaster recovery in a MetroCluster configuration with array LUNs](#) on page 193

## Testing zoning of FC-VI ports in a MetroCluster configuration with array LUNs

To test zoning of the FC-VI ports, you need to check that you have not crossed the FC-VI "a" ports and "b" ports in your zones.

**Related tasks**

[Configuring zoning in a MetroCluster configuration with array LUNs](#) on page 185

## Implementing a stretch or fabric-attached MetroCluster configuration with array LUNs

Switch zoning defines paths between connected nodes based on the node's unique WWN. Single-initiator zoning is recommended for a configuration with array LUNs.

## Verifying proper setup at MetroCluster sites with storage arrays

After installing a MetroCluster configuration with storage arrays, you need to test the paths, FC initiator ports, and switch configuration at each MetroCluster site to ensure that they are set up correctly.

### Steps

1. Enter the following command to display path information from each Data ONTAP system to the array LUNs:

```
storage show disk -p
```

You should see two paths to each array LUN.

2. To test FC initiator ports, enter the following commands for each FC initiator port in the MetroCluster configuration:

- a. Enter privilege mode:

```
priv set advanced
```

- b. Display the state of the HBA port:

```
fcadmin link_state
```

- c. Take an FC port offline to simulate a port failure or a cable pull:

```
fcadmin offline portname
```

```
fcadmin offline 0a
```

- d. Display disk path information:

```
storage show disk -p
```

The display should show only one path.

- e. Bring the FC port online:

```
fcadmin online portname
```

```
fcadmin online 0a
```

- f. Verify that both paths are online:

```
storage show disk -p
```

3. Simulate a switch failure or a storage array controller failure for each switch fabric by entering the following commands:

- a. Take all FC ports offline on both Data ONTAP systems that are attached to one fabric:

```
fcadmin offline portname
```

```
fcadmin offline 0a
```

```
fcadmin offline 0b
```

```
fcadmin offline 0c
```

```
fcadmin offline 0d
```

If your model has more than four FC initiator ports, take the remaining initiator ports offline.

- b. Verify that all HBAs are disabled:

```
fcadmin link_state
```

- c. Initiate a takeover from site A:

```
cf takeover
```

- d. When site B is in takeover mode, initiate a giveback from Site A:



```
cf giveback
```

## Simulating a disaster recovery in a MetroCluster configuration with array LUNs

Testing for proper setup of a MetroCluster configuration with array LUNs includes simulating a disaster recovery. Simulating a site failure and recovery involves disabling and degrading mirrors.

### About this task

You should never simulate disaster site failure and recovery in production environments without prior planning and downtime.

### Steps

1. Disable the HA interconnect between the Data ONTAP systems.

2. Power down one of the systems.

The system you power down simulates the site of the disaster.

3. On the surviving system, complete the following steps:

- a. Enter the following command to activate a forced, manual takeover:

```
cf forcetakeover -d
```

- b. Enter the following command to validate that the forced takeover has occurred:

```
cf status
```

- c. Enter the following commands to validate the status of the aggregates and of the volumes that they include:

```
aggr status
```

```
vol status
```

The display should show an online status for aggregates and volumes, but mirrors, which are disabled, should be displayed as “degraded.”

4. Reconnect the HA interconnect.

5. On the surviving system, enter the following commands:

- a. Enter the following command to rejoin the aggregates:

```
aggr mirror surviving_aggregate -v victim_aggregate
```

**Note:** Data ONTAP uses parentheses to indicate the degraded aggregate. For example, `aggr0_b` indicates the aggregate that is not degraded, and `aggr0_b(1)` is degraded.

- b. Enter the following command to validate that forced takeover has occurred:

```
cf status
```

- c. Enter the following command to validate aggregate status:

```
aggr status
```

The display should show an online and mirrored status for aggregates.

- d. Enter the following command to validate volume status:

```
vol status
```

The display should show an online and mirrored status for volumes.

6. Power up the system that you powered down to simulate the disaster.

7. On the surviving system, enter the following command to reactivate the site that simulated the disaster:

```
cf giveback
```

8. To validate that the HA pair, aggregates, and mirrors are online and operational, enter the following commands at both sites:

**cf status**

**aggr status**

**vol status**

The display should indicate that the HA pair is enabled, and that aggregates and volumes are online and mirrored.

## Implementing a MetroCluster configuration with both disks and array LUNs

The systems that can use both disks and array LUNs in a MetroCluster configuration varies according to type of system and Data ONTAP release.

### Types of systems that can use array LUNs in a MetroCluster configuration

Support for which systems can use array LUNs in a MetroCluster configuration differs according to Data ONTAP release.

| Data ONTAP release | Systems that can use array LUNs in a MetroCluster configuration                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Prior to 8.2       | V-Series systems only, and those systems can use only array LUNs                                                                                                                                                                |
| 8.2                | V-Series systems only, and those systems can use array LUNs, disks, or both                                                                                                                                                     |
| 8.2.1 and later    | <ul style="list-style-type: none"><li>V-Series systems, which can use array LUNs, disks, or both</li><li>FAS systems that can be licensed to use array LUNs</li></ul> These systems can use disks also (or instead) if desired. |

**Note:** V-Series ("V") systems and new FAS platforms released in Data ONTAP 8.2.1 and later can use array LUNs if the proper license is installed. In the documentation, these systems are collectively referred to as Data ONTAP systems.

### Planning a MetroCluster configuration with disks and array LUNs

When planning a MetroCluster configuration with systems using both disks and array LUNs, you need to consider the requirements for using each type of storage in the configuration and the additional guidelines for using both types of storage in a MetroCluster configuration.

### Guidelines for implementing a MetroCluster configuration with disks and array LUNs

When planning your MetroCluster configuration for use with disks and array LUNs, you need to follow the guidelines for setting up a MetroCluster configuration for each type of storage, plus the additional guidelines for setting up a MetroCluster configuration with both types of storage.

The basic requirements for using disks and array LUNs in a MetroCluster configuration are the same whether you are using just one type of storage or both. Likewise, installation, configuration, and testing of each type of storage is the same as if you were using just one type of storage.

However, there are additional guidelines to consider if you are using both disks and array LUNs in your MetroCluster configuration.

Consider the information in the following table when planning your configuration:

| Consideration                             | Guideline                                                                                                                                                                                         |
|-------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Order of setting up access to the storage | You can set up access to either type of storage first. You should complete all setup for that type of storage and verify that it is set up correctly before setting up the other type of storage. |

| Consideration                                                                                 | Guideline                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Root volume location                                                                          | <ul style="list-style-type: none"> <li>• If you are setting up a <i>new</i> MetroCluster deployment with both disks and array LUNs, it is recommended that you create the root volume on native disks.</li> <li>• If you are adding native disks to an <i>existing</i> MetroCluster configuration that uses array LUNs, the root volume can remain on an array LUN.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| FC initiator port usage                                                                       | <p>You need to use different FC initiator ports to connect to a FibreBridge 6500N bridge than you use to connect to the switches that connect to the storage arrays.</p> <p>When FibreBridge 6500N bridges are used in a MetroCluster configuration with disks and array LUNs, four FC connections are required for native disk shelves connected through the FibreBridge 6500N and four FC connections are required for the FC connections to target ports on the storage arrays. Therefore, a minimum of eight FC initiator ports are required for a Data ONTAP system to connect to both native disks and array LUNs. Not all platforms can support eight FC initiator ports. See the <i>Hardware Universe</i> at <a href="http://hww.netapp.com">hww.netapp.com</a> to determine whether your platform can support eight FC initiator ports.</p> |
| Switches and switch port usage in MetroCluster configurations using FibreBridge 6500N bridges | <p>For stretch MetroCluster configurations, only direct connection to a FibreBridge 6500N bridge is supported. Switches are used to connect to the storage arrays.</p> <p>For fabric-attached MetroCluster configurations, the rules for switches are as follows:</p> <ul style="list-style-type: none"> <li>• The same switches are used to connect to the storage arrays and the FibreBridge 6500N bridges.</li> <li>• The switch ports used to connect to a FibreBridge 6500N bridge must be different from the switch ports used to connect to a storage array.</li> </ul>                                                                                                                                                                                                                                                                       |
| Reverting<br>(Transitioning a storage system to a Data ONTAP release in an earlier family)    | <p>You cannot revert a system set up with a mixed disk and array LUN MetroCluster configuration to a release prior to Data ONTAP 8.2 because a mixed disk and array LUN MetroCluster configuration is not supported prior to Data ONTAP 8.2.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

**Related concepts**

[Types of systems that can use array LUNs in a MetroCluster configuration](#) on page 163

Support for which systems can use array LUNs in a MetroCluster configuration differs according to Data ONTAP release.

[Planning a fabric-attached MetroCluster configuration using disks](#) on page 133

Creating a detailed plan for your fabric-attached MetroCluster configuration helps you understand the unique requirements for configuring a fabric-attached MetroCluster configuration with disk. Installing a fabric-attached MetroCluster configuration involves connecting and configuring a number of devices, which might be done by different people.

[Setup requirements and restrictions for fabric-attached MetroCluster configurations with disks](#) on page 133

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

[Planning a stretch MetroCluster configuration using disks](#) on page 118

Creating a detailed plan for your MetroCluster configuration helps you understand the unique requirements for configuring a stretch MetroCluster configuration with disk. For installing a stretch MetroCluster configuration, you should be aware of the requirements (such as the tools and equipment) and restrictions for a stretch MetroCluster configuration.

*Hardware requirements for using disks in a MetroCluster configuration with V-Series systems* on page 119

Starting in Data ONTAP 8.2, a MetroCluster configuration with V-Series systems can include native disks only, array LUNs on storage arrays only, or both. There are specific hardware requirements to set up a MetroCluster with V-Series systems with native disks.

*Setup requirements and restrictions for stretch MetroCluster configurations with disks* on page 118

You must follow certain requirements and restrictions when setting up a new stretch MetroCluster configuration with disks.

*Planning for a MetroCluster configuration with array LUNs* on page 162

Creating a detailed plan for your MetroCluster configuration helps you understand the unique requirements for a MetroCluster configuration that uses LUNs on storage arrays. A plan also helps you communicate with other people involved in the installation. Installing a MetroCluster configuration involves connecting and configuring a number of devices, which might be done by different people.

*Requirements for a MetroCluster configuration with array LUNs* on page 164

There are some unique requirements for setting up a MetroCluster configuration with array LUNs.

*Recommended fabric-attached MetroCluster configuration with array LUNs* on page 166

When configuring a fabric-attached MetroCluster configuration with Data ONTAP systems and storage arrays, you should follow the best practice recommendations for how to connect the hardware components. The Interoperability Matrix at [support.netapp.com](http://support.netapp.com) contains information about the hardware components that are supported for a MetroCluster configuration using array LUNs without a PVR.

*Recommended stretch MetroCluster configuration with array LUNs* on page 169

When configuring a stretch MetroCluster configuration with Data ONTAP systems and storage arrays, you should follow the best practice recommendations for how to connect the hardware components. The Interoperability Matrix at [support.netapp.com](http://support.netapp.com) contains information about the hardware components that are supported for a MetroCluster configuration using array LUNs.

## Supported fabric-attached MetroCluster configuration with disks and array LUNs


When configuring a MetroCluster configuration with disks and storage arrays, you need to follow the best practice recommendations for how to connect the hardware components.

The key points to remember about connections in a fabric-attached MetroCluster configuration with disks and storage arrays are as follows:


- You use the same switches to connect a Data ONTAP system to the FibreBridge 6500N bridges and the storage array.
- You use different switch ports for the connections to the FibreBridge 6500N bridges than for the connections to the storage array.
- You use different FC initiator ports on the Data ONTAP system for the switch ports connecting to the FibreBridge 6500N bridges than for the switch ports connecting to the storage array.  
If your system does not have enough FC initiator ports for separate connections to the bridges and storage array, you must add some FC initiator ports.

The Interoperability Matrix at [support.netapp.com](http://support.netapp.com) contains information about the hardware components that are supported for a MetroCluster configuration with storage arrays. *Configuring a MetroCluster system with SAS disk shelves and FibreBridge 6500N bridges* contains information about the connections between the FibreBridge 6500N bridges and the disk shelves.

The following block diagram shows a fabric-attached MetroCluster configuration that uses only array LUNs:



The following block diagram shows the same fabric-attached MetroCluster, but with disks as well. Notice that the same switches are used for access to the storage arrays and the FibreBridge 6500N bridges, but different switch ports are used for the bridges and the storage arrays. Also notice that different FC initiator ports are used for the two types of storage:



### Supported stretch MetroCluster configuration with disks and array LUNs

When configuring a MetroCluster configuration with both disks and array LUNs, you need to follow the best practice recommendations for how to connect the hardware components.

#### Connections when using FibreBridge 6500N bridges


When using FibreBridge 6500N bridges in a stretch MetroCluster configuration with disks and storage arrays, the key points to remember about the connections are as follows:

- The Data ONTAP system connects directly to the FibreBridge 6500N bridges, but connects to the storage array through switches.
- Connecting the Data ONTAP system to the FibreBridge 6500N bridges through switches is not supported.
- Different FC initiator ports on the Data ONTAP system must be used for the switch ports connecting to the FibreBridge 6500N bridges than the FC initiator ports for connecting to the storage array.

If your Data ONTAP system does not have enough FC initiator ports for separate connections to the bridges and storage array, you must add some FC initiator ports.

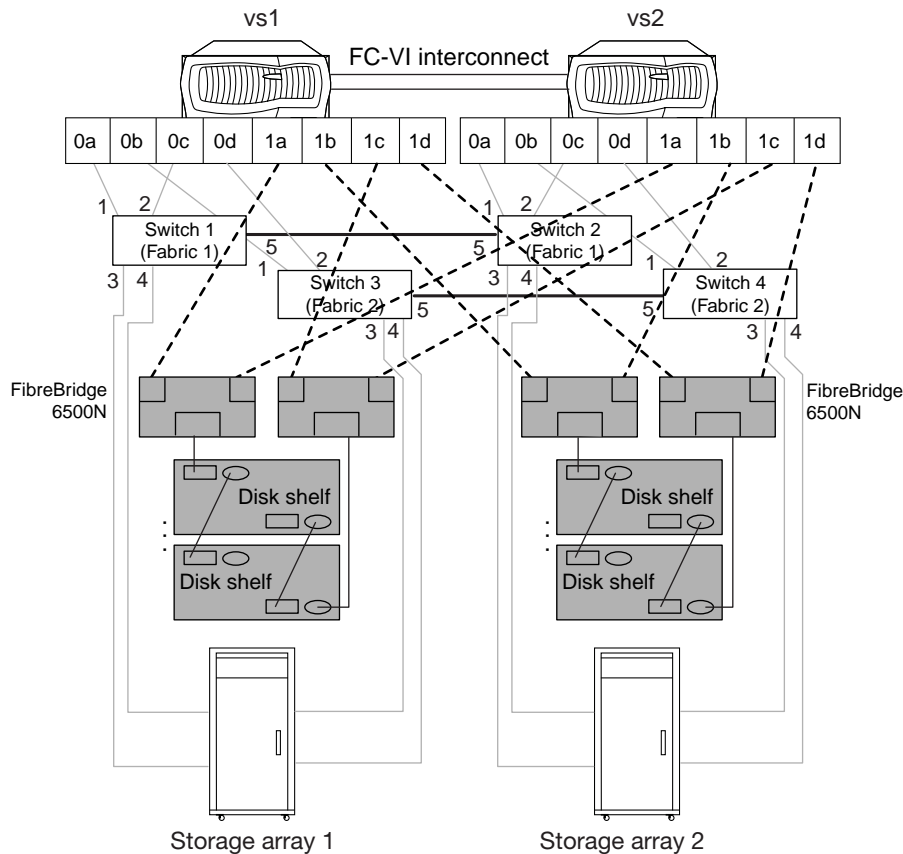
The Interoperability Matrix at [support.netapp.com](http://support.netapp.com) contains information about the hardware components that are supported for a MetroCluster configuration with storage arrays. *Configuring a MetroCluster system with SAS disk shelves and FibreBridge 6500N bridges* contains information about the connections between the FibreBridge 6500N bridges and the disk shelves.

The following block diagram shows a stretch MetroCluster configuration with just array LUNs:



The following block diagram shows the connectivity when FibreBridge 6500N bridges and disk shelves are added to the stretch MetroCluster configuration from the previous illustration. The connections for accessing array LUNs are purposely grayed out to make it easier for you to see the connections for the FibreBridge bridges:





As the illustration shows, the stretch MetroCluster uses direct FC connections between the Data ONTAP systems and the FibreBridge 6500N bridges. Eight FC initiator ports are required on each Data ONTAP system, four for the FibreBridge 6500N bridges and four for the connections to the storage array via the switches.

### Connections when using SAS optical cables

There are no unique requirements when connecting SAS optical cables in a MetroCluster configuration that uses both SAS disk shelves and LUNs on storage arrays. See *Configuring a stretch MetroCluster system with SAS disk shelves and SAS optical cables* on the NetApp Support Site at [support.netapp.com](http://support.netapp.com) for hardware requirements, illustrations, and procedures for installing a new system with SAS optical cables or replacing FibreBridge 6500N bridges with SAS optical cables.

**Note:** SAS optical cables can be used to cable SAS disk shelves in a stretch MetroCluster system to achieve greater connectivity over distance than is possible with FibreBridge 6500N bridges. SAS optical cables can be used in a new system installation or in an existing system in which you want to replace the SAS copper cables and FibreBridge 6500N bridges. SAS optical cables are supported for use with Data ONTAP systems starting in Data ONTAP 8.2.1.

## Testing a MetroCluster configuration with disks and array LUNs

You should test your MetroCluster configuration in stages, testing that each type of storage is working properly before going onto installation for the next type of storage.

### Related concepts

[Testing a MetroCluster configuration with array LUNs](#) on page 191

It is important to test that your MetroCluster configuration is properly set up before putting it into a production environment.

# Recovering from a disaster by using MetroCluster configurations

---

In situations such as prolonged power outages or natural disasters, you can use the 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.

## 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

## 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

There are some failures that do not require any disaster recovery such as, a failure of the interconnect, failure of cables and so on. If you can reestablish the MetroCluster connection after fixing the problem, you should not perform the disaster recovery procedure.

You should not perform the disaster recovery procedure for the following failures:

- A failure of the HA interconnect between the two sites, which 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 configuration returns to normal operation.

- The storage from one site (site A) is not accessible to the node at the other site (site B), which 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 configuration 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 system 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 configuration returns to normal operation.
- If you are using FibreBridge 6500N bridge, a failure of the FibreBridge 6500N bridge.

## 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 node 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 ensure 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 for 7-Mode*.

### Steps

1. [Restricting access to the disaster site node](#) on page 204
2. [Forcing a node into takeover mode](#) on page 204
3. [Remounting volumes of the failed node](#) on page 205
4. [Recovering LUNs of the failed node](#) on page 205
5. [Fixing failures caused by the disaster](#) on page 206
6. [Reestablishing the MetroCluster configuration](#) on page 206

## 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.

### Steps

1. [Restricting access to the node by turning off power](#) on page 204
2. [Restricting access to the node by fencing off](#) on page 204

## 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

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 HA 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:

| If you are using...  | Then fencing is achieved by...                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Application failover | 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. |
| IP failover          | 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

Enter the following command on the surviving node:

```
cf forcetakeover -d
```

### Result

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_fsid` 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 *Clustered Data ONTAP File Access and Protocols Management Guide*.

**Note:** You can disable the `change_fsid` 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 by entering the following command.

```
mkdir directory_path
mkdir /n/toaster/home
```

2. Mount the volume by entering the following command.

```
mount volume_name
mount toaster:/vol/vol0/home /n/toaster/home
```

### Related tasks

[Disabling the change\\_fsid option in MetroCluster configurations](#) on page 56

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.

## 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_fsid` 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_fs_id` 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 SAN Administration Guide for 7-Mode*.

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.

```
lun online /vol/vol1/lun5
```

```
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.

### About this task

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 partner for a MetroCluster configuration at a different site.

### Step

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

You can reestablish a MetroCluster configuration 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.

### Rejoining the mirrored aggregates to reestablish a MetroCluster configuration

You must rejoin the mirrored aggregates if the mirrored aggregate was in a normal state before the forced takeover.

#### About this task

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
```

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 is shown, in this case after access to the partner node at the surviving site is reestablished:

```
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.snapcrtok: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
```



### Related tasks

[Re-creating mirrored aggregates to return a MetroCluster configuration to normal operation](#) on page 209

To return the MetroCluster configuration to normal operation, you must re-create the mirrored aggregates.

## Re-creating mirrored aggregates to return a MetroCluster configuration to normal operation

To return the MetroCluster configuration to normal operation, you must re-create the mirrored aggregates.

### 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...(Press Ctrl-C to abort wait)
```

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
```

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. 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.

The following output shows the aggregate status of the disaster site's storage after reestablishing access to the partner at the surviving site:

```
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.
```

### Related tasks

[Rejoining the mirrored aggregates to reestablish a MetroCluster configuration](#) on page 207

You must rejoin the mirrored aggregates if the mirrored aggregate was in a normal state before the forced takeover.

## Feature update record

This record provides the history of changes made to this guide. When a change is made, 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 HA pairs, 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 7-Mode 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         |
| <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        |

| Feature updates                                                                                                                                                                                                                                                                                                                                                                                                                                         | Feature first implemented in | Feature release date |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|----------------------|
| <ul style="list-style-type: none"> <li>• Updated and extended HA pair 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 HA pairs.</li> </ul>                                                                                                                          | Data ONTAP 7.1.1             | June 2006            |
| <ul style="list-style-type: none"> <li>• Generalized standard and mirrored HA pair cabling instructions</li> <li>• Updated standard and mirrored HA pair 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 HA pairs 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             |
| <ul style="list-style-type: none"> <li>• Added information about Multipath Storage for HA pairs</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           |

| Feature updates                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Feature first implemented in | Feature release date |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|----------------------|
| <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 HA pair</li> </ul>                                                                                                                                                                                             | Data ONTAP 7.2.4             | November 2007        |
| <ul style="list-style-type: none"> <li>• Changed name of document from <i>Active/Active Configuration Guide</i> to <i>High-Availability Configuration Guide</i></li> <li>• Moved from <i>active/active configuration</i> to <i>High Availability pair</i></li> <li>• Added information about configuration on V-Series systems</li> <li>• Added support for 672 disks in MetroClusters</li> <li>• Added MetroCluster support for the Brocade 300 and 5100 switches</li> </ul> | Data ONTAP 8.0 RC1           | June 2009            |
| <ul style="list-style-type: none"> <li>• Added references to the DS4243 disk shelf documentation</li> </ul>                                                                                                                                                                                                                                                                                                                                                                   | Data ONTAP 8.0 GA            | January 2010         |
| <ul style="list-style-type: none"> <li>• Added support for 32xx and 62xx systems</li> <li>• Added support for the 8-Gbps FC-VI adapter on MetroCluster configurations</li> </ul>                                                                                                                                                                                                                                                                                              | Data ONTAP 8.0.1             | August 2010          |
| <ul style="list-style-type: none"> <li>• Added support for 32xx and 62xx systems.</li> <li>• Added support for the 8-Gbps FC-VI adapter on MetroCluster configurations.</li> </ul>                                                                                                                                                                                                                                                                                            | Data ONTAP 7.3.5             | October 2010         |

| Feature updates                                                                                                                                                                                                                                                                                                                                                                          | Feature first implemented in           | Feature release date |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------|----------------------|
| <ul style="list-style-type: none"> <li>Added more SAS disk shelf information and references to SAS disk shelf documentation</li> <li>Updated multipath HA as a requirement</li> <li>Removed procedures for non-multipath HA cabling</li> <li>Replaced the term <i>Multipath Storage</i> with <i>multipath HA</i> for consistency with other documentation</li> </ul>                     | Data ONTAP 8.0.2<br>Data ONTAP 7.3.5.1 | May 2011             |
| <ul style="list-style-type: none"> <li>Added support for MetroCluster with SAS disk shelves</li> <li>Added support for MetroCluster shared-switches configurations</li> <li>Added procedures for MetroCluster installation with third-party storage</li> <li>Added support for FAS2240 systems</li> </ul>                                                                                | Data ONTAP 8.1 RC1                     | September 2011       |
| <ul style="list-style-type: none"> <li>Removed references to ESH2 and LRC modules</li> </ul>                                                                                                                                                                                                                                                                                             | Data ONTAP 8.1 GA                      | December 2011        |
| <ul style="list-style-type: none"> <li>Changed title to <i>Data ONTAP High Availability and MetroCluster Configuration Guide for 7-Mode</i></li> <li>Added MetroCluster support for Cisco switches</li> </ul>                                                                                                                                                                            | Data ONTAP 8.1.1 RC1                   | May 2012             |
| <ul style="list-style-type: none"> <li>Added support for automatic giveback after takeover-on-panic</li> </ul>                                                                                                                                                                                                                                                                           | Data ONTAP 8.1.2 RC1                   | October 2012         |
| <ul style="list-style-type: none"> <li>Added support for 3220 and 3250 systems</li> </ul>                                                                                                                                                                                                                                                                                                | Data ONTAP 8.1.2 GA                    | November 2012        |
| <ul style="list-style-type: none"> <li>Added support for the <code>cf.mode</code> option</li> <li>Added coverage for the MetroCluster configuration requirement for both the <code>cf.mode</code> and <code>cf.remote_syncmirror.enable</code> options</li> <li>Added support for native disks in addition to LUNs on storage arrays for V-Series MetroCluster configurations</li> </ul> | Data ONTAP 8.2 RC1                     | April 2013           |

| Feature updates                                                                                                                                                                                                                               | Feature first implemented in | Feature release date |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|----------------------|
| <ul style="list-style-type: none"> <li>• Added coverage for disk pool assignment in a stretch MetroCluster configuration</li> </ul>                                                                                                           | Data ONTAP 8.2 GA            | May 2013             |
| <ul style="list-style-type: none"> <li>• Added coverage for nondisruptive shelf removal</li> <li>• Added support for SAS optical cables</li> <li>• Transferred troubleshooting topics to NetApp Support Site</li> </ul>                       | Data ONTAP 8.2.1 RC1         | November 2013        |
| <ul style="list-style-type: none"> <li>• Added support for FAS80xx systems</li> <li>• Added coverage for FlexArray Virtualization Software transition</li> <li>• Added the <i>Verifying the MetroCluster Configuration</i> chapter</li> </ul> | Data ONTAP 8.2.1 GA          | February 2014        |

## Copyright

---

Copyright © 2020 NetApp, Inc. All rights reserved. Printed in the U.S.

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. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).



## Trademark

---

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<http://www.netapp.com/us/legal/netapptmlist.aspx>

## How to send comments about documentation and receive update notifications

---

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

[\*doccomments@netapp.com\*](mailto:doccomments@netapp.com)

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 1395 Crossman Ave., Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

# Index

## A

- About
  - unidirectional port mirroring [158](#)
- active/active configurations
  - status messages [73](#)
- active/passive configuration [27](#)
- adapters
  - MetroCluster configuration requirements [123](#), [140](#)
  - quad-port Fibre Channel HBA [36](#), [41](#)
- aggregates
  - ownership change [21](#), [78](#)
  - re-creating mirrored after disaster [209](#)
  - rejoining after disaster [207](#)
  - root [23](#)
- automatic giveback
  - enhancing speed of [93](#)
- automatic takeover
  - disabling [83](#)
  - enabling [83](#)
  - reasons for [79](#)

## B

- background disk firmware update [22](#)
- best practices
  - cf options [55](#)
  - HA configuration [23](#)
- bring up
  - manually setting options for [54](#)
- Brocade switch configuration [145](#)
- Brocade switches
  - Cisco switches
    - fabric-attached MetroCluster configurations
      - use of Brocade and Cisco switches [137](#)
    - how used by fabric-attached MetroCluster configurations [137](#)
  - how used by fabric-attached MetroCluster configurations [137](#)

## C

- cabinets
  - preparing for cabling [35](#)
- cable [34](#)
- cabling
  - Channel A
    - for mirrored HA pairs [42](#)
    - for standard HA pairs [37](#)
  - Channel B
    - for mirrored HA pairs [44](#)
    - for standard HA pairs [38](#)
  - cluster interconnect for fabric-attached MetroCluster configurations

- cabling (*continued*)
  - cluster interconnect for fabric-attached MetroCluster configurations (*continued*)
    - with software-based disk ownership [151](#)
  - cross-cabled cluster interconnect [40](#), [48](#)
  - cross-cabled HA interconnect [40](#), [47](#)
  - error message, cross-cabled cluster interconnect [40](#), [47](#), [48](#)
  - fabric-attached MetroCluster configurations [143](#)
  - FC-VI adapter for fabric-attached MetroCluster configurations
    - with software-based disk ownership [151](#)
  - guidelines for MetroCluster configurations with array LUNs [172](#)
  - HA interconnect for fabric-attached MetroCluster configurations [148](#)
  - HA interconnect for standard HA pair [40](#), [47](#)
  - HA interconnect for standard HA pair, 32xx systems [40](#), [48](#)
  - HA pairs [32](#)
  - local controller in fabric-attached MetroCluster configuration
    - with software-based disk ownership [145](#)
  - local disk shelves in fabric-attached MetroCluster configurations
    - with software-based disk ownership [146](#)
  - local node [145](#)
  - Node A [145](#)
  - Node B [149](#)
  - preparing equipment racks for [35](#)
  - preparing system cabinets for [35](#)
  - remote controller in fabric-attached MetroCluster configuration
    - with software-based disk ownership [149](#)
  - remote disk shelves in fabric-attached MetroCluster configuration [150](#)
  - remote node [149](#)
  - requirements [34](#)
  - requirements for MetroCluster configurations [123](#), [140](#)
  - stretch MetroCluster configurations [125](#)
  - with software-based disk ownership [150](#)
- cabling disk shelves to switch
  - with software-based disk ownership [150](#)
- cf forcegiveback command
  - using to force giveback [91](#)
- cf.giveback.auto.cifs.terminate.minutes options [93](#)
- cf.giveback.check.partner option [89](#)
- cf.takeover.use\_mcrc\_file [71](#)
- change\_fsid option [56](#)
- Channel A
  - cabling [37](#), [42](#)
  - defined [25](#)
- Channel B
  - cabling [38](#), [44](#)
- chassis configurations, single or dual [28](#)
- CIFS clients and giveback delay [93](#)
- CIFS sessions terminated on takeover [21](#), [78](#)
- cluster interconnect, cabling [148](#), [151](#)
- collecting information
  - about various components of MetroCluster configuration [159](#)

- command exceptions
  - emulated node [87](#)
- commands
  - accessing takeover node with partner [86](#)
  - cf (enables and disables takeover) [82](#)
  - cf forcesgiveback (forces giveback) [91](#)
  - cf forcetakeover -d (forces takeover) [204](#)
  - cf forcetakeover (forces takeover) [81](#)
  - cf giveback (enables giveback) [71](#), [160](#)
  - cf partner (displays partner's name) [76](#)
  - cf status (displays status) [84](#)
  - cf takeover (initiates takeover) [81](#)
  - cf takeover (initiates takeover) [71](#), [160](#)
  - disabling HA mode [95](#)
  - enabling non-HA mode [95](#)
  - exceptions for emulated node [87](#)
  - for monitoring HA pair status [73](#)
  - ha-config modify [52](#)
  - ha-config show [52](#)
  - halt (halts system without takeover) [82](#)
  - ifconfig [66](#)
  - options cf.mode ha | non-ha [53](#)
  - partner (accesses emulated node) [85](#)
  - storage show disk -p (displays paths) [104](#)
  - sysconfig [77](#)
  - takeover (description of all takeover commands) [81](#)
- comments
  - how to send feedback about documentation [218](#)
- comparison of types of HA pairs [13](#)
- components of MetroCluster configuration
  - gathering information about [159](#)
- Config Advisor
  - checking for common configuration errors with [71](#)
  - downloading and running [71](#)
- configuration
  - limitations for fabric-attached MetroCluster with disks [135](#)
- configuration speeds
  - changing stretch MetroCluster configuration default [128](#)
  - changing stretch MetroCluster default [190](#)
- configuration variations
  - fabric-attached MetroCluster configurations [116](#)
  - mirrored HA pairs [14](#)
  - standard HA pairs [27](#)
  - stretch MetroCluster configurations [115](#)
- configurations
  - HA differences between supported system [29](#)
  - MetroCluster
    - overview of implementing with array LUNs [163](#)
    - reestablishing MetroCluster configuration [206](#)
    - systems that support array LUNs in MetroCluster [163](#), [195](#)
    - testing takeover and giveback [71](#), [160](#)
- configuring
  - automatic giveback [92](#)
  - interfaces [66](#)
  - shared interfaces [65](#)
- connections

- connections (*continued*)
  - between a V-Series system and a FibreBridge 6500N bridge [119](#), [135](#)
- considerations
  - for running setup on HA pairs [50](#)
- controller failover
  - benefits of [9](#)
  - monitoring status of [73](#)
- controller failovers
  - events that trigger [16](#)
- controller-to-switch cabling, fabric-attached MetroCluster configurations [145](#), [149](#)

## D

- Data ONTAP
  - in MetroCluster configurations [114](#)
  - upgrading nondisruptively [108](#)
  - upgrading nondisruptively, documentation for [108](#)
- dedicated interfaces
  - configuring using setup [51](#)
  - described [60](#)
  - diagram [63](#)
- default configuration speed
  - managing [128](#)
- delay, specifying before takeover [83](#)
- disabling takeover (cf) [82](#)
- disaster recovery
  - MetroCluster configurations
    - testing for disaster recovery [161](#)
  - testing the MetroCluster configuration for [161](#)
  - using MetroCluster configurations [202](#)
  - when not to perform [202](#)
- disasters
  - determining whether one occurred [202](#)
  - recognizing [202](#)
  - recovering from, with MetroCluster configuration [203](#)
  - recovery from
    - forcing takeover [204](#)
    - manually fencing off the disaster site node [204](#)
    - reestablishing MetroCluster configuration [206](#)
    - restricting access to the failed node [204](#)
- disk firmware update [22](#)
- disk information, displaying [77](#)
- disk paths, verifying in a fabric-attached MetroCluster configuration
  - with software-based disk ownership [153](#)
- disk shelf pool assignments, fabric-attached MetroCluster configurations [152](#)
- disk shelves
  - about modules for [103](#)
  - adding to an HA pair with multipath HA [102](#)
  - hot swapping modules in [106](#)
  - hot-removing [108](#)
  - managing in a MetroCluster configuration [102](#)
  - managing in an HA pair [102](#)
- disk-shelf-to-switch cabling
  - fabric-attached MetroCluster configurations [146](#)

disk-shelf-to-switch cabling, fabric-attached MetroCluster configurations [150](#)

distances between nodes in the HA configuration [13](#)

documentation

- how to receive automatic notification of changes to [218](#)
- how to send feedback about [218](#)
- required for MetroCluster configurations [122](#), [139](#)

documentation, required [33](#)

dual-chassis HA configurations

- diagram of [28](#)
- interconnect [28](#)

Duplicate Address Detection (DAD) [64](#)

## E

e0M management interface [64](#)

eliminating single point of failure [10](#)

EMS message, takeover impossible [23](#)

emulated node

- accessing from the takeover node [85](#)
- backing up [88](#)
- description of [85](#)
- dumps and restores [88](#)
- managing [85](#)

emulated node, accessing with Remote Shell [87](#)

emulated nodes

- command exceptions for [87](#)

enabling HA mode capability and controller failover [53](#)

enabling takeover (cf) [82](#)

equipment racks

- installation in [32](#)
- preparation of [35](#)

events

- table of failover triggering [16](#)

exceptions

- emulated node command [87](#)

## F

fabric-attached MetroCluster configuration

- assigning disk pools [152](#)
- cabling cluster interconnect for
  - cabling FC-VI adapter for
    - with software-based disk ownership [151](#)
    - with software-based disk ownership [151](#)
  - cabling disk shelves to switch
    - with software-based disk ownership [146](#)
  - cabling HA interconnect for
    - cabling FC-VI adapter for [148](#)
- for disks [137](#)
- implementing, with disks [133](#)
- local node
  - cabling controller to switch
    - with software-based disk ownership [145](#)
- planning [133](#)
- remote node
  - cabling controller to switch

fabric-attached MetroCluster configuration (*continued*)

- remote node (*continued*)
  - cabling controller to switch (*continued*)
    - with software-based disk ownership [149](#)
  - verifying disk paths
    - with software-based disk ownership [153](#)
- fabric-attached MetroCluster configurations
  - advantages of [116](#)
  - array LUNs
    - recommended configuration [166](#)
    - requirements [164](#)
  - Brocade switch configuration [145](#)
  - cabling [143](#), [145](#), [146](#), [149](#), [150](#)
  - illustration of [143](#)
  - limitations [135](#)
  - planning worksheet [138](#)
  - restrictions [133](#)
  - setup requirements for [133](#)
  - storage arrays
    - recommended configuration [197](#)
  - variations [116](#)
- failover
  - benefits of controller [9](#)
  - determining status (cf status) [84](#)
  - monitoring status of [73](#)
- failovers
  - events that trigger [16](#)
- failures
  - table of failover triggering [16](#)
- fault tolerance [8](#)
- FC-VI adapter, cabling [148](#), [151](#)
- feedback
  - how to send comments about documentation [218](#)
- fencing, manual [204](#)
- Fibre Channel ports
  - identifying for HA pair [36](#), [41](#)
- Fibre Channel switches
  - MetroCluster configuration requirements [123](#), [140](#)
- FibreBridge failure
  - verifying the behavior [161](#)
- forcing
  - giveback [91](#)
  - takeover [81](#)
- FRU replacement, nondisruptive
  - documentation for [108](#)

## G

giveback

- cf.giveback.check.partner option and [89](#)
- configuring automatic [92](#)
- definition of [15](#)
- delay time for CIFS clients [93](#)
- disabling automatic, after takeover on panic [94](#)
- enabling automatic, after takeover on panic [94](#)
- enhancing speed of automatic [93](#)
- forcing [91](#)

giveback (*continued*)  
  initiating normal [90](#)  
  interrupted [92](#)  
  managing [89](#)  
  partial-giveback [92](#)  
  performing a [89](#)  
  setting to terminate long-running processes [94](#)  
  testing [71](#), [160](#)  
  what happens during [21](#), [78](#)  
guidelines  
  MetroCluster configuration with disks and array LUNs [195](#)

## H

HA configurations  
  benefits of [8](#)  
  converting to MetroCluster configuration [142](#)  
  definition of [8](#)  
  differences between supported system [29](#)  
  single- and dual-chassis [28](#)  
HA interconnect  
  cabling [40](#), [47](#)  
  cabling, 32xx dual-chassis HA configurations [40](#), [48](#)  
  single-chassis and dual-chassis HA configurations [28](#)  
HA mode  
  disabling [95](#)  
  enabling [53](#)  
HA pairs  
  cabling [32](#), [35](#)  
  cabling mirrored [41](#)  
  changing nodes to stand-alone [95](#), [97](#), [99](#)  
  events that trigger failover in [16](#)  
  installation [32](#)  
  managing disk shelves in [102](#)  
  required connections for using UPSs with [48](#)  
  setup requirements [24](#)  
  setup restrictions [24](#)  
  status messages [73](#)  
  types of  
    compared [13](#)  
    installed in equipment racks [32](#)  
    installed in system cabinets [32](#)  
    mirrored [14](#)  
HA state  
  chassis [52](#)  
  controller modules [52](#)  
ha-config modify command  
  modifying the HA state [52](#)  
ha-config show command  
  verifying the HA state [52](#)  
halting system without takeover [82](#)  
hardware  
  components described [12](#)  
  HA components described [12](#)  
  single point of failure [9](#)  
hardware assisted takeover [20](#), [27](#), [57](#)

hardware replacement, nondisruptive  
  documentation for [108](#)  
hardware-assisted takeover  
  checking statistics [75](#)  
  checking status of [74](#)  
hardware-assisted takeovers  
  setting partner IP address for [58](#)  
history of changes to this guide [211](#)  
hot-removing  
  disk shelves [108](#)

## I

ifconfig command [64](#)  
implementing  
  fabric-attached MetroCluster configuration [133](#)  
information  
  how to send feedback about improving documentation [218](#)  
installation  
  equipment rack [32](#)  
  HA pairs [32](#)  
  system cabinet [32](#)  
interface configurations  
  dedicated [60](#)  
  shared [60](#)  
  standby [60](#)  
interfaces  
  configuration for takeover [63](#)  
  configuring [66](#)  
  configuring dedicated [51](#)  
  configuring shared [51](#)  
  configuring standby [52](#)  
  dedicated, diagram [63](#)  
  IPv6 considerations [64](#)  
  shared, diagram [62](#)  
  standby, diagram [63](#)  
  types and configurations [64](#)  
internode distance [13](#)  
IP address  
  configuring [66](#)  
  partner IP, specifying [67](#)  
IPv6 considerations [64](#)

## L

licenses  
  cf [53](#)  
  not required [53](#)  
LIF configuration, best practice [23](#)  
limitations  
  configuration, for fabric-attached MetroCluster configurations  
  with disks [135](#)  
local node  
  cabling [145](#)  
long-running processes  
  setting giveback to terminate [94](#)

lun commands, lun online [205](#)

LUNs (array)

- cabling guidelines for MetroCluster configurations with [172](#)
- MetroCluster configuration implementation overview [163](#)
- MetroCluster configuration with disks and array LUNs [195](#)
- systems that support, in MetroCluster configurations [163, 195](#)

LUNs, bringing online [205](#)

## M

mailbox disks in the HA pair [8](#)

manual fencing [204](#)

messages

- active/active configuration status [73](#)
- HA pair status [73](#)

MetroCluster

- managing disk shelves in [102](#)

MetroCluster configuration

- about [114](#)
- enabling unidirectional port mirroring [157, 189](#)
- gathering information about [159](#)
- resetting the default speed of stretch MetroCluster configuration [131](#)
- types [114](#)
- unidirectional port mirroring [158](#)
- verifying during bridge failure [161](#)
- verifying that a switch failure does not cause node failover [160](#)
- verifying the configuration [159](#)

MetroCluster configurations

- advantages of fabric-attached [116](#)
- advantages of stretch [115](#)
- array LUNs
  - cabling guidelines [172](#)
  - connecting devices in [175](#)
  - connecting local systems in [176](#)
  - connecting remote systems in [179](#)
  - connecting the fabric and storage array [183](#)
  - connecting the switch fabric [182](#)
  - FC-VI ports, testing zoning of [191](#)
  - implementation overview [163](#)
  - planning zoning with [173](#)
  - requirements and restrictions [164](#)
  - tasks after connecting devices [191](#)
  - testing setup [192](#)
  - testing zoning of FC-VI ports [191](#)
  - testing, simulating disaster recovery [193](#)
  - zoning [185](#)

changing the default speed of stretch [128, 190](#)

converting to, from standard or mirrored HA pair [142](#)

equipment requirements [123, 140](#)

events that trigger failover in [16](#)

How Data ONTAP works with [114](#)

how mirroring works in [114](#)

LUNs and [205](#)

reestablishing configuration after disaster [206](#)

required documentation [122, 139](#)

MetroCluster configurations (*continued*)

SAS optical cables

- use with disk shelves on Data ONTAP systems [199](#)
- use with disk shelves on V-Series systems [119, 135](#)

setup requirements and restrictions for stretch [118](#)

shared-switches configurations

- cabling with storage arrays [186](#)
- requirements with array LUNs [166](#)

storage arrays

- recommended fabric-attached configuration [197](#)
- recommended stretch MetroCluster configuration [169](#)
- supported stretch configuration [199](#)

systems that support array LUNs in [163, 195](#)

third-party storage

- recommended fabric-attached configuration [166](#)
- See also* array LUNs

V-Series systems

- FibreBridge 6500N bridge with [119, 135](#)
- hardware requirements for disks [119, 135](#)

mirrored aggregates

- re-creating after disaster [209](#)

mirrored HA pairs

- about [14](#)
- advantages of [14](#)
- cabling [41](#)
- cabling Channel A [42](#)
- cabling Channel B [44](#)
- restrictions [25](#)
- setup requirements for [25](#)
- variations [14](#)

mirroring, NVMEM or NVRAM log [8](#)

modules, disk shelf

- about [103](#)
- best practices for changing types [104](#)
- hot-swapping [106](#)
- restrictions for changing types [104](#)
- testing [104](#)

monitoring in normal mode [73](#)

multipath HA

- advantages of [26](#)
- connection types used by [26](#)
- description of [25](#)

multipath HA loop

- adding disk shelves to [102](#)

## N

negotiated failover [84](#)

network interface

- automatic takeover [84](#)
- nfo [84](#)

network interfaces

- configuration for takeover [63](#)
- IPv6 considerations [64](#)
- types and configurations [64](#)

Node A

Node A (*continued*)  
  cabling [145](#)

Node B  
  cabling [149](#)

nodes  
  accessing takeover, with partner command [86](#)  
  command exceptions for emulated [87](#)

nodes in an HA pair  
  reconfiguring for stand-alone operation  
    for nodes using array LUNs [99](#)

Non-HA mode  
  enabling [95](#)

nondisruptive hardware replacement  
  documentation for [108](#)

nondisruptive operations [8](#)

nondisruptive upgrades  
  Data ONTAP [108](#)  
  Data ONTAP, documentation for [108](#)

normal giveback  
  initiating [90](#)

normal mode  
  monitoring in [73](#)

NVMEM log mirroring [8](#)

NVRAM adapter [34](#)

NVRAM adapters  
  MetroCluster configuration requirements [123](#), [140](#)

NVRAM log mirroring [8](#)

**O**

options  
  cf  
    best practices [55](#)

options, matching [54](#)

options, setting [53](#)

**P**

parameters  
  change\_fsid [56](#)  
  required to be identical between nodes [54](#)  
  setting [53](#)

partner command [85](#)

partner commands  
  accessing takeover node with [86](#)

partner IP addresses  
  setting for hardware-assisted takeover [58](#)

partner name, displaying (cf partner) [76](#)

planning worksheet for fabric-attached MetroCluster configurations [138](#)

plexes, requirements for in the HA pair [25](#)

pool assignments, fabric-attached MetroCluster configurations [152](#)

port list  
  creating for mirrored HA pairs [42](#)

ports  
  identifying which ones to use [36](#), [41](#)

power supply best practice [23](#)

preferred primary port  
  removing the settings [157](#), [189](#)

preparing equipment racks [35](#)

primary connections, in multipath HA [26](#)

## R

racking the HA pair  
  in a system cabinet [32](#)  
  in telco-style racks [32](#)

reasons for automatic takeover [79](#)

redundant connections, in multipath HA [26](#)

reestablishing MetroCluster configuration [206](#)

remote node  
  cabling [149](#)

removing  
  disk shelves [108](#)

removing an HA pair [95](#)

requirement, multipath HA [25](#)

requirements  
  documentation [33](#)  
  equipment [34](#)  
  HA pair setup [24](#)  
  hot-swapping a disk shelf module [106](#)  
  MetroCluster configuration documentation [122](#), [139](#)  
  MetroCluster configuration equipment [123](#), [140](#)  
  stretch MetroCluster configuration setup [118](#)  
  tools [33](#), [123](#), [140](#)

restrictions  
  fabric-attached MetroCluster configuration [133](#)  
  HA pair setup [24](#)  
  in mirrored HA pairs [25](#)  
  stretch MetroCluster configuration setup [118](#)

rsh, using to access node after takeover [78](#)

## S

SAS optical cables  
  use with disk shelves on Data ONTAP systems [199](#)  
  use with disk shelves on V-Series systems [119](#), [135](#)

setting options and parameters [53](#)

setup  
  considerations for running on active/active configurations [50](#)

SFP modules  
  MetroCluster configuration requirements [123](#), [140](#)

shared interfaces  
  configuring using ifconfig [65](#)  
  configuring using setup [51](#)  
  described [60](#)  
  diagram [62](#)

shared-switch [153](#), [186](#)

shared-switch configuration  
  requirements [136](#)

shared-switches  
  setting preferred primary port [157](#), [189](#)

shared-switches configuration



- shared-switches configuration (*continued*)
    - cabling FC-VI adapter and ISL [154](#)
  - sharing storage loops or stacks [27](#)
  - shelves
    - hot-removing [108](#)
    - managing in a MetroCluster configuration [102](#)
    - managing in an HA pair [102](#)
  - shorting giveback time [89](#)
  - single point of failure
    - analysis [10](#)
    - definition of [9](#)
    - eliminating [10](#)
  - single-chassis HA configurations
    - diagram of [28](#)
    - interconnect [28](#)
  - SNMP protocol and takeover mode [85](#)
  - software-based disk management [152](#)
  - software-based disk ownership [145](#), [149](#), [150](#)
  - spare disks in the HA pair [8](#), [25](#)
  - stand-alone operation
    - changing an HA pair node to [95](#), [97](#), [99](#)
  - standard HA pair
    - cabling Channel A [37](#)
    - cabling Channel B [38](#)
    - cabling HA interconnect for [40](#), [47](#)
    - cabling HA interconnect for, 32xx systems [40](#), [48](#)
    - variations [27](#)
  - standby connections, in multipath HA [26](#)
  - standby interfaces
    - configuring using setup [52](#)
    - described [60](#)
    - diagram [63](#)
  - status
    - active/active configuration message [73](#)
    - HA pair message [73](#)
    - monitoring HA pair [73](#)
    - of hardware-assisted takeover [74](#)
  - storage arrays
    - recommended fabric-attached MetroCluster configuration [166](#), [197](#)
    - recommended MetroCluster configuration [169](#)
    - recommended stretch MetroCluster configuration [199](#)
  - stretch MetroCluster configuration
    - disk shelf pool assignments
      - fabric-attached MetroCluster configurations [126](#)
    - implementing [118](#)
    - planning [118](#)
    - pool assignments [126](#)
    - resetting default speed [131](#)
  - stretch MetroCluster configurations
    - advantages of [115](#)
    - array LUNs
      - requirements [164](#)
    - cabling [125](#)
    - changing the default speed of [128](#), [190](#)
    - connections required [121](#)
    - stretch MetroCluster configurations (*continued*)
      - illustration of [121](#)
      - on dual-controller systems [126](#)
      - setup requirements and restrictions [118](#)
      - storage arrays
        - recommended configuration [169](#), [199](#)
        - variations [115](#)
    - suggestions
      - how to send feedback about documentation [218](#)
    - switch configuration, for fabric-attached MetroCluster configurations [145](#)
    - switch failure
      - verifying behavior [160](#)
    - switch zoning
      - planning for MetroCluster configuration with array LUNs [173](#)
    - switches
      - requirements in MetroCluster configurations with array LUNs [164](#)
    - SyncMirror
      - requirements for a MetroCluster configuration [164](#)
    - system cabinets
      - installation in [32](#)
      - preparing for cabling [35](#)
    - system configurations
      - HA differences between supported [29](#)
- ## T
- takeover
    - automatic [15](#), [77](#)
    - CIFS sessions and [21](#), [78](#)
    - configuring when it occurs [79](#)
    - configuring with dedicated and hot standby interfaces [63](#)
    - definition of [15](#)
    - determining why one occurred [84](#)
    - disabling [82](#)
    - disabling automatic [83](#)
    - enabling [82](#)
    - enabling automatic [83](#)
    - forcing [81](#)
    - forcing for disaster recovery [204](#)
    - hardware assisted [20](#), [27](#), [57](#)
    - manual [15](#), [77](#)
    - reasons for [79](#)
    - rsh access after [78](#)
    - SNMP settings and [85](#)
    - specifying delay before [83](#)
    - statistics [85](#)
    - Telnet access after [78](#)
    - testing [71](#), [160](#)
    - using /etc/mcrc file at takeover [71](#)
    - what happens after [78](#)
    - what happens during [21](#), [78](#)
    - when it occurs [15](#), [77](#)
  - takeover impossible EMS message [23](#)
  - takeover mode

- takeover mode (*continued*)
  - managing in [84](#)
  - statistics in [85](#)
- takeover nodes
  - accessing with partner command [86](#)
- takeovers
  - setting partner IP address for hardware-assisted [58](#)
- Telnet, using to access node after takeover [78](#)
- testing
  - hardware-assisted takeover [75](#)
  - takeover and giveback [71](#), [160](#)
- testing MetroCluster configuration
  - during switch failure [160](#)
- testing the behavior of MetroCluster configuration
  - during bridge failure [161](#)
- third-party storage
  - MetroCluster configurations with requirements and restrictions [164](#)
- tools, required [33](#), [123](#), [140](#)
- triggers for automatic takeover [79](#)
- Twitter
  - how to receive automatic notification of documentation changes [218](#)

## U

- UDP [157](#), [158](#), [189](#)
- unconfiguring an HA pair [95](#)
- unidirectional port mirroring
  - enabling in a MetroCluster configuration [157](#), [189](#)
  - what is [158](#)
- uninterruptible power supplies, *See* UPSs
- update history for this guide [211](#)
- UPS
  - using with MetroCluster configurations [120](#), [136](#)
- UPSs
  - required connections with HA pairs [48](#)
- utilities
  - checking for common configuration errors with Config Advisor [71](#)
  - downloading and running Config Advisor [71](#)

## V

- V-Series MetroCluster configurations, *See* MetroCluster configurations
- verifying
  - MetroCluster configuration [159](#)
  - takeover and giveback [71](#), [160](#)
- VIF configuration, best practice in an HA configuration [23](#)

## Z

- zoning
  - MetroCluster configurations, array LUNs [185](#)