



Clustered Data ONTAP® 8.3

MetroCluster Installation and Configuration Guide



NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089
U.S.

Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 463-8277
Web: www.netapp.com
Feedback: doccomments@netapp.com

Part number: 215-09156_B0
January 2015

Contents

MetroCluster documentation	6
Preparing for the MetroCluster installation	8
Differences between 7-Mode and clustered Data ONTAP MetroCluster configurations	8
Understanding the parts of the MetroCluster configuration	9
Local HA pair illustration	12
Redundant FC-to-SAS bridges	12
Redundant FC switch fabrics	13
The cluster peering network	14
Considerations for MetroCluster configurations with native disk shelves or array LUNs	14
Required MetroCluster components and naming guidelines	15
Considerations when transitioning from 7-Mode to clustered Data ONTAP	17
Configuration of new MetroCluster systems	17
Hardware setup checklist	18
Software setup checklist	20
Choosing the correct installation procedure for your system	25
Configuring the MetroCluster hardware components	26
Gathering required information and reviewing the workflow	26
FC switch and FC-to-SAS bridge worksheet	26
Hardware installation workflow	28
Installing and cabling MetroCluster components	29
Racking the hardware components	30
Cabling the FC-VI and HBA adapters to the FC switches	31
Cabling the ISLs between MetroCluster sites	36
Cabling the cluster interconnect	37
Cabling the cluster peering connections	37
Cabling the HA interconnect, if necessary	38
Cabling the management and data connections	39
Recommended port assignments for FC switches	39
Installing FC-to-SAS bridges and SAS disk shelves	41
Preparing for the installation	42

Installing the FC-to-SAS bridge and SAS shelves	44
Configuring the FC switches	49
Configuring the FC switches by running a configuration file	50
Configuring the Cisco or Brocade FC switches manually	51
Configuring hardware for sharing a Brocade 6510 FC fabric during transition	96
Reviewing Brocade license requirements	97
Racking the hardware components	97
Cable the new MetroCluster controllers to the existing FC fabrics	98
Configuring switch fabrics sharing between the 7-Mode and clustered MetroCluster configuration	100
Disabling one of the switch fabrics	100
Deleting TI zoning and configuring IOD settings	102
Ensuring ISLs are in the same port group and configuring zoning	104
Reenabling the switch fabric and verify operation	105
Configuring the MetroCluster software in Data ONTAP	107
Gathering required information and reviewing the workflow	107
Worksheet for IP network information for site A	107
Worksheet for IP network information for site B	110
Software configuration workflow	113
Similarities and differences between regular cluster and MetroCluster configurations	113
Verifying disk assignment in Maintenance mode	114
Assigning disk ownership shelf by shelf	116
Verifying the HA state of components is mcc in Maintenance mode	118
Running System Setup to configure the nodes and clusters	119
Configuring the clusters into a MetroCluster configuration	122
Peering the clusters	122
Mirroring the root aggregates	130
Creating a mirrored data aggregate on each node	130
Implementing the MetroCluster configuration	132
Configuring MetroCluster components for health monitoring	134
Checking the MetroCluster configuration	135
Checking for MetroCluster configuration errors with Config Advisor	137
Verifying local HA operation	138
Verifying switchover, healing, and switchback	140

Protecting configuration backup files	140
Planning and installing a MetroCluster configuration with array	
LUNs	141
Planning for a MetroCluster configuration with array LUNs	141
Requirements for a MetroCluster configuration with array LUNs	141
Implementation overview for a MetroCluster configuration with array	
LUNs	143
Connecting devices in a MetroCluster configuration with array LUNs	144
Switch zoning for a MetroCluster configuration with array LUNs	147
Setting up Data ONTAP after connecting devices in a MetroCluster	
configuration with array LUNs	149
Implementing a MetroCluster configuration with both disks and array LUNs	150
Planning a MetroCluster configuration with disks and array LUNs	150
Example of a MetroCluster configuration with disks and array LUNs	151
Using the OnCommand management tools for further configuration	
and monitoring	154
Requirements and limitations when using Data ONTAP in a	
MetroCluster configuration	155
Job schedules in a MetroCluster configuration	155
Cluster peering from the MetroCluster sites to a third cluster	155
Volume creation on a root aggregate	156
Networking and LIF creation guidelines for MetroCluster configurations	156
Volume or FlexClone command VLDB errors	157
Modifying volumes to set NVFAIL in case of switchover	158
Monitoring and protecting database validity by using NVFAIL	158
How NVFAIL protects database files	158
Commands for monitoring data loss events	159
Accessing volumes in NVFAIL state after a switchover	160
Recovering LUNs in NVFAIL states after switchover	160
Glossary of MetroCluster terms	162
Copyright information	165
Trademark information	166
How to send comments about documentation and receive update	
notification	167
Index	168

MetroCluster documentation

There are a number of documents that can help you configure, operate, and monitor a MetroCluster configuration.

MetroCluster and Data ONTAP libraries

Library	Content
NetApp Documentation: MetroCluster in clustered Data ONTAP	<ul style="list-style-type: none">• All MetroCluster guides
NetApp Documentation: Clustered Data ONTAP Express Guides	<ul style="list-style-type: none">• All Data ONTAP express guides
NetApp Documentation: Data ONTAP 8 (current releases)	<ul style="list-style-type: none">• All Data ONTAP guides

MetroCluster and miscellaneous guides

Guide	Content
Clustered Data ONTAP 8.3 MetroCluster Installation Express Guide	<p>How to install a MetroCluster system that has been received from the factory. You should use this guide only if the following is true:</p> <ul style="list-style-type: none">• The MetroCluster configuration has been received from the factory.• The configuration is using Brocade FC storage switches. This guide does not document configuration of the Cisco FC storage switches.• The configuration is not using array LUNs (FlexArray Virtualization).• The configuration is not sharing existing FC fabrics with a 7-Mode fabric MetroCluster during transition.

Guide	Content
<i>Clustered Data ONTAP 8.3 MetroCluster Management and Disaster Recovery Guide</i>	<ul style="list-style-type: none"> • Understanding the MetroCluster configuration • Switchover, healing and switchback • Disaster recovery
<i>MetroCluster Service Guide</i>	<ul style="list-style-type: none"> • Guidelines for maintenance in a MetroCluster configuration • Hardware replacement and firmware upgrade procedures for FC-to-SAS bridges and FC switches • Hot-adding a disk shelf • Hot-removing a disk shelf • Replacing hardware at a disaster site
<i>MetroCluster Tiebreaker Software Installation and Configuration Guide</i>	<ul style="list-style-type: none"> • Monitoring the MetroCluster configuration with the MetroCluster Tiebreaker software
<i>Clustered Data ONTAP 8.3 Data Protection Guide</i>	<ul style="list-style-type: none"> • How mirrored aggregates work • SyncMirror • SnapMirror • SnapVault
<i>NetApp Documentation: OnCommand Unified Manager Core Package (current releases)</i>	<ul style="list-style-type: none"> • Monitoring the MetroCluster configuration
<i>NetApp Documentation: OnCommand Performance Manager for Clustered Data ONTAP</i>	<ul style="list-style-type: none"> • Monitoring MetroCluster performance
<i>7-Mode Transition Tool 2.0 Data and Configuration Transition Guide</i>	<ul style="list-style-type: none"> • Transitioning data from 7-Mode storage systems to clustered storage systems

Preparing for the MetroCluster installation

As you prepare for the MetroCluster installation, you should understand the MetroCluster hardware architecture and required components. If you are familiar with MetroCluster configurations in a 7-mode environment, you should understand the key MetroCluster differences you find in a clustered Data ONTAP environment.

Differences between 7-Mode and clustered Data ONTAP MetroCluster configurations

There are key differences between clustered Data ONTAP MetroCluster configurations and configurations with Data ONTAP operating in 7-Mode.

In clustered Data ONTAP, the MetroCluster configuration includes two HA pairs, each in a separate cluster at physically separated sites.

Feature or component	Clustered Data ONTAP MetroCluster configuration	Data ONTAP 7-Mode MetroCluster configuration
Number of storage controllers	Four The controllers are configured as two HA pairs, one HA pair at each site.	Two The two controllers are configured as a HA pair with one controller at each site.
Local failover available?	Yes A failover can occur at either site without triggering an overall switchover of the configuration.	No In the event of a problem at the local site, the system fails over to the partner site.
Single command for failover or switchover?	Yes The command for local failover is: cf takeover The command for switchover is: metrocluster switchover or metrocluster switchover -forced-on-disaster true	Yes cf takeover or cf forcetakeover -d
DS14 disk shelves supported?	No	Yes

Feature or component	Clustered Data ONTAP MetroCluster configuration	Data ONTAP 7-Mode MetroCluster configuration
Two FC switch fabrics?	Yes	Yes

Related concepts

[Understanding the parts of the MetroCluster configuration](#) on page 9

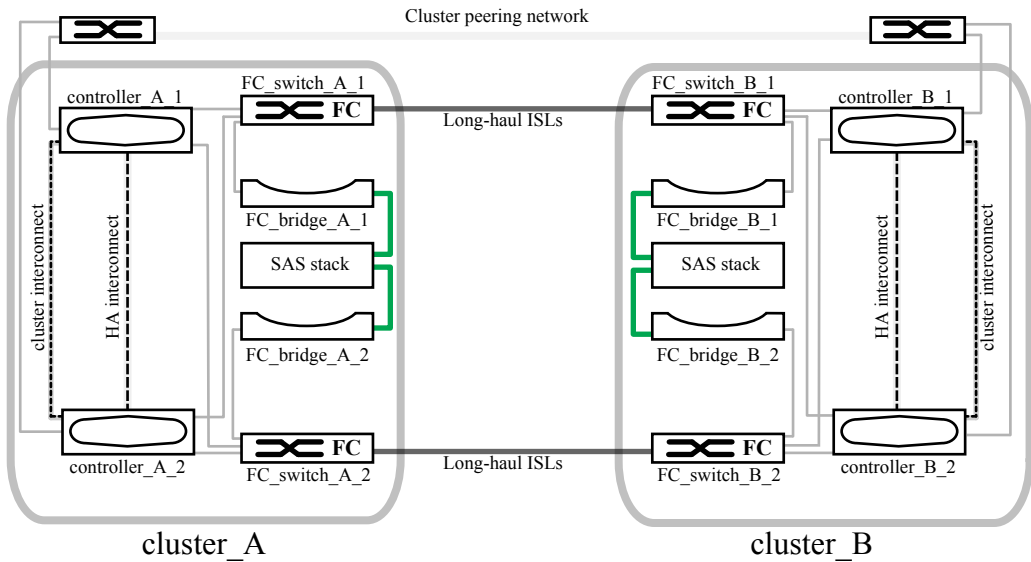
Understanding the parts of the MetroCluster configuration

The MetroCluster configuration consists of a *disaster recovery (DR) group* that includes two HA pairs, each in a separate cluster at physically separated sites. FC switches and long distance Inter-Switch Links (ISLs) provide a backbone connection between the clusters. The clusters are also in a peering relationship, with each cluster's configuration information mirrored to the partner.

The MetroCluster configuration includes the following key hardware elements:


- **Storage controllers**
The storage controllers are not connected directly to the storage but connect to two redundant FC switch fabrics.
- **FC-to-SAS bridges**
The FC-to-SAS bridges connect the SAS storage stacks to the FC switches, providing bridging between the two protocols.
- **FC switches**
The FC switches provide the long haul backbone ISL between the two sites. The switches provide the two storage fabrics that allow data mirroring to the remote storage pools.
- **Cluster peering network**
The cluster peering network provides connectivity for mirroring of Storage Virtual Machine (SVM) configuration information. The configuration of all SVMs on one cluster is mirrored to the partner cluster.

The following illustration shows a simplified view of the MetroCluster configuration. For some connections, a single line represents multiple, redundant connections between the components. Data and management network connections are not shown.



- The configuration consists of two clusters, one at each geographically separated site.
- cluster_A is located at one MetroCluster site.
- cluster_B is located at the second MetroCluster site.
- Each site has one stack of SAS storage.
Additional storage stacks are supported, but only one is shown at each site.
- The HA pairs are configured as switchless clusters, without cluster interconnect switches.
A switched configuration is supported but not shown.

The following illustration shows a more detailed view of the connectivity in a single MetroCluster cluster (both clusters have the same configuration):



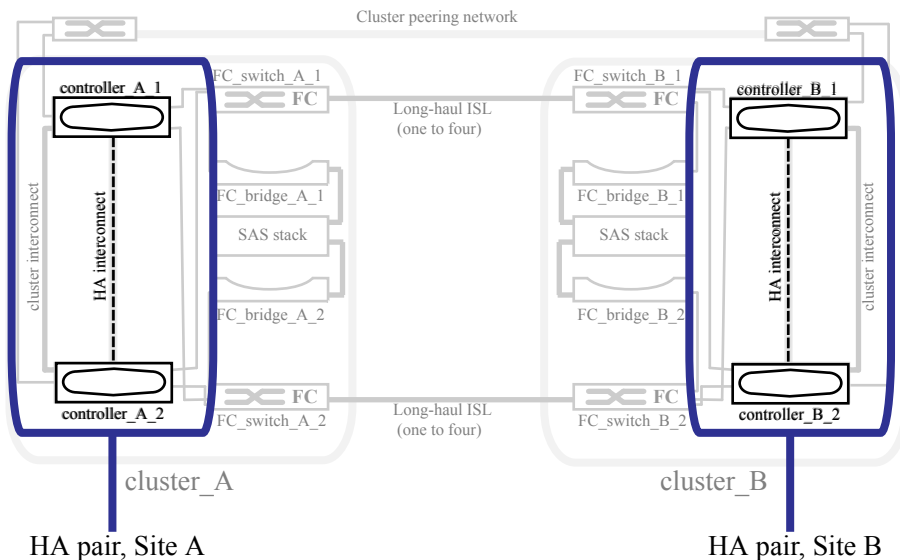
The configuration includes the following connections:

- FC connections from each controller's HBAs and FC-VI adapters to each of the FC switches.
- An FC connection from each FC-to-SAS bridge to an FC switch.
- SAS connections between each SAS shelf and from the top and bottom of each stack to an FC-to-SAS bridge.
- An HA interconnect between each controller in the local HA pair.
If the controllers support a single-chassis HA pair, the HA interconnect is internal, occurring through the backplane, meaning an external interconnect is not required.
- Ethernet connections from the controllers to the customer-provided network used for cluster peering.
SVM configuration is replicated over the cluster peering network.
- A cluster interconnect between each controller in the local HA pair.
If the controllers are configured as a switched cluster, each controller would connect to two cluster interconnect switches.

Local HA pair illustration

Each site consists of two storage controllers configured as an HA pair. This allows local redundancy so that if one storage controller fails, its local HA partner can take over. Such failures can be handled without a MetroCluster switchover operation.

Local HA failover and giveback operations are performed with the `storage failover` commands, in the same manner as a non-MetroCluster configuration.




Related information

[Clustered Data ONTAP 8.3 High-Availability Configuration Guide](#)

Redundant FC-to-SAS bridges


FC-to-SAS bridges provide protocol bridging between SAS attached disks and the FC switch fabric.

Each disk shelf stack requires two FC-to-SAS bridges.



Redundant FC switch fabrics

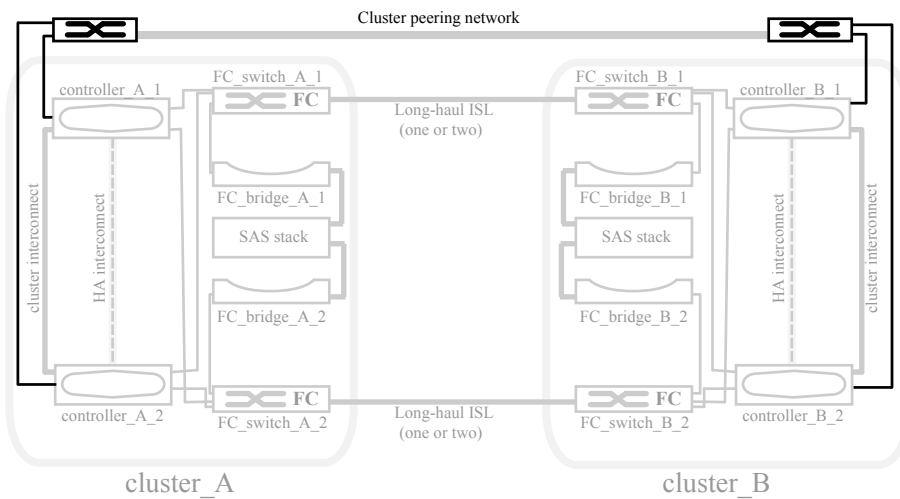
Each switch fabric includes interswitch links (ISLs) that connect the sites. Data is replicated from site-to-site over the ISL.



The cluster peering network

The two clusters in the MetroCluster configuration are peered through a customer-provided cluster peering network. Cluster peering supports the synchronous mirroring of Storage Virtual Machines (SVMs, formerly known as Vservers) between the sites.

Intercluster LIFs must be configured on each node in the MetroCluster configuration, and the clusters must be configured for peering. The ports with the intercluster LIFs are connected to the customer-provided cluster peering network. Replication of SVM configuration is carried out over this network through the Configuration Replication Service.



Related tasks

[Cabling the cluster peering connections](#) on page 37

Considerations for MetroCluster configurations with native disk shelves or array LUNs

The MetroCluster configuration supports installations with native (NetApp) disk shelves only, array LUNs only, or a combination of both.

Related concepts

[Planning and installing a MetroCluster configuration with array LUNs](#) on page 141

Related tasks

[Configuring the MetroCluster hardware components](#) on page 26

Related information

[FlexArray Virtualization Installation Requirements and Reference Guide](#)

Required MetroCluster components and naming guidelines

The MetroCluster configuration requires a variety of hardware components. For convenience and clarity, standard names for components are used throughout these procedures. Also, one site is referred to as Site A and the other site is referred to as Site B. Such terms are relative.

Supported software and hardware

The hardware and software must be supported for the MetroCluster configuration.

[NetApp Interoperability Matrix Tool](#)

[NetApp Hardware Universe](#)

Required components

Because of the hardware redundancy in the MetroCluster configuration, there are two of each component at each site. The sites are arbitrarily assigned the letters A and B and the individual components the numbers 1 and 2.

The MetroCluster configuration also includes SAS storage shelves that connect to the FC-to-SAS bridges.

Component	Example names	
	Site A	Site B
Two Data ONTAP clusters, one at each MetroCluster site. Naming must be unique within the MetroCluster configuration.	cluster_A	cluster_B

Component	Example names	
	Site A	Site B
<p>Four Fibre Channel switches (supported Brocade or Cisco models).</p> <p>The four switches form two switch storage fabrics that provide the ISL between each of the clusters in the MetroClusters.</p> <p>Up to four ISLs per fabric are supported.</p> <p>Naming must be unique within the MetroCluster configuration.</p>	FC_switch_A_1	FC_switch_B_1
	FC_switch_A_2	FC_switch_B_2
<p>Four storage controllers.</p> <p>The two controllers at each site form an HA pair. Each controller has a DR partner at the other site.</p> <p>Naming must be unique within the MetroCluster configuration.</p>	controller_A_1	controller_B_1
	controller_A_2	controller_B_2
<p>Four cluster interconnect switches (if not using two-node switchless clusters).</p> <p>These switches provide cluster communication among the storage controllers in each cluster. The switches are not required if the storage controllers at each site are configured as a two-node switchless cluster.</p> <p>Naming must be unique within the MetroCluster configuration.</p>	clus_sw_A_1	clust_sw_B_1
	clus_sw_A_2	clust_sw_B_2
<p>Two FC-to-SAS bridges (ATTO 6500N FibreBridges) per storage stack.</p> <p>These connect the SAS disk shelves to the FC switch fabric.</p> <p>Naming must be unique within the MetroCluster configuration.</p> <p>The suggested names used as examples in this guide identify the controller that the bridge connects to and the port.</p>	bridge_A_1_port- number	bridge_B_1_po rt-number
	bridge_A_2_port- number	bridge_B_2_po rt-number

Component	Example names	
	Site A	Site B
At least eight SAS disk shelves (recommended). Four shelves are recommended at each site to allow disk ownership on a per-shelf basis. A minimum of two shelves at each site is supported. Note: FlexArray systems support array LUNs and have different storage requirements. Requirements for a MetroCluster configuration with array LUNs on page 141	shelf_A_1_1	shelf_B_1_1
	shelf_A_2_1	shelf_B_2_1
	shelf_B_1_2	shelf_A_1_2
	shelf_B_2_2	shelf_A_2_2

Considerations when transitioning from 7-Mode to clustered Data ONTAP

You must have the new MetroCluster configuration fully configured and operating before you use the transition tools to move data from a 7-Mode MetroCluster configuration to a Data ONTAP configuration. If the 7-Mode configuration uses Brocade 6510 switches, the new configuration can share the existing fabrics to reduce the hardware requirements.

If you have Brocade 6510 switches and plan on sharing the switch fabrics between the 7-Mode fabric MetroCluster and the MetroCluster running in clustered Data ONTAP, you must use the specific procedure for configuring the MetroCluster components.

FMC-MCC transition: Configuring the MetroCluster hardware for sharing a 7-Mode Brocade 6510 FC fabric during transition on page 96

Configuration of new MetroCluster systems

New MetroCluster components are preconfigured and MetroCluster settings are enabled in the software. In most cases, you will not need to perform the detailed procedures provided in this guide.

Hardware racking and cabling

Depending on the configuration you ordered, you might need to rack the systems and complete the cabling.

Configuring the MetroCluster hardware components in systems with native disk shelves on page 26

FC switch and FC-to-SAS bridge configuration

FC-to-SAS bridges received with the new MetroCluster configuration will be configured and will not require additional configuration unless you want to change the names and IP addresses.

In most cases, FC switch fabrics received with the new MetroCluster configuration are preconfigured for two ISLs. If you are using three or four ISLs, you must manually configure the switches.

Configuring the FC switches on page 49

MetroCluster configuration in Data ONTAP

Nodes and clusters received with the new MetroCluster configuration are configured and the MetroCluster is enabled.

Configuring the FC switches on page 49

Hardware setup checklist

You need to know which hardware setup steps were completed at the factory and those you need to complete at each MetroCluster site.

Step	Completed at factory	Completed by you
Mount components in one or more cabinets.	Yes	No
Position cabinets in the desired location.	No	Yes Position them in the original order to ensure that the supplied cables are long enough.
Connect multiple cabinets to each other, if applicable.	No	Yes Use the cabinet interconnect kit if it is included in the order. The kit box is labeled.
Secure the cabinets to the floor, if applicable.	No	Yes Use the universal bolt-down kit if it is included in the order. The kit box is labeled.

Step	Completed at factory	Completed by you
Cable the components within the cabinet.	Yes Cables 5 meters and longer are removed for shipping and placed in the accessories box.	No
Connect the cables between cabinets, if applicable.	No	Yes Cables are in the accessories box.
Connect management cables to the customer's network.	No	Yes Connect them directly or through the CN1601 management switches, if present. Attention: To avoid address conflicts, do not connect management ports to the customer's network until after you change the default IP addresses to the customer's values.
Connect console ports to the customer's terminal server, if applicable.	No	Yes
Connect the customer's data cables to the cluster.	No	Yes
Connect the long-distance ISLs between the MetroCluster sites.	No	Yes <i>Connecting the ISLs between the MetroCluster sites</i> on page 36

Step	Completed at factory	Completed by you
Connect the cabinets to power and power on the components.	No	Yes Power them on in the following order: <ol style="list-style-type: none"> 1. PDUs 2. Disk shelves and FC-to-SAS bridges 3. FC switches 4. Nodes
Assign IP addresses to the management ports of the cluster switches and to the management ports of the management switches if present.	No	Yes, for switched clusters only Connect to the serial console port of each switch and log in with user name “admin” with no password. Suggested addresses are 10.10.10.81, 10.10.10.82, 10.10.10.83, and 10.10.10.84.
Verify cabling by running the Config Advisor tool.	No	Yes <i>Verifying the MetroCluster configuration</i> on page 137

Software setup checklist

You need to know which software setup steps were completed at the factory and those you need to complete at each MetroCluster site.

Step	Completed at factory	Completed by you
Install the clustered Data ONTAP software.	Yes	No

Step	Completed at factory	Completed by you
<p>Create the cluster on the first node at the first MetroCluster site. This includes the following:</p> <ul style="list-style-type: none"> • Name the cluster. • Set the admin password. • Set up the private cluster interconnect. • Install all purchased license keys. • Create the cluster management interface. • Create the node management interface. • Configure the FC switches 	Yes	No
Join the remaining nodes to the cluster.	Yes	No
Enable storage failover on one node of each HA pair and configure cluster high availability.	Yes	No
Enable the switchless-cluster option on a two-node switchless cluster.	Yes	No
Create a test SVM on node cluster-01.	Yes The SVM is configured for NFS and the volume demo is exported.	No
Repeat the steps to configure the second MetroCluster site.	Yes	No
Configure the clusters for peering.	Yes	No
Enable the MetroCluster configuration.	Yes	No

Step	Completed at factory	Completed by you
Configure user credentials and management IP addresses on the management and cluster switches.	Yes, if ordered. User IDs are “admin” with no password.	No
Thoroughly test the MetroCluster configuration.	Yes	No, although you must perform verification steps at your site as described below.
Complete the cluster setup worksheet.	No	Yes Information gathering worksheet for FC switches on page 26
Connect to the cluster management port of node cluster1-01 using SSH and log in.	No	Yes <ol style="list-style-type: none"> 1. Set your laptop to an unused address in the 10.10.10.x subnet (for example, 10.10.10.111) with netmask 255.255.255.0. 2. If the optional management switches are installed, connect your laptop to an open port on one of the management switches. 3. If there are no management switches, connect directly to the following port of node 1 for your controller: e0a for 32xx and 62xx controllers; e0e for FAS8020 controllers; e0i for FAS8040, FAS8060, and FAS8080 controllers. 4. Log in to the cluster as “admin” with password “netapp!123”. The cluster management IP address is 10.10.10.10 with netmask 255.255.255.0.

Step	Completed at factory	Completed by you
Change the password for the admin account to the customer's value.	No	Yes
Configure each node with the customer's values.	No	Yes
Discover the clusters in OnCommand System Manager.	No	Yes <i>Using the OnCommand management tools for further configuration and monitoring on page 154</i>
Configure an NTP server for each cluster.	No	Yes
Verify the cluster peering.	No	Yes
Verify the health of the cluster and that the cluster is in quorum.	No	Yes
Verify basic operation of the MetroCluster sites.	No	Recommended, if the customer allows you to access the network.
Check the MetroCluster configuration.	No	Yes <i>Checking the MetroCluster configuration on page 135</i>
Test storage failover.	No	Yes <i>Verifying local HA operation on page 138</i>
Add the MetroCluster switches and bridges for health monitoring.	No	Yes <i>Configuring MetroCluster components for health monitoring on page 134</i>
Test switchover, healing and switchback.	No	Yes <i>Clustered Data ONTAP 8.3 MetroCluster Management and Disaster Recovery Guide</i>
Set the destination for configuration backup files.	No	Yes <i>Protecting configuration backup files on page 140</i>

Step	Completed at factory	Completed by you
Optional: Change the cluster name if desired, for example, to better distinguish the clusters.	No	Yes Use the Data ONTAP CLI <code>cluster identity modify</code> command.
Optional: Change the node name, if desired.	No	Yes Use the Data ONTAP CLI <code>system node rename</code> command to rename the node.
Configure AutoSupport.	No	Yes

Choosing the correct installation procedure for your system

You need to choose the correct installation procedure based on whether you are using array LUNs, native disks, or sharing existing FC switch fabrics used by a 7-Mode fabric MetroCluster.

For this installation type...	Use these procedures...
Installation with native disks	<ol style="list-style-type: none"> 1. <i>Configuring the MetroCluster hardware components in systems with native disk shelves</i> on page 26 2. <i>Configuring the MetroCluster software in Data ONTAP (native disk shelves only)</i> on page 107
Installation with array LUNs	<i>Planning and installing a MetroCluster configuration with array LUNs</i> on page 141
Installation when sharing with an existing FC switch fabric This is supported only as a temporary configuration with a 7-Mode fabric MetroCluster using Brocade 6510 switches.	<ol style="list-style-type: none"> 1. <i>Configuring the MetroCluster hardware for sharing a 7-Mode Brocade 6510 FC fabric during transition</i> on page 96 2. <i>Configuring the MetroCluster software in Data ONTAP (native disk shelves only)</i> on page 107

Configuring the MetroCluster hardware components

The MetroCluster components must be physically installed, cabled, and configured at both geographic sites of the configuration. The steps are slightly different for a system with native disk shelves as opposed to a system with array LUNs.

Steps

- 1. [Gathering required information and reviewing the workflow](#) on page 26
- 2. [Installing and cabling MetroCluster components](#) on page 29
- 3. [Installing FC-to-SAS bridges and SAS disk shelves](#) on page 41
- 4. [Configuring the FC switches](#) on page 49

Gathering required information and reviewing the workflow

You need to gather the required network information, identify the switch ports you will be using, and review the hardware installation workflow before you begin cabling your system.

Related information

[NetApp Interoperability Matrix Tool](#)

FC switch and FC-to-SAS bridge worksheet

Before beginning to configure the MetroCluster sites, you should gather required configuration information.

Site A, FC switch one (FC_switch_A_1)

Switch configuration parameter	Your value
FC_switch_A_1 IP address	
FC_switch_A_1 Username	
FC_switch_A_1 Password	

Site A, FC switch two (FC_switch_A_2)

Switch configuration parameter	Your value
FC_switch_A_2 IP address	
FC_switch_A_2 Username	
FC_switch_A_2 Password	

Site A, FC-to-SAS bridge 1 (FC_bridge_A_1_port-number)

Each SAS stack requires two FC-to-SAS bridges. One bridge connects to FC_switch_A_1_port-number and the second connects to FC_switch_A_2_port-number.

Site A	Your value
Bridge_A_1_port-number IP address	
Bridge_A_1_port-number Username	
Bridge_A_1_port-number Password	

Site A, FC-to-SAS bridge 2 (FC_bridge_A_2_port-number)

Each SAS stack requires two FC-to-SAS bridges. One bridge connects to FC_switch_A_1_port-number and the second connects to FC_switch_A_2_port-number.

Site A	Your value
Bridge_A_2_port-number IP address	
Bridge_A_2_port-number Username	
Bridge_A_2_port-number Password	

Site B, FC switch one (FC_switch_B_1)

Site B	Your value
FC_switch_B_1 IP address	
FC_switch_B_1 Username	
FC_switch_B_1 Password	

Site B, FC switch two (FC_switch_B_2)

Site B	Your value
FC_switch_B_2 IP address	
FC_switch_B_2 Username	
FC_switch_B_2 Password	

Site B, FC-to-SAS bridge 1 (FC_bridge_B_1_port-number)

Each SAS stack requires two FC-to-SAS bridges. One bridge connects to FC_switch_B_1_port-number and the second connects to FC_switch_B_2_port-number.

Site B	Your value
Bridge_B_1_port-number IP address	
Bridge_B_1_port-number Username	
Bridge_B_1_port-number Password	

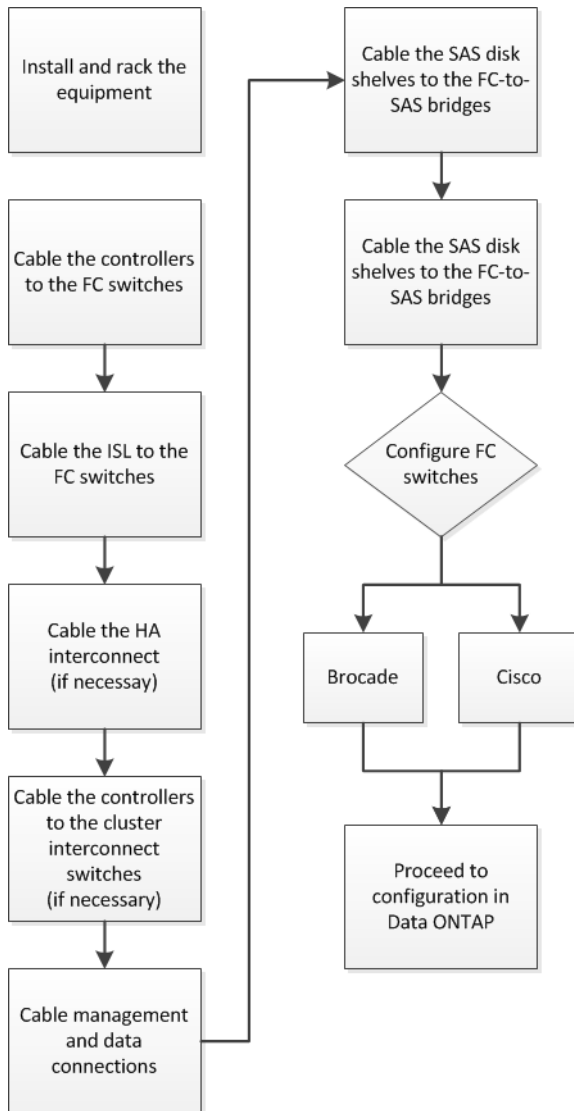
Site B, FC-to-SAS bridge 2 (FC_bridge_B_2_port-number)

Each SAS stack requires two FC-to-SAS bridges. One bridge connects to FC_switch_B_1_port-number and the second connects to FC_switch_B_2_port-number.

Site B	Your value
Bridge_B_2_port-number IP address	
Bridge_B_2_port-number Username	
Bridge_B_2_port-number Password	

Hardware installation workflow

Configuring the MetroCluster components involves some preliminary preparation, setting up the physical components, cabling those components, and configuring the FC-to-SAS bridges and FC switches. This workflow does not include configuring the MetroCluster components in Data ONTAP.



Installing and cabling MetroCluster components

The storage controllers must be cabled to the FC switches and the ISLs must be cabled to link the MetroCluster sites. The storage controllers must also be cabled to the data and management network.

Steps

1. [Racking the hardware components](#) on page 30

2. [Cabling the FC-VI and HBA adapters to the FC switches](#) on page 31
3. [Cabling the ISLs between MetroCluster sites](#) on page 36
4. [Cabling the cluster interconnect](#) on page 37
5. [Cabling the cluster peering connections](#) on page 37
6. [Cabling the HA interconnect, if necessary](#) on page 38
7. [Cabling the management and data connections](#) on page 39
8. [Recommended port assignments for FC switches](#) on page 39

Related tasks

[Connecting devices in a MetroCluster configuration with array LUNs](#) on page 144

Racking the hardware components

If you have not received the equipment already installed in cabinets, you must rack the components.

About this task

This task must be performed on both MetroCluster sites.

Steps

1. Plan out the positioning of the MetroCluster components.

The rack space will depend on the platform model of the storage controllers, the switch types, and the number of disk shelf stacks in your configuration.

2. Properly ground yourself.

3. Install the storage controllers in the rack or cabinet.

[Installation and Setup Instructions FAS8040/FAS8060 Systems](#)

[Installation and setup Instructions FAS80xx Systems with I/O Expansion Modules](#)

[Installation and Setup Instructions FAS8020 systems](#)

[Installation and Setup Instructions 62xx Systems](#)

[Installation and Setup Instructions 32xx Systems](#)

4. Install the FC switches in the rack or cabinet.

5. Install the disk shelves, power them on, and set the shelf IDs.

[SAS Disk Shelves Installation and Service Guide for DS4243, DS2246, DS4486, and DS4246](#)

- You must power-cycle each disk shelf.
- Shelf IDs must be unique for each SAS disk shelf within the entire MetroCluster configuration (including both sites).

6. Install each FC-to-SAS bridge:

- a. Secure the “L” brackets on the front of the bridge to the front of the rack (flush-mount) with the four screws.

The openings in the bridge “L” brackets are compliant with rack standard ETA-310-X for 19-inch (482.6 mm) racks.

For more information and an illustration of the installation, see the *ATTO FibreBridge 6500N Installation and Operation Manual*.

- b. Connect each bridge to a power source that provides a proper ground.
- c. Power on each bridge.

Note: For maximum resiliency, ensure that bridges attached to the same stack of disk shelves are connected to different power sources.

The bridge Ready LED might take up to 30 seconds to illuminate, indicating that the bridge has completed its power-on self test sequence.

Cabling the FC-VI and HBA adapters to the FC switches

The FC-VI adapter and HBAs must be cabled to the site FC switches on each controller module in the MetroCluster configuration.

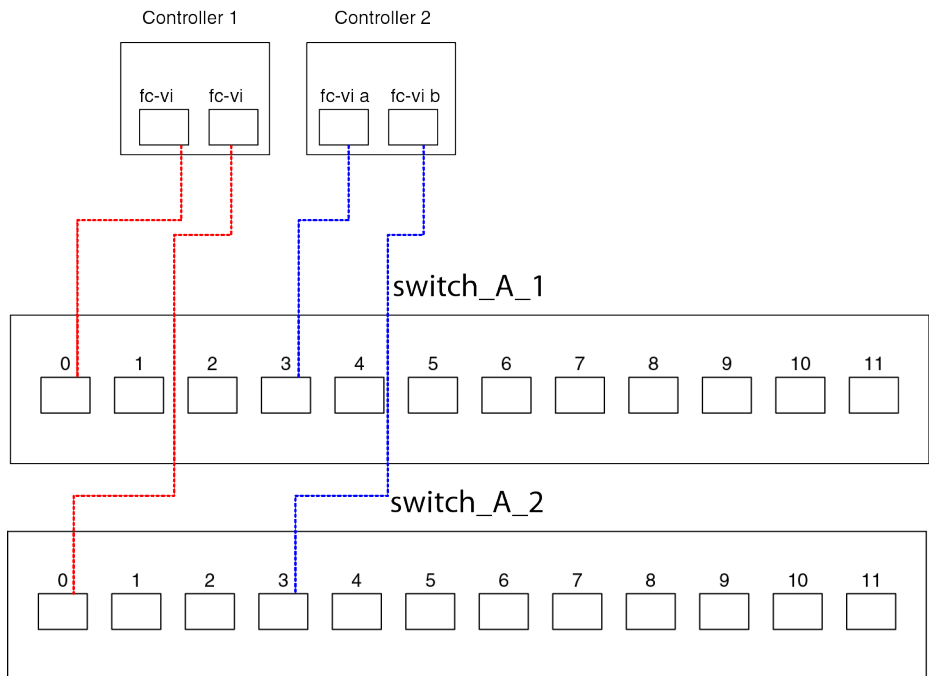
About this task

This task must be performed on both MetroCluster sites.

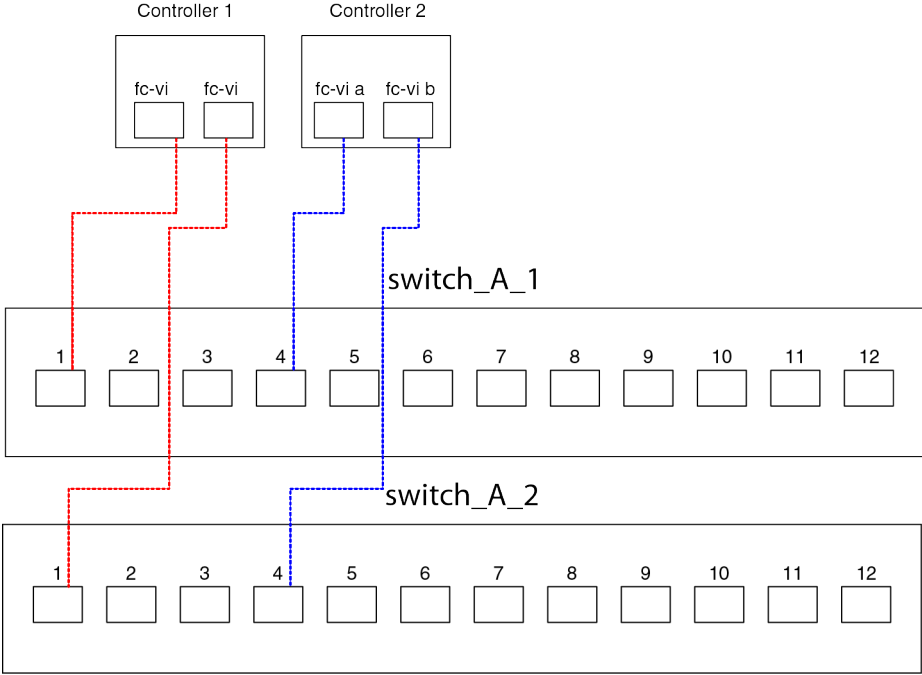
Steps

1. Cable the FC-VI ports.

The following illustration shows Brocade switches:



The following illustration shows Cisco switches:



Note: The Brocade and Cisco switches use different port numbering:

- On Brocade switches, the first port is numbered 0.
- On Cisco switches, the first port is numbered 1.

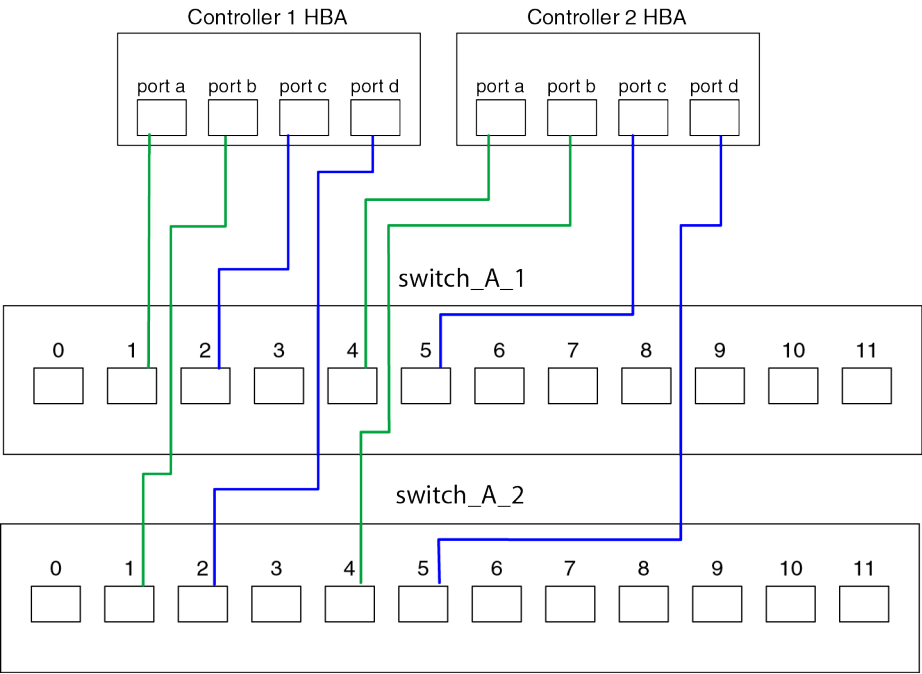
The following tables show both Brocade and Cisco port numbering:

Connect this site component and port...	To this port on FC_switch_x_1...	
	Brocade	Cisco
controller_x_1 FC-VI port a	0	1
controller_x_2 FC-VI port a	3	4

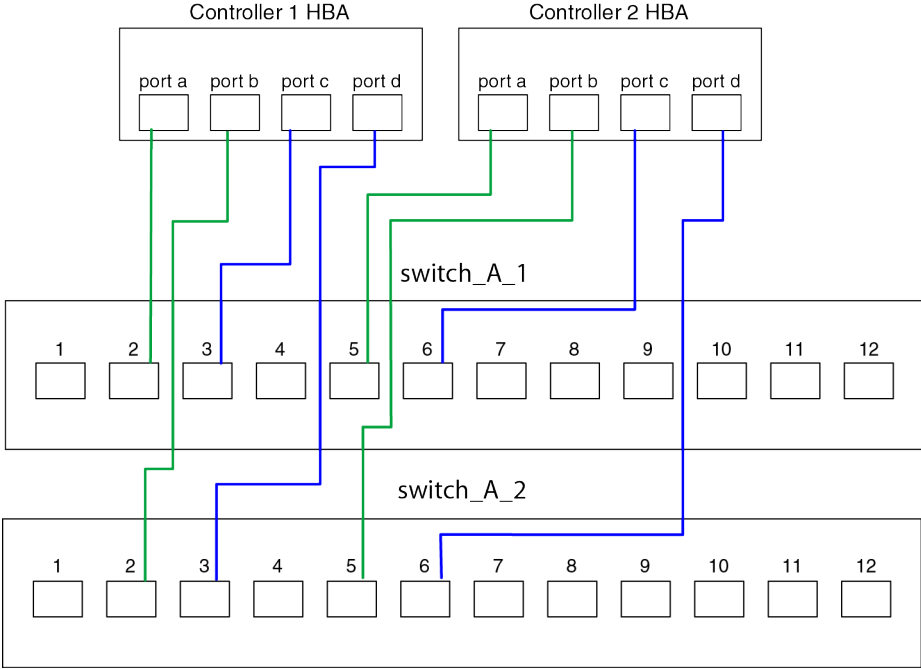
Connect this site component and port...	To this port on FC_switch_x_2...	
	Brocade	Cisco
controller_x_1 FC-VI port b	0	1
controller_x_2 FC-VI port b	3	4

2. Cable the HBA ports.

The following illustration shows Brocade switches:



The following illustration shows Cisco switches:



Note: The Brocade and Cisco switches use different port numbering:

- On Brocade switches, the first port is numbered 0.
- On Cisco switches, the first port is numbered 1.

The following tables show both Brocade and Cisco port numbering:

Connect this site component and port...	To this port on FC_switch_x_1...	
	Brocade	Cisco
controller_x_1 HBA port a	1	2
controller_x_1 HBA port c	2	3
controller_x_2 HBA port a	4	5
controller_x_2 HBA port c	5	6

Connect this site component and port...	To this port on FC_switch_x_2...	
	Brocade	Cisco
controller_x_1 HBA port b	1	2
controller_x_1 HBA port d	2	3

Connect this site component and port...	To this port on FC_switch_x_2...	
	Brocade	Cisco
controller_x_2 HBA port b	4	5
controller_x_2 HBA port d	5	6

Related concepts

Recommended port assignments for FC switches on page 39

Cabling the ISLs between MetroCluster sites

You must connect the FC switches at each site through the fiber-optic Inter-Switch Links (ISLs) to form the switch fabrics that connect the MetroCluster components.

About this task

This task includes steps performed at each MetroCluster site.

Up to four ISLs per FC switch fabric are supported.

Step

1. Connect the FC switches at each site to the ISL or ISLs.

This must be done for both switch fabrics.

Note: The Brocade and Cisco switches use different port numbering:

- On Brocade switches, the first port is numbered 0.
- On Cisco switches, the first port is numbered 1.

The port numbering is shown in the following table:

ISL	Site A			Site B		
	To this port on FC_switch_A_x...			To this port on FC_switch_B_x...		
	Brocade 6505	Brocade 6510	Cisco 9148	Brocade 6505	Brocade 6510	Cisco 9148
First ISL	8	20	36	8	20	36
Second ISL	9	21	40	9	21	40
Third ISL	10	22	44	10	22	44
Four ISL	11	23	48	11	23	48

Related concepts

[Recommended port assignments for FC switches](#) on page 39

Cabling the cluster interconnect

At each site, you must cable the cluster interconnect between the local controllers. If the local controllers are not configured as two-node switchless clusters, two-cluster switches are required.

About this task

This task must be performed at both MetroCluster sites.

Step

1. Cable the cluster interconnect from one controller to the other, or, if cluster interconnect switches are used, from each controller to the switches.

If the storage controller is a...	Do this if using cluster interconnect switches...	Do this if using a switchless cluster...
FAS80xx	Connect ports e0a through e0d on each controller to a local cluster interconnect switch.	Connect ports e0a through e0d on each controller to the same ports on its HA partner.
62xx	Connect ports e0c and e0e on each controller to a local cluster interconnect switch.	Connect ports e0c and e0e on each controller to the same ports on its HA partner.
32xx	Connect ports e1a and e2a on each controller to a local cluster interconnect switch.	Connect ports e1a and e2a on each controller to the same ports on its HA partner.

Related information

[Clustered Data ONTAP 8.3 Network Management Guide](#)

Cabling the cluster peering connections

You must ensure that the controller ports used for cluster peering have connectivity with the cluster on the partner site.

About this task

This task must be performed at both MetroCluster sites.

Step

1. Identify the ports you want to use for cluster peering and ensure they have network connectivity with the partner cluster.

Cluster peering can be done on dedicated ports or on data ports. Using dedicated ports ensures higher throughput for the cluster peering traffic.

Cabling the HA interconnect, if necessary

If the storage controllers in the HA pair are in separate chassis, you must cable the HA interconnect between the controllers.

About this task

This task must be performed at both MetroCluster sites.

The HA interconnect must be cabled only if the storage controllers in the HA pair are in separate chassis. Some storage controller models support two controllers in a single chassis, in which case they use an internal HA interconnect.

Steps

1. Cable the HA interconnect if the storage controller's HA partner is in a separate chassis.

If the storage controller is a...	Do this...
FAS80xx with I/O Expansion Module	<ol style="list-style-type: none">a. Connect port ib0a on the first controller in the HA pair to port ib0a on the other controller.b. Connect port ib0b on the first controller in the HA pair to port ib0b on the other controller.
62xx	<ol style="list-style-type: none">a. Connect port 2a (top port of NVRAM8 card in vertical slot 2) on the first controller in the HA pair to port 2a on the other controller.b. Connect port 2b (bottom port of NVRAM8 card in vertical slot 2) on the first controller in the HA pair to port 2b on the other controller.
32xx	<ol style="list-style-type: none">a. Connect port c0a on the first controller in the HA pair to port c0a on the other controller.b. Connect port c0b on the first controller in the HA pair to port c0b on the other controller.

2. Repeat this task at the MetroCluster partner site.

Related information

[*Installation and Setup Instructions FAS8040/FAS8060 Systems*](#)

[*Installation and setup Instructions FAS80xx Systems with I/O Expansion Modules*](#)

[*Installation and Setup Instructions FAS8020 systems*](#)

[*Installation and Setup Instructions 62xx Systems*](#)

[*Installation and Setup Instructions 32xx Systems*](#)

Cabling the management and data connections

You must cable the management and data ports on each storage controller to the site networks.

About this task

This task must be repeated for each controller at both MetroCluster sites.

Step

1. Cable the controller's management and data ports to the management and data networks at the local site.

[*Installation and Setup Instructions 62xx Systems*](#)

[*Installation and Setup Instructions 32xx Systems*](#)

Recommended port assignments for FC switches

You need to verify that you are using the recommended port assignments when you cable the FC switches.

The Brocade and Cisco switches use different port numbering:

- On Brocade switches, the first port is numbered 0.
- On Cisco switches, the first port is numbered 1.

This difference is reflected in the following table.

Note: The cabling is the same for each FC switch in the switch fabric.

Component and port	FC_switch_x_1			FC_switch_x_2		
	Brocade 6505	Brocade 6510	Cisco 9148	Brocade 6505	Brocade 6510	Cisco 9148
controller_x_1 FC-VI port a	0	0	1	-	-	-
controller_x_1 FC-VI port b	-	-	-	0	0	1

	FC_switch_x_1			FC_switch_x_2		
Component and port	Brocade 6505	Brocade 6510	Cisco 9148	Brocade 6505	Brocade 6510	Cisco 9148
controller_x_1 HBA port a	1	1	2	-	-	-
controller_x_1 HBA port b	-	-	-	1	1	2
controller_x_1 HBA port c	2	2	3	-	-	-
controller_x_1 HBA port d	-	-	-	2	2	3
controller_x_2 FC-VI port a	3	3	4	-	-	-
controller_x_2 FC-VI port b	-	-	-	3	3	4
controller_x_2 HBA port a	4	4	5	-	-	-
controller_x_2 HBA port b	-	-	-	4	4	5
controller_x_2 HBA port c	5	5	6	-	-	-
controller_x_2 HBA port d	-	-	-	5	5	6
bridge_x_1_port-number port 1	6	6	7	6	6	7
bridge_x_1_port-number port 1	7	7	8	7	7	8
bridge_x_1_port-number port 1	12	8	9	12	8	9
bridge_x_1_port-number port 1	13	9	10	13	9	10
ISL port 1	8	20	36	8	20	36
ISL port 2	9	21	40	9	21	40
ISL port 3	10	22	44	10	22	44

	FC_switch_x_1			FC_switch_x_2		
Component and port	Brocade 6505	Brocade 6510	Cisco 9148	Brocade 6505	Brocade 6510	Cisco 9148
ISL port 4	11	23	48	11	23	48

Related tasks

[Cabling the FC-VI and HBA adapters to the FC switches](#) on page 31

[Cabling the ISLs between MetroCluster sites](#) on page 36

Installing FC-to-SAS bridges and SAS disk shelves

You install and cable ATTO FibreBridge FC-to-SAS bridges and SAS disk shelves as part of a new MetroCluster installation.

About this task

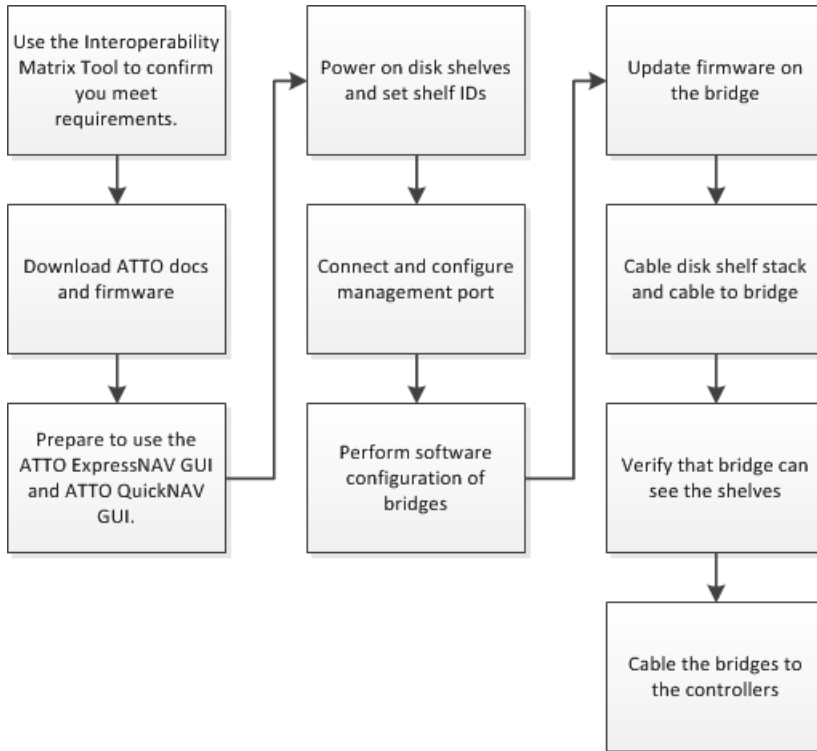
For systems received from the factory, the FC-to-SAS bridges are preconfigured and do not require additional configuration.

This procedure is written with the assumption that you are using the recommended bridge management interfaces: the ATTO ExpressNAV GUI and ATTO QuickNAV utility.

You use the ATTO ExpressNAV GUI to configure and manage a bridge, and to update the bridge firmware. You use the ATTO QuickNAV utility to configure the bridge Ethernet management 1 port.

You can use other management interfaces instead, if needed, such as a serial port or Telnet to configure and manage a bridge and to configure the Ethernet management 1 port, and FTP to update the bridge firmware.

This procedure uses the following workflow:



Steps

1. [Preparing for the installation](#) on page 42
2. [Installing the FC-to-SAS bridge and SAS shelves](#) on page 44

Related concepts

[Example of a MetroCluster configuration with disks and array LUNs](#) on page 151

Preparing for the installation

When you are preparing to install the FC-to-SAS bridges as part of your new MetroCluster system, you must ensure that your system meets certain requirements, including meeting setup and configuration requirements for the bridges. Other requirements include downloading the necessary documents, the ATTO QuickNAV utility, and the bridge firmware.

Before you begin

- Your system must already be installed in a rack if it was not shipped in a system cabinet.

- Your system must be using supported disk shelves and firmware for your version of Data ONTAP.

[*NetApp Interoperability Matrix Tool*](#)

- Your ATTO bridge must be running supported firmware for your version of Data ONTAP.

[*NetApp Interoperability Matrix Tool*](#)

- Your configuration must meet all configuration requirements for the bridge.

[*NetApp Interoperability Matrix Tool*](#)

- Each FC switch must have one FC port available for one bridge to connect to it.
- The computer you are using to set up the bridges must be running an ATTO-supported web browser to use the ATTO ExpressNAV GUI.

The ATTO-supported web browsers are Internet Explorer 8 and 9, and Mozilla Firefox 3.

The *ATTO Product Release Notes* have an up-to-date list of supported web browsers. You can access this document from the ATTO web site as described in the following steps.

Steps

1. Download the following documents:

- [*SAS Disk Shelves Installation and Service Guide for DS4243, DS2246, DS4486, and DS4246*](#)

2. Download content from the ATTO website and from the NetApp website:

- a. From [*NetApp Support*](#), navigate to the **ATTO FibreBridge Description** page by clicking **Software**, scrolling to **Protocol Bridge**, choosing **ATTO FibreBridge** from the drop-down menu, clicking **Go!**, and then clicking **View & Download**.

- b. Access the ATTO web site using the link provided and download the following:

- *ATTO FibreBridge 6500N Installation and Operation Manual*
- ATTO QuickNAV utility (to the computer you are using for setup)

- c. Go to the **ATTO FibreBridge 6500N Firmware Download** page and do the following:

- Navigate to the ATTO FibreBridge 6500N Firmware Download page by clicking **Continue** at the end of the ATTO FibreBridge Description page.
- Download the bridge firmware file using Steps 1 through 3 of that procedure. You update the firmware on each bridge later in this procedure.
- Make a copy of the ATTO FibreBridge 6500N Firmware Download page and release notes for reference when you are instructed to update the firmware on each bridge.

3. Gather the hardware and information needed to use the recommended bridge management interfaces, the ATTO ExpressNAV GUI, and the ATTO QuickNAV utility:

- a. Acquire a standard Ethernet cable (which connects from the bridge Ethernet management 1 port to your network).
- b. Determine a non-default user name and password (for accessing the bridges).
You should change the default user name and password.
- c. Obtain an IP address, subnet mask, and gateway information for the Ethernet management 1 port on each bridge.
- d. Disable VPN clients on the computer you are using for setup.

Active VPN clients cause the QuickNAV scan for bridges to fail.

Installing the FC-to-SAS bridge and SAS shelves

After ensuring that the system meets all the requirements in the *Preparing for the installation* section, you can install your new system.

About this task

- You should use an equal number of disk shelves at each site.
- The system connectivity requirements for maximum distances for disk shelves, FC switches, and backup tape devices using 50-micron, multimode fiber-optic cables, also apply to FibreBridges. The *Site Requirements Guide* has detailed information about system connectivity requirements.

Steps

1. Connect the Ethernet management 1 port on each bridge to your network using an Ethernet cable.

Note: The Ethernet management 1 port enables you to quickly download the bridge firmware (using ATTO ExpressNAV or FTP management interfaces) and to retrieve core files and extract logs.

2. Configure the Ethernet management 1 port for each bridge by following the procedure in the *ATTO FibreBridge 6500N Installation and Operation Manual*, section 2.0.

Note: When running QuickNAV to configure an Ethernet management port, only the Ethernet management port that is connected by the Ethernet cable is configured. For example, if you also wanted to configure the Ethernet management 2 port, you would need to connect the Ethernet cable to port 2 and run QuickNAV.

3. Configure the bridges.

Be sure to make note of the user name and password that you designate.

- a. Configure the bridge name.

The bridges should each have a unique name within the MetroCluster configuration.

- b. Configure the data rate/speed of the bridge FC ports.

The most current information on supported distance can be found in the [NetApp Interoperability Matrix Tool](#).

- c. Configure the connection mode that the bridges use to communicate across the FC network.

You must set the bridge connection mode to *ptp* (point-to-point).

For example, if you were to use the command line interface (CLI) to set the bridge FC 1 port's basic required configuration, you would enter the following commands; the last command saves the configuration changes:

```
set ipaddress
set subnet
set ipgateway
set FCDataRate 1 8Gb
set FCConnMode 1 ptp
set SNMP enabled
set bridgename
SaveConfiguration
```

Note: To set the IP address without the Quicknav utility, you need to have a serial connection to the FibreBridge.

Result: You are prompted to restart the bridge.

The *ATTO FibreBridge 6500N Installation and Operation Manual* has the most current information on available commands and how to use them.

4. Update the firmware on each bridge to the latest version by following the instructions—starting with Step 4—on the FibreBridge 6500N Download page.
5. Cable the disk shelves to the bridges by completing the following substeps:
 - a. Daisy-chain the disk shelves in each stack.


For information about daisy-chaining disk shelves, see the *Installation and Service Guide* for your disk shelf model.

- b. For each stack of disk shelves, cable IOM A square port of the first shelf to SAS port A on FibreBridge A.
 - c. For each stack of disk shelves, cable IOM B circle port of the last shelf to SAS port A on FibreBridge B.

Each bridge has one path to its stack of disk shelves; bridge A connects to the A-side of the stack through the first shelf, and bridge B connects to the B-side of the stack through the last shelf.

Note: The bridge SAS port B is disabled.

The following illustration shows a set of bridges cabled to a stack of three disk shelves:



6. Verify that each bridge can detect all disk drives and disk shelves it is connected to.

If you are using the...	Then...
ATTO ExpressNAV GUI	<ol style="list-style-type: none"> In a supported web browser, enter the IP address of a bridge in the browser box. You are brought to the ATTO FibreBridge 6500N home page, which has a link. Click the link and enter your user name and the password that you designated when you configured the bridge. The ATTO FibreBridge 6500N status page appears with a menu to the left. Click Advanced in the menu. Enter the following command, and then click Submit: sastargets
Serial port connection	Enter the following command: sastargets

Example

The output shows the devices (disks and disk shelves) that the bridge is connected to. Output lines are sequentially numbered so you can quickly count the devices. For example, the following output shows that 10 disks are connected.

Tgt	VendorID	ProductID	Type	SerialNumber
0	NETAPP	X410_S15K6288A15	DISK	3QP1CLE300009940UHJV
1	NETAPP	X410_S15K6288A15	DISK	3QP1ELF600009940V1BV
2	NETAPP	X410_S15K6288A15	DISK	3QP1G3EW00009940U2M0
3	NETAPP	X410_S15K6288A15	DISK	3QP1EWMP00009940U1X5
4	NETAPP	X410_S15K6288A15	DISK	3QP1FZLE00009940G8YU
5	NETAPP	X410_S15K6288A15	DISK	3QP1FZLF00009940TZKZ
6	NETAPP	X410_S15K6288A15	DISK	3QP1CEB400009939MGXL
7	NETAPP	X410_S15K6288A15	DISK	3QP1G7A900009939FNNTT
8	NETAPP	X410_S15K6288A15	DISK	3QP1FY0T00009940G8PA
9	NETAPP	X410_S15K6288A15	DISK	3QP1FXW600009940VERQ

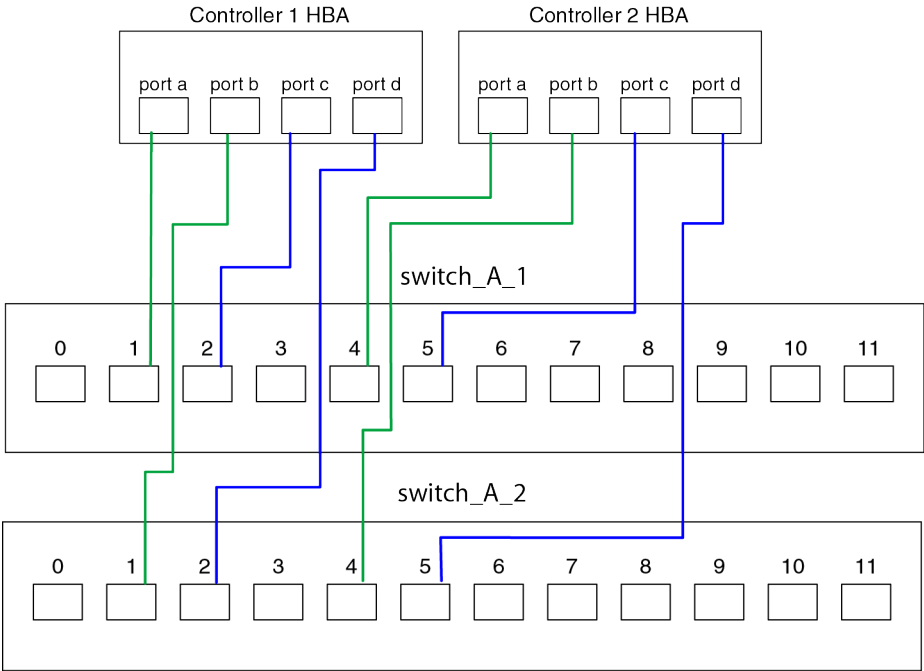
Note: If the text response truncated appears at the beginning of the output, you can use Telnet to connect to the bridge and enter the same command to see all the output.

7. Verify the command output shows the bridge is connected to all disks and disk shelves in the stack that it is supposed to be connected to.

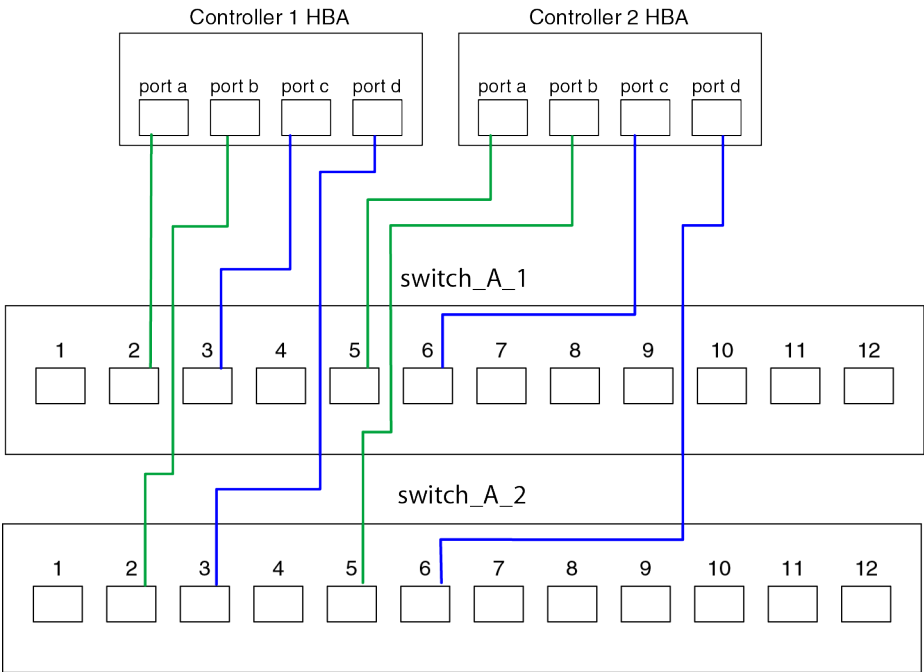
If the output is...	Then...
Correct	Repeat step 6 on page 46 for each remaining bridge.
Not correct	<div>a. Check for loose SAS cables or correct the SAS cabling by repeating step 5 on page 45.</div> <div>b. Repeat step 6 on page 46.</div>

8. Cable each bridge in the first switch fabric to the FC switches:

This illustration shows Brocade switches.



This illustration shows Cisco switches.



Note: The Brocade and Cisco switches use different port numbering:

- On Brocade switches, the first port is numbered “0”
- On Cisco switches, the first port is numbered “1”

This is reflected in the following table:

Cable this bridge...	To this port on FC_switch_x_1...		
	Brocade 6505	Brocade 6510	Cisco
bridge_A_1_port number	6	6	7
bridge_A_1_port number	7	7	8
bridge_A_1_port number	-	8	9
bridge_A_1_port number	-	9	10

9. Repeat the previous step on the bridges at the partner site.

Configuring the FC switches

For systems that were not pre-configured in the factory, you must configure each FC switch in the DR group. This is done manually, or, depending on the switch, can optionally be done with a configuration file.

About this task

For new systems, the FC switch fabrics are typically configured for two ISLs and do not require additional configuration unless you want to change the preconfigured IP addresses.

Choices

- [Configuring the FC switches by running a configuration file](#) on page 50
- [Configuring the Cisco or Brocade FC switches manually](#) on page 51

Related tasks

[Connecting devices in a MetroCluster configuration with array LUNs](#) on page 144

Configuring the FC switches by running a configuration file

If you want to simplify switch configuration, you can download and execute switch configuration files that provide the complete switch settings for certain configurations.

Choices

- [Configuring Brocade FC switches with configuration files](#) on page 50
- [Configuring the Cisco FC switches with configuration files](#) on page 50

Configuring Brocade FC switches with configuration files

When you configure a Brocade FC switch, you can download and execute switch configuration files that provide the complete switch settings for certain configurations.

Before you begin

You must have access to an FTP server. The switches must have connectivity with the FTP server.

About this task

Each configuration file is different and must be used with the correct switch. Only one of the configuration files for each switch fabric contains zoning commands.

Steps

1. Go to the software download page.
[NetApp Downloads: Software](#)
2. In the list of products, find the row for Fibre Channel Switch, and in the drop-down list select **Brocade**.
3. On the **Fibre Channel Switch for Brocade** page, click **View & Download**.
4. On the **Fibre Channel Switch - Brocade** page, click the MetroCluster link.
5. Follow the directions on the **MetroCluster Configuration Files for Brocade Switches** description page to download and run the files.

Configuring the Cisco FC switches with configuration files

To configure a Cisco FC switch, you can download and execute switch configuration files that provide the complete switch settings for certain configurations.

Steps

1. Go to the software download page.
[NetApp Downloads: Software](#)

2. In the list of products, find the row for **Fibre Channel Switch**, and then select Cisco from the drop-down list.
3. On the **Fibre Channel Switch for Cisco** page, click the **View & Download** button.
4. On the **Fibre Channel Switch - Cisco** page, click the **MetroCluster** link.
5. Follow the directions on the **MetroCluster Configuration Files for Cisco Switches Description** page to download and run the files.

Configuring the Cisco or Brocade FC switches manually

Switch configuration procedures and commands are different, depending on the switch vendor.

Choices

- [Configuring the Brocade FC switches](#) on page 51
- [Configuring the Cisco FC switches](#) on page 73

Configuring the Brocade FC switches

You must configure each of the Brocade switch fabrics in the MetroCluster configuration.

Before you begin

- You must have a PC or UNIX workstation with Telnet or SSH access to the FC switches.
- You must be using four supported Brocade switches of the same model with the same Brocade Fabric Operating System (FOS) version and licensing.
[NetApp Interoperability Matrix Tool](#)
- You must have four switches; the MetroCluster configuration requires four switches. The four switches must be connected to two fabrics of two switches each, with each fabric spanning both sites.
- Two initiator ports must be connected from each storage controller to each fabric. Each storage controller must have four initiator ports available to connect to the switch fabrics.

About this task

- ISL trunking is supported.
- ISLs should have the same length and same speed ISLs in one fabric. Different lengths and speeds can be used in the different fabrics.
- Metro-E and TDM (SONET/SDH) are not supported. Any non-FC native framing or signaling is not supported.
Metro-E means Ethernet framing/signaling either natively over a Metro distance or through some TDM, MPLS or WDM.

- TDMs, FCR (native FC Routing) or FCIP extensions are not supported for the MetroCluster FC switch fabric.
- Third-party encryption devices are not supported on any link in the MetroCluster FC switch fabric, including the ISL links across the WAN.
- Certain switches in the MetroCluster FC switch fabric support encryption or compression, and sometimes support both.
[NetApp Interoperability Matrix Tool](#)
- The Brocade Virtual Fabric (VF) feature is not supported.

Steps

1. [Reviewing Brocade license requirements](#) on page 52
2. [Setting the Brocade FC switch values to factory defaults](#) on page 53
3. [Configuring the basic switch settings](#) on page 56
4. [Configuring the E-ports on a Brocade FC switch](#) on page 58
5. [Configuring the non-E-ports on the Brocade switch](#) on page 64
6. [Configuring zoning on Brocade FC switches](#) on page 65
7. [Setting ISL encryption on Brocade 6510 switches](#) on page 69

Related information

[NetApp Interoperability Matrix Tool](#)

Reviewing Brocade license requirements

You need certain licenses for the switches in a MetroCluster configuration. You must install these licenses on all four switches.

The MetroCluster configuration has the following Brocade license requirements:

- Trunking license for systems using more than one ISL, as recommended.
- Extended Fabric license (for ISL distances over 6 km)
- Enterprise license for sites with more than one ISL and an ISL distance greater than 6 km
The Enterprise license includes Brocade Network Advisor and all licenses except for additional port licenses.

You can verify that the licenses are installed by using the `licenseshow` command. If you do not have these licenses, contact your sales representative before proceeding.

Setting the Brocade FC switch values to factory defaults

You must set the switch to its factory defaults to ensure a successful configuration. You must also assign each switch a unique name.

About this task

In the examples in this procedure, the fabric consists of BrocadeSwitchA and BrocadeSwitchB.

Steps

1. Make a console connection and log in to both switches in one fabric.
2. Disable the switch persistently:

switchcfgpersistentdisable

This ensures the switch will remain disabled after a reboot or fastboot. If this command is not available, use the `switchdisable` command.

Example

The following example shows the command on BrocadeSwitchA:

```
BrocadeSwitchA:admin> switchcfgpersistentdisable
```

The following example shows the command on BrocadeSwitchB:

```
BrocadeSwitchA:admin> switchcfgpersistentdisable
```

3. Enter **switchname** *switch_name* to set the switch name.

The switches should each have a unique name. After setting the name, the prompt changes accordingly.

Example

The following example shows the command on BrocadeSwitchA:

```
BrocadeSwitchA:admin> switchname "FC_switch_A_1"
FC_switch_A_1:admin>
```

The following example shows the command on BrocadeSwitchB:

```
BrocadeSwitchB:admin> switchname "FC_Switch_B_1"
FC_switch_B_1:admin>
```

4. Set all ports to their default values by issuing the following command for each port:

portcfgdefault

This must be done for all ports on the switch.

Example

The following example shows the commands on FC_switch_A_1:

```
FC_switch_A_1:admin> portcfgdefault 0
FC_switch_A_1:admin> portcfgdefault 1
...
FC_switch_A_1:admin> portcfgdefault 39
```

The following example shows the commands on FC_switch_B_1:

```
FC_switch_B_1:admin> portcfgdefault 0
FC_switch_B_1:admin> portcfgdefault 1
...
FC_switch_B_1:admin> portcfgdefault 39
```

5. Clear the zoning information by issuing the following commands:

cfgdisable

cfgclear

cfgsave

Example

The following example shows the commands on FC_switch_A_1:

```
FC_switch_A_1:admin> cfgdisable
FC_switch_A_1:admin> cfgclear
FC_switch_A_1:admin> cfgsave
```

The following example shows the commands on FC_switch_B_1:

```
FC_switch_B_1:admin> cfgdisable
FC_switch_B_1:admin> cfgclear
FC_switch_B_1:admin> cfgsave
```

6. Set the general switch settings to default:

configdefault

Example

The following example shows the command on FC_switch_A_1:

```
FC_switch_A_1:admin> configdefault
```

The following example shows the command on FC_switch_B_1:

```
FC_switch_B_1:admin> configdefault
```

7. Set all ports to non-trunking mode:

switchcfgtrunk 0

Example

The following example shows the command on FC_switch_A_1:

```
FC_switch_A_1:admin> switchcfgtrunk 0
```

The following example shows the command on FC_switch_B_1:

```
FC_switch_B_1:admin> switchcfgtrunk 0
```

8. On Brocade 6510 switches, disable the Brocade Virtual Fabrics (VF) feature:

fosconfig options

Example

The following example shows the command on FC_switch_A_1:

```
FC_switch_A_1:admin> fosconfig --disable vf
```

The following example shows the command on FC_switch_B_1:

```
FC_switch_B_1:admin> fosconfig --disable vf
```

9. Clear the Administrative Domain (AD) configuration:

ad options

Example

The following example shows the commands on FC_switch_A_1:

```
FC_switch_A_1:admin> switch:admin> ad --select AD0
FC_switch_A_1:> defzone --noaccess
FC_switch_A_1:> cfgsave
FC_switch_A_1:> exit
```

```
FC_switch_A_1:admin> ad --clear -f
FC_switch_A_1:admin> ad --apply
FC_switch_A_1:admin> ad --save
FC_switch_A_1:admin> exit
```

The following example shows the commands on FC_switch_B_1:

```
FC_switch_B_1:admin> switch:admin> ad --select AD0
FC_switch_A_1:> defzone --noaccess
FC_switch_A_1:> cfgsave
FC_switch_A_1:> exit
FC_switch_B_1:admin> ad --clear -f
FC_switch_B_1:admin> ad --apply
FC_switch_B_1:admin> ad --save
FC_switch_B_1:admin> exit
```

10. Reboot the switch by issuing the following command:

reboot

Example

The following example shows the command on FC_switch_A_1:

```
FC_switch_A_1:admin> reboot
```

The following example shows the command on FC_switch_B_1:

```
FC_switch_B_1:admin> reboot
```

Configuring the basic switch settings

You must configure basic global settings, including the domain ID, for the switches.

About this task

This task contains steps that must be performed on each switch at both MetroCluster sites.

In this procedure, you set the unique domain ID for each switch as shown in the following examples:

- FC_switch_A_1 is assigned to domain ID 5
- FC_switch_A_2 is assigned to domain ID 6
- FC_switch_B_1 is assigned to domain ID 7
- FC_switch_B_2 is assigned to domain ID 8

Using that example, domain IDs 5 and 7 form fabric_1 and domain IDs 6 and 8 form fabric_2.

Steps

1. Issue the following command to enter configuration mode:
configure
2. As you proceed through the prompts, set the domain ID for the switch, press ENTER in response to the prompts until you get to RSCN Transmission Mode, set that value to y, and press ENTER until you return to the switch prompt.

Example

```
FC_switch_A_1:admin> configure
Fabric parameters = y
Domain_id = 5
.
.
RSCN Transmission Mode (yes, y, no, n): [no] y
```

3. If you are using two or more ISLs per fabric, configure in-order-delivery of frames:
These steps must be performed on each switch fabric.

- a. Enable in-order-delivery:

```
iodset
```

- b. Set the Advanced Performance Tuning (APT) policy to 1:

```
aptpolicy 1
```

- c. Disable Dynamic Load Sharing (DLS):

```
dlsreset
```

- d. Verify the IOD settings using the **iodshow**, **aptpolicy** and **dlsshow** commands.

Example

For example, issue the following commands on FC_switch_A_1:

```
FC_switch_A_1:admin> iodshow
IOD is set

FC_switch_A_1:admin> aptpolicy
Current Policy: 1 0(ap)

3 0(ap) : Default Policy
1: Port Based Routing Policy
3: Exchange Based Routing Policy
    0: AP Shared Link Policy
    1: AP Dedicated Link Policy
```

```
command aptpolicy completed

FC_switch_A_1:admin> dlsshow
DLS is not set
```

e. Repeat these steps on the second switch fabric.

4. Reboot the switch:

reboot

Example

On FC_switch_A_1:

```
FC_switch_A_1:admin> reboot
```

On FC_switch_B_1:

```
FC_switch_B_1:admin> reboot
```

5. Persistently enable the switch:

switchcfgpersistenable

Example

On FC_switch_A_1:

```
FC_switch_A_1:admin> switchcfgpersistenable
```

On FC_switch_B_1:

```
FC_switch_B_1:admin> switchcfgpersistenable
```

Configuring the E-ports on a Brocade FC switch

On each switch fabric, you must configure the switch ports that connect the Inter-Switch Link (ISL). These ISL ports are otherwise known as the E-ports.

Before you begin

Review the following guidelines before configuring the E-ports:

- All ISLs in an FC switch fabric must be configured with the same speed and distance.
- The supported speeds are 4 Gbps, 8 Gbps, and 16 Gbps.
The combination of the switch port and SFP must support the speed.

- The distance supported can be as far as 200 km, depending on the FC switch model.
NetApp Interoperability Matrix Tool
- The ISL link must have a dedicated lambda and the link must be supported by Brocade for the distance, switch type and FOS.
- You must not use the L0 setting when issuing the `portCfgLongDistance` command. Instead, you should use the LE or LS setting to configure the distance on the Brocade switches with a minimum of LE.
- You must not use the LD setting when issuing the `portCfgLongDistance` command when working with xWDM/TDM equipment. Instead, you should use the LE or LS setting to configure the distance on the Brocade switches.

About this task

This task must be performed for each FC switch fabric.

The following tables show the ISL ports for the different switches and different numbers of ISLs:

Ports for dual ISL configurations (all switches)	
FC_switch_x_1 ISL ports	FC_switch_B_1 ports
10	10
11	11

Ports for three or four ISL configurations (Brocade 6505)	
FC_switch_A_1 ISL ports	FC_switch_B_1 ports
8	8
9	9
10	10
11	11

Ports for three or four ISL configurations (Brocade 6510)	
FC_switch_A_1 ISL ports	FC_switch_B_1 ports
20	20
21	21
22	22
23	23

Steps

1. Configure the port speed:

```
portcfgspeed port octet_combo
```

You must use the highest common speed supported by the components in the path.

Example

In the following example, there is one ISL for each fabric:

```
FC_switch_A_1:admin> portcfgspeed 10 16
FC_switch_B_1:admin> portcfgspeed 10 16
```

In the following example, there are two ISLs for each fabric:

```
FC_switch_A_1:admin> portcfgspeed 10 16
FC_switch_A_1:admin> portcfgspeed 11 16

FC_switch_B_1:admin> portcfgspeed 10 16
FC_switch_B_1:admin> portcfgspeed 11 16
```

2. If more than one ISL for each fabric is used, enable trunking by issuing the following command for each ISL port:

```
portcfgtrunkport port-number 1
```

Example

```
FC_switch_A_1:admin> portcfgtrunkport 10 1
FC_switch_A_1:admin> portcfgtrunkport 11 1

FC_switch_B_1:admin> portcfgtrunkport 10 1
FC_switch_B_1:admin> portcfgtrunkport 11 1
```

3. Enable QoS traffic by issuing the following command for each of the ISL ports:

```
portcfgqos --enable port-number
```

Example

In the following example, there is one ISL per switch fabric:

```
FC_switch_A_1:admin> portcfgqos --enable 10
FC_switch_B_1:admin> portcfgqos --enable 10
```

Example

In the following example, there are two ISLs per switch fabric:

```
FC_switch_A_1:admin> portcfgqos --enable 10
FC_switch_A_1:admin> portcfgqos --enable 11

FC_switch_B_1:admin> portcfgqos --enable 10
FC_switch_B_1:admin> portcfgqos --enable 11
```

4. Verify the settings using the portCfgShow command.

Example

The following example shows the output for a configuration that uses two ISLs cabled to port 10 and port 11:

Ports of Slot	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Speed	AN	AN	AN	AN	AN	AN	8G	AN	AN	AN	16G	16G	AN	AN	AN	AN
Fill Word	0	0	0	0	0	0	3	0	0	0	3	3	3	0	0	0
AL_PA Offset 13
Trunk Port	ON	ON
Long Distance
VC Link Init
Locked L_Port
Locked G_Port	ON
Disabled E_Port
Locked E_Port
ISL R_RDY Mode
RSCN Suppressed
Persistent Disable
LOS TOV enable
NPV capability	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON
NPV PP Limit	126	126	126	126	126	126	126	126	126	126	126	126	126	126	126	126
QOS E_Port	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE
Mirror Port
Rate Limit
Credit Recovery	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON
Fport Buffers
Port Auto Disable
CSCTL mode
Fault Delay	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

5. Calculate the ISL distance.

Due to the behavior of FC-VI, the distance must be set to 1.5 times the real distance with a minimum of 10 (LE).

The distance for the ISL is calculated as follows, rounded up to the next full kilometer:

$$1.5 \times \text{real_distance} = \text{distance}$$

Example

The distance is 3 km, then $1.5 \times 3\text{km} = 4.5$. This is lower than 10, so the ISL must be set to LE.

Example

The distance is 20 km, then $1.5 \times 20\text{km} = 30$. The ISL must be set to LS 30.

6. Set the distance on each ISL port:

portcfglongdistance

port level vc_link_init distance

A *vc_link_init* value of 1 uses the ARB fill word (default). A value of 0 uses IDLE. The required value may depend on the link being used. The commands must be repeated for each ISL port.

Example

Because the distance is assumed to be 20 km, the setting is 30 with the default *vc_link_init* value of 1:

```
FC_switch_A_1:admin> portcfglongdistance 10 LS 1 30
FC_switch_B_1:admin> portcfglongdistance 10 LS 1 30
```

7. Verify the distance setting:

portbuffershow

A distance setting of LE appears as 10km.

Example

The following example shows output is a configuration that uses ISLs on port 10 and port 11:

```
FC_switch_A_1:admin> portbuffershow
```

User Port	Port Type	Lx Mode	Max/Resv Buffers	Buffer Usage	Needed Buffers	Link Distance	Remaining Buffers
----	----	----	-----	-----	-----	-----	-----
...							
10	E	-	8	67	67	30km	
11	E	-	8	67	67	30km	
...							
23		-	8	0	-	-	466

8. Verify that both switches form one fabric:

switchshow

Example

The following example shows the output for a configuration that uses ISLs on port 10 and port 11:

```

FC_switch_A_1:admin> switchshow
switchName: FC_switch_A_1
switchType: 71.2
switchState:Online
switchMode: Native
switchRole: Subordinate
switchDomain:      5
switchId:   fffc01
switchWwn:  10:00:00:05:33:86:89:cb
zoning:      OFF
switchBeacon: OFF

Index Port Address Media Speed State Proto
=====
...
10  10  010C00  id    16G  Online FC  E-Port  10:00:00:05:33:8c:2e:9a "FC_switch_B_1"
11  11  010B00  id    16G  Online FC  E-Port  10:00:00:05:33:8c:2e:9a "FC_switch_B_1"
(upstream)
...

FC_switch_B_1:admin> switchshow
switchName: FC_switch_B_1
switchType: 71.2
switchState:Online
switchMode: Native
switchRole: Principal
switchDomain:      7
switchId:   fffc03
switchWwn:  10:00:00:05:33:8c:2e:9a
zoning:      OFF
switchBeacon: OFF

Index Port Address Media Speed State Proto
=====
...
10  10  030A00  id    16G  Online FC  E-Port  10:00:00:05:33:86:89:cb "FC_switch_A_1"
11  11  030B00  id    16G  Online FC  E-Port  10:00:00:05:33:86:89:cb "FC_switch_A_1"
(downstream)
...

```

9. Confirm the configuration of the fabrics:

fabricshow

Example

```

FC_switch_A_1:admin> fabricshow
Switch ID      Worldwide Name      Enet IP Addr FC IP Addr Name
-----
1: fffc01 10:00:00:05:33:86:89:cb 10.10.10.55 0.0.0.0 "FC_switch_A_1"
3: fffc03 10:00:00:05:33:8c:2e:9a 10.10.10.65 0.0.0.0 >"FC_switch_B_1"

```

```

FC_switch_B_1:admin> fabricshow
Switch ID      Worldwide Name      Enet IP Addr FC IP Addr Name
-----
1: fffc01 10:00:00:05:33:86:89:cb 10.10.10.55 0.0.0.0 "FC_switch_A_1"
3: fffc03 10:00:00:05:33:8c:2e:9a 10.10.10.65 0.0.0.0 >"FC_switch_B_1"

```

10. Repeat the previous steps for the second FC switch fabric.

Related concepts

[Recommended port assignments for FC switches](#) on page 39

Configuring the non-E-ports on the Brocade switch

You must configure the non-E-ports on the FC switch. In the MetroCluster configuration, these are the ports that connect the switch to the HBA initiators, FC-VI interconnects, and FC-to-SAS bridges. These steps must be done for each port.

About this task

In the following example, the ports connect an FC-to-SAS bridge:

- Port 6 on FC_FC_switch_A_1 at Site_A
- Port 6 on FC_FC_switch_B_1 at Site_B

Steps

1. Configure the port speed for each non-E-port:

`portcfgspeed port speed`

You should use the highest common speed, which is the highest speed supported by all components in the data path: the SFP, the switch port that the SFP is installed on, and the connected device (HBA, bridge, etc).

For example, the components might have the following supported speeds:

- The SFP is capable of 1/2/4 GB.
- The switch port is capable of 1/2/4/8 GB.
- The connected HBA maximum speed is 4 GB.

The highest common speed in this case is 4 GB, so the port should be configured for a speed of 4 GB.

Example

```
FC_switch_A_1:admin> portcfgspeed 6 4
FC_switch_B_1:admin> portcfgspeed 6 4
```

2. Verify the settings using the `portcfgshow` command:

Example

```
FC_switch_A_1:admin> portcfgshow
FC_switch_B_1:admin> portcfgshow
```


In the example output, port 6 has the following settings:

- Speed is set to 4G

Ports of Slot	0	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Speed	AN	AN	AN	AN	AN	AN	AN	4G	AN	AN	AN	AN	AN	AN	AN	AN	AN
Fill Word	0	0	0	0	0	0	0	3	0	0	0	0	0	0	0	0	0
AL_PA Offset 13
Trunk Port
Long Distance
VC Link Init
Locked L_Port
Locked G_Port
Disabled E_Port
Locked E_Port
ISL R_RDY Mode
RSCN Suppressed
Persistent Disable
LOS TOV enable
NPIV capability	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON
NPIV PP Limit	126	126	126	126	126	126	126	126	126	126	126	126	126	126	126	126	126
QOS E_Port	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE	AE
Mirror Port
Rate Limit
Credit Recovery	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON
Fport Buffers
Port Auto Disable
CSCTL mode
Fault Delay	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Configuring zoning on Brocade FC switches

You must assign the switch ports to separate zones to separate controller and storage traffic, one zone for the FC-VI ports and one for the storage ports.

About this task

For the FC-VI ports, a Quality of Service (QoS) zone is required. A QoS zone has a special name to differentiate it from a regular zone. The FC-VI zone must have the name start with the prefix **QOSH1_**, followed by a user-defined portion of the name.

Zone type	Contains	Priority	Example names
FC-VI	Four ports, one for each FC-VI cable from each controller	Always high priority	QOSH1_FCVI_1

Zone type	Contains	Priority	Example names
Storage	Nine ports: <ul style="list-style-type: none"> • Eight HBA initiator ports (two HBA connections for each controller). • One port connecting to an FC-to-SAS bridge. 	Does not apply Storage ports use standard zoning.	STOR_A_1_6

Note: Zoning for the fabric can be configured from one switch in the fabric. In this example, it is configured on Switch_A_1.

The examples in the following steps use these ports and zones:

Zone	Switch Domain	Port	Port usage
QOSH1_FCVI_1	5	0	FC-VI
		3	
	7	0	
		3	
STOR_A_1_6	5	1	HBA
		2	
		4	
		5	
		6	ATTO
STOR_B_1_6	7	1	HBA
		2	
		4	
		5	
		6	ATTO

Steps

1. Create the FC-VI QoS zone by using the `zonecreate` command:

```
zonecreate "QOSH1_FCVI_1", member;member ...
```

Example

In this example, ports 0 and 3 of domain 1 (Switch_A_1) and ports 0 and 4 of domain 3 (Switch_B_1) are members of the FC-VI zone.

```
Switch_A_1:admin> zonecreate "QOSH1_FCVI_1", "5,0; 5,3; 7,0; 7,3"
```

2. Create the storage zone for switch domains:

```
zonecreate name, member;member ...
```

Example

In the following example, ports 1, 2, 4 and 5 of domain 5 (Switch_A_1) connect with the HBAs on the storage controllers. Port 6 in domain 5 (Switch_A_1) connects to an FC-to-SAS bridge.

Note: You should give each zone a descriptive name. In this example, it is "STOR_A_1_6" that identifies the zone as a storage zone for the target port 6 at Site_A.

```
Switch_A_1:admin> zonecreate "STOR_A_1_6", "5,1; 5,2; 5,4; 5,5; 7,1;
7,2; 7,4; 7,5; 7,6"
Switch_A_1:admin> zonecreate "STOR_B_1_6", "5,1; 5,2; 5,4; 5,5; 7,1;
7,2; 7,4; 7,5; 7,6"
```

3. Create the configuration: `cfgcreate config_name , zone ; zone ...`

Example

```
Switch_A_1:admin> cfgcreate "CFG_1", "QOSH1_FCVI_1; STOR_A_1_6;
STOR_B_1_6"
```

4. Enter the `cfgadd config_name zone ; zone ...` command if you want to add more zones to the configuration.
5. Enable the configuration: `cfgenable config_name`

Example

```
Switch_A_1:admin> cfgenable "CFG_1"
```

6. Save the configuration: `cfgsave`

Example

```
Switch_A_1:admin> cfgsave
```

7. Verify the zoning configuration: zone --validate**Example**

```
Switch_A_1:admin> zone --validate
Defined configuration:
cfg: CFG_1 QOSH1_FCVI_1; STOR_A_1_6; STOR_B_1_6
zone: QOSH1_FCVI_1
      5,0; 7,0; 5,3; 7,3
zone: STOR_A_1_6
      5,1; 5,2; 5,4; 5,5; 7,1; 7,2; 7,4; 7,5; 5,6
zone: storage_zone_5_7
      5,1; 5,2; 5,4; 5,5; 7,1; 7,2; 7,4; 7,5; 7,6

Effective configuration:
cfg: CFG_1
zone: QOSH1_FCVI_1
      5,0
      7,0
      5,3
      7,3
zone: STOR_A_1_6
      5,1
      5,2
      5,4
      5,5
      7,1
      7,2
      7,4
      7,5
      5,6
zone: STOR_B_1_6
      5,1
      5,2
      5,4
      5,5
      7,1
      7,2
      7,4
      7,5
      7,6
-----
~ - Invalid configuration
* - Member does not exist
# - Invalid usage of broadcast zone
```

After you finish

You must configure ISL encryption, or, if you are not using ISL encryption, repeat the configuration tasks on the second FC switch fabric.

Setting ISL encryption on Brocade 6510 switches

On Brocade 6510 switches, you can optionally use the Brocade encryption feature on the ISL connections. If you want to use the encryption feature, you must perform additional configuration steps on each switch in the MetroCluster configuration.

Before you begin

- You must have Brocade 6510 switches.
- You must have selected two switches from the same fabric.

About this task

The steps must be performed on both the switches in the same fabric.

Disabling virtual fabric

In order to set the ISL encryption, you must disable the virtual fabric on all the four switches being used in a MetroCluster configuration.

Step

1. Disable the virtual fabric by entering the following command at the switch console:

```
fosconfig --disable vf
```

After you finish

Reboot the switch.

Setting the payload

After disabling the virtual fabric, you must set the payload or the data field size on both switches in the fabric.

About this task

The data field size must not exceed 2048.

Steps

1. Disable the switch:

```
switchdisable
```

2. Configure and set the payload:

```
configure
```

3. Set the following switch parameters:

- a. Set the `Fabric` parameter as follows:

```
y
```

- b. Set the other parameters such as `Domain`, `WWN Based persistent PID` and so on.

- c. Set the data field size as follows:

```
2048
```

4. Enable the switch:

```
switchenable
```

Setting the authentication policy

You must set the authentication policy, type, group and secret.

About this task

The commands must be executed at the switch console.

Steps

1. Disable the switch by entering the following command:

```
switchdisable
```

2. Set the authentication policy on the switch to **on** by entering the following command:

```
authUtil --policy -sw on
```

It displays the following output:

```
Warning: Activating the authentication policy requires either DH-CHAP
secrets or PKI certificates depending on the protocol selected.
Otherwise, ISLs will be segmented during next E-port bring-up.
ARE YOU SURE (yes, y, no, n): [no] yes
Auth Policy is set to ON
```

3. Set the authentication type to **dhchap** by entering the following command:

```
authUtil --set -a dhchap
```

It displays the following output:

```
Authentication is set to dhchap.
```

4. Set the authentication group to 4 by entering the following command:

```
authUtil --set -g 4
```

5. Set the authentication secret by performing the following steps:

- a. Provide the wwn of the other switch in the fabric for the parameter, Enter peer WWN, Domain, or switch name.
- b. Provide the peer secret for the parameter Enter peer secret.
- c. Provide the local secret for the parameter Enter local secret.
- d. Enter the following value for the parameter Are you done.

y

Example

The following is an example of setting authentication secret.

```
brod> secAuthSecret --set
```

This command is used to set up secret keys for the DH-CHAP authentication.

The minimum length of a secret key is 8 characters and maximum 40 characters. Setting up secret keys does not initiate DH-CHAP authentication. If switch is configured to do DH-CHAP, it is performed whenever a port or a switch is enabled.

Warning: Please use a secure channel for setting secrets. Using an insecure channel is not safe and may compromise secrets.

Following inputs should be specified for each entry.

1. WWN for which secret is being set up.
2. Peer secret: The secret of the peer that authenticates to peer.
3. Local secret: The local secret that authenticates peer.

Press enter to start setting up secrets > <cr>

Enter peer WWN, Domain, or switch name (Leave blank when done):
10:00:00:05:33:76:2e:99

Enter peer secret: <hidden>

Re-enter peer secret: <hidden>

Enter local secret: <hidden>

Re-enter local secret: <hidden>

Enter peer WWN, Domain, or switch name (Leave blank when done):

Are you done? (yes, y, no, n): [no] yes

Saving data to key store... Done.

6. Enable the switch by entering the following command:

```
switchenable
```

Enabling ISL encryption in a Brocade switch

After setting the authentication policy and secret, you must enable ISL encryption on the Brocade 6510 switches.

About this task

- These steps should be performed on one switch fabric at a time.
- The commands must be run at the switch console.

Steps

1. Disable the switch:

```
switchdisable
```

2. Enable encryption on all the ISL ports:

```
portCfgEncrypt --enable port_number
```

Example

In the following example, the encryption is enabled on ports 8 and 12:

```
portCfgEncrypt --enable 8
```

```
portCfgEncrypt --enable 12
```

3. Enable the switch:

```
switchenable
```

4. Verify that the ISL is up and working:

```
islshow
```

5. Verify that the encryption is enabled:

```
portenccompshow
```

Example

The following example shows the encryption is enabled on ports 8 and 12:

User	Encryption	
Port	configured	Active
----	-----	-----
8	yes	yes

9	No	No
10	No	No
11	No	No
12	yes	yes

After you finish

Perform all the steps on the switches in the other fabric in a MetroCluster configuration.

Configuring the Cisco FC switches

You must configure each Cisco switch in the MetroCluster configuration.

About this task

The following requirements apply to the Cisco FC switches:

- You must be using four supported Cisco switches of the same model with the same NX-OS version and licensing.
[NetApp Interoperability Matrix Tool](#)
- The MetroCluster configuration requires four switches.
The four switches must be connected into two fabrics of two switches each, with each fabric spanning both sites.
- Not all switches are supported for connectivity to the ATTO FiberBridge model.
[NetApp Interoperability Matrix Tool](#)
- Encryption and compression in the Cisco FC storage fabric is not supported in the MetroCluster configuration.

The following requirements apply to the Inter-Switch Link (ISL) connections:

- Fibre Channel over IP (FCIP) is not supported for ISL connections in a MetroCluster environment.
- ISLs of different speeds and lengths are supported between switches in the same fabric.

The following requirements apply to the storage connections:

- Two initiator ports must be connected from each storage controller to each fabric.
Each storage controller must have four initiator ports available to connect to the switch fabrics.

Steps

1. [Reviewing Cisco license requirements](#) on page 74
2. [Setting the Cisco FC switch to factory defaults](#) on page 74
3. [Configure the Cisco FC switch basic settings and community string](#) on page 75
4. [Acquiring licenses for ports](#) on page 76

5. [Enabling ports in a Cisco MDS 9148 switch](#) on page 78
6. [Configuring the F-ports on a Cisco FC switch](#) on page 79
7. [Assigning buffer-to-buffer credits to F-Ports in the same port group as the ISL](#) on page 81
8. [Creating and configuring the VSANs on Cisco FC switches](#) on page 83
9. [Configuring the E-ports on the Cisco FC switch](#) on page 88
10. [Configuring zoning on a Cisco FC switch](#) on page 92
11. [Ensuring the FC switch configuration is saved](#) on page 95

Related information

[NetApp Interoperability Matrix Tool](#)

Reviewing Cisco license requirements

Certain feature-based licenses might be required for the switches in a MetroCluster configuration. These licenses enable you to use features such as QoS or long-distance mode credits on the switches. You must install these licenses on all four switches in a MetroCluster configuration.

The following are the feature-based licenses that might be required in a MetroCluster configuration:

- **ENTERPRISE_PKG**
This enables you to use the QoS feature in Cisco switches.
- **PORT_ACTIVATION_PKG**
You can use this license for Cisco 9148 switches. This license enables you to activate or deactivate ports on the switches as long as only 16 ports are active at any given time. By default, 16 ports are enabled in Cisco MDS 9148 switches.
- **FM_SERVER_PKG**
This enables you to manage fabrics simultaneously and to manage switches through a web browser.
The FM_SERVER_PKG license also enables performance management features such as performance thresholds, threshold monitoring, and so on. For more information about this license, see the Cisco Fabric Manager Server Package.

You can verify that the licenses are installed by using the `show license usage` command. If you do not have these licenses, contact your sales representative before proceeding.

Setting the Cisco FC switch to factory defaults

To ensure a successful configuration, you must set the switch to its factory defaults. This ensures that the switch is starting from a clean configuration.

About this task

This task must be performed on all switches in the MetroCluster configuration.

Steps

1. Make a console connection and log in to both switches in the same fabric.
2. Issue the following command to set the switch back to its default settings:

write erase

You can respond **y** when prompted to confirm the command. This erases all licenses and configuration information on the switch.

3. Issue the following command to reboot the switch:

reload

You can respond **y** when prompted to confirm the command.

4. Repeat the `write erase` and `reload` commands on the other switch.

After issuing the `reload` command, the switch reboots and then prompts with setup questions. At that point, proceed to the next section.

The following example shows the process on a fabric consisting of FC_switch_A_1 and FC_switch_B_1.

```
FC_Switch_A_1# write erase
Warning: This command will erase the startup-configuration.
Do you wish to proceed anyway? (y/n) [n] y
FC_Switch_A_1# reload
This command will reboot the system. (y/n)? [n] y

FC_Switch_B_1# write erase
Warning: This command will erase the startup-configuration.
Do you wish to proceed anyway? (y/n) [n] y
FC_Switch_B_1# reload
This command will reboot the system. (y/n)? [n] y
```

Configure the Cisco FC switch basic settings and community string

You must specify the basic settings with the `setup` command or after issuing the `reload` command.

Steps

1. If the switch does not display the setup questions, issue the following command to configure the basic switch settings:

setup

2. Accept the default responses to the setup questions until you are prompted for the SNMP community string.

3. Set the community string to **public** (all lowercase) to allow access from the Data ONTAP Health Monitors.

Example

The following example shows the commands on FC_switch_A_1:

```
FC_switch_A_1# setup
  Configure read-only SNMP community string (yes/no) [n]: y
  SNMP community string : public
  Note: Please set the SNMP community string to "Public" or
  another value of your choosing.
  Configure default switchport interface state (shut/noshut)
[shut]: noshut
  Configure default switchport port mode F (yes/no) [n]: n
  Configure default zone policy (permit/deny) [deny]: deny
  Enable full zoneset distribution? (yes/no) [n]: yes
```

The following example shows the commands on FC_switch_B_1:

```
FC_switch_B_1# setup
  Configure read-only SNMP community string (yes/no) [n]: y
  SNMP community string : public
  Note: Please set the SNMP community string to "Public" or
  another value of your choosing.
  Configure default switchport interface state (shut/noshut)
[shut]: noshut
  Configure default switchport port mode F (yes/no) [n]: n
  Configure default zone policy (permit/deny) [deny]: deny
  Enable full zoneset distribution? (yes/no) [n]: yes
```

Acquiring licenses for ports

You do not have to use Cisco switch licenses on a continuous range of ports; instead, you can acquire licenses for specific ports that are used and remove licenses from unused ports. You should verify the number of licensed ports in the switch configuration and, if necessary, move licenses from one port to another as needed.

Steps

1. Issue the following command to show license usage for a switch fabric:

```
show port-resources module 1
```

Determine which ports require licenses. If some of those ports are unlicensed, determine if you have extra licensed ports and consider removing the licenses from them.

2. Issue the following command to enter configuration mode:

```
config t
```

3. Remove the license from the selected port:

- a. Issue the following command to select the port to be unlicensed:

```
interface interface-name
```

- b. Remove the license from the port using the following command:

```
no port-license acquire
```

- c. Exit the port configuration interface:

```
exit
```

4. Acquire the license for the selected port:

- a. Issue the following command to select the port to be unlicensed:

```
interface interface-name
```

- b. Make the port eligible to acquire a license using the "port license" command:

```
port-license
```

- c. Acquire the license on the port using the following command:

```
port-license acquire
```

- d. Exit the port configuration interface:

```
exit
```

5. Repeat for any additional ports.

6. Issue the following command to exit configuration mode:

```
exit
```

Removing and acquiring a license on a port

This example shows a license being removed from port fc1/2, port fc1/1 being made eligible to acquire a license, and the license being acquired on port fc1/1:

```
Switch_A_1# conf t
Switch_A_1(config)# interface fc1/2
Switch_A_1(config)# shut
Switch_A_1(config-if)# no port-license acquire
Switch_A_1(config-if)# exit
Switch_A_1(config)# interface fc1/1
Switch_A_1(config-if)# port-license
Switch_A_1(config-if)# port-license acquire
Switch_A_1(config-if)# no shut
Switch_A_1(config-if)# end
Switch_A_1# copy running-config startup-config

Switch_B_1# conf t
```

```
Switch_B_1(config)# interface fc1/2
Switch_B_1(config)# shut
Switch_B_1(config-if)# no port-license acquire
Switch_B_1(config-if)# exit
Switch_B_1(config)# interface fc1/1
Switch_B_1(config-if)# port-license
Switch_B_1(config-if)# port-license acquire
Switch_B_1(config-if)# no shut
Switch_B_1(config-if)# end
Switch_B_1# copy running-config startup-config
```

The following example shows port license usage being verified:

```
Switch_A_1# show port-resources module 1
Switch_B_1# show port-resources module 1
```

Enabling ports in a Cisco MDS 9148 switch

In Cisco MDS 9148 switches, you must manually enable the ports required in a MetroCluster configuration.

About this task

- You can manually enable 16 ports in a Cisco MDS 9148 switch.
- The Cisco switches enable you to apply the POD license on random ports, as opposed to applying them in sequence.
- Cisco switches require that you use one port from each port group, unless you need more than 12 ports.

Steps

1. View the port groups available in a Cisco switch:
2. License and acquire the required port in a port group by entering the following commands in sequence:

```
show port-resources module blade_number
```

```
config t
```

```
interface port_number
```

```
shut
```

```
port-license acquire
```

```
no shut
```

Example

For example, the following command licenses and acquires Port fc 1/45:

```
switch# config t
switch(config)#
switch(config)# interface fc 1/45
switch(config-if)#
switch(config-if)# shut
switch(config-if)# port-license acquire
switch(config-if)# no shut
switch(config-if)# end
```

- 3. Save the configuration:
`copy running-config startup-config`

Configuring the F-ports on a Cisco FC switch

You must configure the F-ports on the FC switch. In a MetroCluster configuration, the F-ports are the ports that connect the switch to the HBA initiators, FC-VI interconnects and FC-to-SAS bridges. Each port must be configured individually.

About this task

The following table lists the ports that must be configured as F-ports (switch-to-node) and shows what each port connects to:

Configure this port to F-mode:	Port connects to...
1	controller_x_1 FC-VI port 1
2	controller_x_1 HBA port 1
3	controller_x_1 HBA port 2
4	controller_x_2 FC-VI port 1
5	controller_x_2 HBA 1
6	controller_x_2 HBA 2
7	FC-to-SAS bridge

This task must be performed on each switch in the MetroCluster configuration.

Steps

- 1. Issue the following command to enter configuration mode:
`config t`
- 2. Enter interface configuration mode for the port:

```
interface port-ID
```

3. Shut down the port:

```
shutdown
```

4. Set the ports to F mode by issuing the following command:

```
switchport mode F
```

5. Set the ports to fixed speed by issuing the following command:

```
switchport speed speed
```

speed is either **8000** or **16000**

6. Set the rate mode of the switch port to dedicated by issuing the following command:

```
switchport rate-mode dedicated
```

7. Restart the port:

```
no shutdown
```

8. Issue the following command to exit configuration mode:

```
end
```

The following example shows the commands on the two switches:

```
Switch_A_1# config t
FC_switch_A_1(config)# interface fc 1/1
FC_switch_A_1(config-if)# shutdown
FC_switch_A_1(config-if)# switchport mode F
FC_switch_A_1(config-if)# switchport speed 8000
FC_switch_A_1(config-if)# switchport rate-mode dedicated
FC_switch_A_1(config-if)# no shutdown
FC_switch_A_1(config-if)# end
FC_switch_A_1# copy running-config startup-config

FC_switch_B_1# config t
FC_switch_B_1(config)# interface fc 1/1
FC_switch_B_1(config-if)# switchport mode F
FC_switch_B_1(config-if)# switchport speed 8000
FC_switch_B_1(config-if)# switchport rate-mode dedicated
FC_switch_B_1(config-if)# no shutdown
FC_switch_B_1(config-if)# end
FC_switch_B_1# copy running-config startup-config
```


Assigning buffer-to-buffer credits to F-Ports in the same port group as the ISL

You must assign the buffer-to-buffer credits to the F-ports if they are in the same port group as the ISL. If the ports do not have the required buffer-to-buffer credits, the ISL could be inoperative. This task is not required if the F-ports are not in the same port group as the ISL port.

About this task

If the F-Ports are in a port group that contains the ISL, this task must be performed on each FC switch in the MetroCluster configuration.

Steps

1. Issue the following command to enter configuration mode:
`config t`
2. Enter the following command to set the interface configuration mode for the port:
`interface port-ID`
3. Disable the port:
`shut`
4. If the port is not already in F mode, set the port to F mode by entering the following command:
`switchport mode F`
5. Set the buffer-to-buffer credit of the non-E ports to 1 by using the following command:
`switchport fcrxbbcredit 1`
6. Re-enable the port:
`no shut`
7. Exit configuration mode:
`exit`
8. Copy the updated configuration to the startup configuration:
`copy running-config startup-config`
9. Verify the buffer-to-buffer credit assigned to a port by entering the following commands:
`show port-resources module 1`
10. Issue the following command to exit configuration mode:
`exit`
11. Repeat these steps on the other switch in the fabric.

12. Verify the settings:

```
show port-resource module 1
```

In this example, port fc1/40 is the ISL. Ports fc1/37, fc1/38 and fc1/39 are in the same port group and must be configured.

The following commands show the port range being configured for fc1/37 through fc1/39:

```
FC_switch_A_1# conf t
FC_switch_A_1(config)# interface fc1/37-39
FC_switch_A_1(config-if)# shut
FC_switch_A_1(config-if)# switchport mode F
FC_switch_A_1(config-if)# switchport fcrxbbcredit
1FC_switch_A_1(config-if)# no shut
FC_switch_A_1(config-if)# exit
FC_switch_A_1# copy running-config startup-config

FC_switch_B_1# conf t
FC_switch_B_1(config)# interface fc1/37-39
FC_switch_B_1(config-if)# shut
FC_switch_B_1(config-if)# switchport mode F
FC_switch_B_1(config-if)# switchport fcrxbbcredit 1
FC_switch_A_1(config-if)# no shut
FC_switch_A_1(config-if)# exit
FC_switch_B_1# copy running-config startup-config
```

The following commands and system output show that the settings are properly applied:

```
FC_switch_A_1# show port-resource module 1
...
Port-Group 11
  Available dedicated buffers are 93
```

Interfaces in the Port-Group Mode	B2B Credit	Bandwidth	Rate
	Buffers	(Gbps)	
fc1/37	32	8.0	dedicated
fc1/38	1	8.0	dedicated
fc1/39	1	8.0	dedicated
...			

```
FC_switch_B_1# port-resource module
...
Port-Group 11
  Available dedicated buffers are 93
```

Interfaces in the Port-Group	B2B Credit Buffers	Bandwidth (Gbps)	Rate	Mode
------------------------------	-----------------------	---------------------	------	------

fc1/37	32	8.0	dedicated
fc1/38	1	8.0	dedicated
fc1/39	1	8.0	dedicated
...			

Creating and configuring the VSANs on Cisco FC switches

You must create a VSAN for the FC-VI ports and a VSAN for the storage ports on each FC switch in the MetroCluster configuration. The VSANs should have a unique number and name. You must do additional configuration if you are using two ISLs with in-order delivery of frames.

About this task

The examples here use the following naming conventions:

Switch fabric	VSAN name	ID number
1	FCVI_1_10	10
	STOR_1_20	20
2	FCVI_2_30	30
	STOR_2_20	40

This task must be performed on each FC switch fabric.

Steps

1. Configure the FC-VI VSAN:
- a. Issue the following command to enter configuration mode if you have not done so already:

`config t`

b. Issue the following command to edit the VSAN database:

`vsan database`

c. Set the VSAN ID using the following command:

`vsan vsan-ID`

d. Set the VSAN name using the following command:

`vsan vsan-ID name vsan_name`

Example

The following example shows the commands on FC_switch__A_1:

```

FC_switch_A_1# conf t
FC_switch_A_1(config)# vsan database
FC_switch_A_1(config-vsan-db)# vsan 10
FC_switch_A_1(config-vsan-db)# vsan 10 name FCVI_1_10
FC_switch_A_1(config-vsan-db)# end
FC_switch_A_1# copy running-config startup-config

```

The following example shows the commands on FC_switch__B_1:

```

FC_switch_B_1# conf t
FC_switch_B_1(config)# vsan database
FC_switch_B_1(config-vsan-db)# vsan 10
FC_switch_B_1(config-vsan-db)# vsan 10 name FCVI_1_10
FC_switch_B_1(config-vsan-db)# end
FC_switch_B_1# copy running-config startup-config

```

2. Add ports to the FC-VI VSAN:
 - a. Issue the following command for each port in the VSAN:

vsan vsan-ID interface interface_name

For the FC-VI VSAN, the ports connecting the two local FC-VI ports will be added. In the following example, the ports are fc1/1 and fc1/13:

Example

```

FC_switch_A_1# conf t
FC_switch_A_1(config)# vsan database
FC_switch_A_1(config)# vsan 10 interface fc1/1
FC_switch_A_1(config)# vsan 10 interface fc1/13
FC_switch_A_1(config)# end
FC_switch_A_1# copy running-config startup-config
FC_switch_B_1# conf t
FC_switch_B_1(config)# vsan database
FC_switch_B_1(config)# vsan 10 interface fc1/1
FC_switch_B_1(config)# vsan 10 interface fc1/13
FC_switch_B_1(config)# end
FC_switch_B_1# copy running-config startup-config

```

3. Verify the members of the VSAN using the following command:

show vsan member

Example

```

FC_switch_A_1# show vsan member
FC_switch_B_1# show vsan member

```

4. Configure the VSAN to guarantee in-order-delivery of frames:

- a. Enable the in-order-guarantee of exchanges for the VSAN by entering the following command:

```
in-order-guarantee vsan vsan-ID
```

- b. Enable load balancing for the VSAN by entering the following command:

```
vsan vsan-ID loadbalancing src-dst-id
```

Example

The following example shows the commands on FC_switch__A_1:

```
FC_switch_A_1# config t
FC_switch_A_1(config)# in-order-guarantee vsan 10
FC_switch_A_1(config)# vsan database
FC_switch_A_1(config-vsan-db)# vsan 10 loadbalancing src-dst-id
FC_switch_A_1(config-vsan-db)# end
FC_switch_A_1# copy running-config startup-config
```

The following example shows the commands on FC_switch__B_1:

```
FC_switch_B_1# config t
FC_switch_B_1(config)# in-order-guarantee vsan 10
FC_switch_B_1(config)# vsan database
FC_switch_B_1(config-vsan-db)# vsan 10 loadbalancing src-dst-id
FC_switch_B_1(config-vsan-db)# end
FC_switch_B_1# copy running-config startup-config
```

5. Set QoS policies for the FC-VI VSAN:

- a. Enter configuration mode:
- b. Enable the QoS and create a class map by entering the following commands in sequence:

```
conf t
```

```
qos enable
```

```
qos class-map class_name match-any
```

- c. Add the class map created in a previous step to the policy map by entering the following command:

```
class class_name
```

- d. Set the priority by entering the following command:

```
priority high
```

- e. Add the VSAN to the policy map created in step 2 by entering the following command:

```
qos service policy policy_name vsan vsanid
```

- f. Copy the updated configuration to the startup configuration:

```
copy running-config startup-config
```

Example

The following example shows the commands on FC_switch__A_1:

```
FC_switch_A_1# conf t
FC_switch_A_1(config)# qos enable
FC_switch_A_1(config)# qos class-map FCVI_1_10_Class match-any
FC_switch_A_1(config)# qos policy-map FCVI_1_10_Policy
FC_switch_A_1(config-pmap)# class FCVI_1_10_Class
FC_switch_A_1(config-pmap-c)# priority high
FC_switch_A_1(config-pmap-c)# exit
FC_switch_A_1(config)# exit
FC_switch_A_1(config)# qos service policy FCVI_1_10_Policy vsan 10
FC_switch_A_1(config)# end
FC_switch_A_1# copy running-config startup-config
```

The following example shows the commands on FC_switch__B_1:

```
FC_switch_B_1# conf t
FC_switch_B_1(config)# qos enable
FC_switch_B_1(config)# qos class-map FCVI_1_10_Class match-any
FC_switch_B_1(config)# qos policy-map FCVI_1_10_Policy
FC_switch_B_1(config-pmap)# class FCVI_1_10_Class
FC_switch_B_1(config-pmap-c)# priority high
FC_switch_B_1(config-pmap-c)# exit
FC_switch_B_1(config)# exit
FC_switch_B_1(config)# qos service policy FCVI_1_10_Policy vsan 10
FC_switch_B_1(config)# end
FC_switch_B_1# copy running-config startup-config
```

6. Configure the storage VSAN:

- a. Set the VSAN ID using the following command:

```
vsan vsan-ID
```

- b. Set the VSAN name using the following command:

```
vsan vsan-ID name vsan_name
```

Example

The following example shows the commands on FC_switch_A_1:

```

FC_switch_A_1# conf t
FC_switch_A_1(config)# vsan database
FC_switch_A_1(config-vsan-db)# vsan 20
FC_switch_A_1(config-vsan-db)# vsan 20 name STOR_1_20
FC_switch_A_1(config-vsan-db)# end
FC_switch_A_1# copy running-config startup-config

```

The following example shows the commands on FC_switch_B_1:

```

FC_switch_B_1# conf t
FC_switch_B_1(config)# vsan database
FC_switch_B_1(config-vsan-db)# vsan 20
FC_switch_B_1(config-vsan-db)# vsan 20 name STOR_1_20
FC_switch_B_1(config-vsan-db)# end
FC_switch_B_1# copy running-config startup-config

```

7. Add ports to the storage VSAN.

For the storage VSAN, all ports connecting HBA or FC-to-SAS bridges must be added. In this example fc1/5, fc1/9, fc1/17, fc1/21, fc1/25, fc1/29, fc1/33 and fc1/37 are being added.

Example

The following example shows the commands on FC_switch_A_1:

```

FC_switch_A_1# conf t
FC_switch_A_1(config)# vsan database
FC_switch_A_1(config)# vsan 20 interface fc1/5
FC_switch_A_1(config)# vsan 20 interface fc1/9
FC_switch_A_1(config)# vsan 20 interface fc1/17
FC_switch_A_1(config)# vsan 20 interface fc1/21
FC_switch_A_1(config)# vsan 20 interface fc1/25
FC_switch_A_1(config)# vsan 20 interface fc1/29
FC_switch_A_1(config)# vsan 20 interface fc1/33
FC_switch_A_1(config)# vsan 20 interface fc1/37
FC_switch_A_1(config)# end
FC_switch_A_1# copy running-config startup-config

```

The following example shows the commands on FC_switch_B_1:

```

FC_switch_B_1# conf t
FC_switch_B_1(config)# vsan database
FC_switch_B_1(config)# vsan 20 interface fc1/5
FC_switch_B_1(config)# vsan 20 interface fc1/9
FC_switch_B_1(config)# vsan 20 interface fc1/17
FC_switch_B_1(config)# vsan 20 interface fc1/21
FC_switch_B_1(config)# vsan 20 interface fc1/25
FC_switch_B_1(config)# vsan 20 interface fc1/29

```

```

FC_switch_B_1(config)# vsan 20 interface fc1/33
FC_switch_B_1(config)# vsan 20 interface fc1/37
FC_switch_B_1(config)# end
FC_switch_B_1# copy running-config startup-config

```

Configuring the E-ports on the Cisco FC switch

You must configure the FC switch ports that connect the ISL. These are the E-ports, and configuration must be done for each port. To do so, the correct number of buffer-to-buffer credits (BBCs) must be calculated.

About this task

All ISLs in the fabric must be configured with the same speed and distance settings.

This task must be performed on each ISL port.

Steps

1. Use the following table to determine the adjusted required BBCs per kilometer for possible port speeds.

To determine the correct number of BBCs, you multiply the *Adjusted BBCs required* (determined from the table below) by the *distance in kilometers between the switches*. The adjustment factor of 1.5 is required to account for FC-VI framing behavior.

Speed in Gbps	BBCs required per kilometer	Adjusted BBCs required (BBCs per km x 1.5)
1	0.5	0.75
2	1	1.5
4	2	3
8	4	6
16	8	12

Example

For example, to compute the required number of credits for a distance of 30 km on a 4-Gbps link, make the following calculation:

- *Speed in Gbps* is 4
- *Adjusted BBCs required* is 3.
- *Distance in kilometers between switches* is 30 km.
- $3 \times 30 = 90$

2. Issue the following command to enter configuration mode:
`config t`
3. Specify the port you are configuring by entering the following command:
`interface port-name`
4. Shut down the port:
`shutdown`
5. Set the rate mode of the port to **dedicated**:
`switchport rate-mode dedicated`
6. Set the speed for the port:
`switchport speed speed`
7. Set the buffer-to-buffer credits for the port:
`switchport fcrxbbcredit number of buffers`
8. Set the port to E mode:
`switchport mode E`
9. Enable the trunk mode for the port:
`switchport trunk mode on`
10. Add the ISL VSANs to the trunk:
`switchport trunk allowed vsan 10`
`switchport trunk allowed vsan add 20`
11. Add the port to port channel 1:
`channel-group 1`
12. Repeat the previous steps for the matching ISL port on the partner switch in the fabric.

Example

The following example shows port fc1/41 configured for a distance of 30 km and 8 Gbps:

```
FC_switch_A_1# conf t
FC_switch_A_1# shutdown
FC_switch_A_1# switchport rate-mode dedicated
FC_switch_A_1# switchport speed 8000
FC_switch_A_1# switchport fcrxbbcredit 60
FC_switch_A_1# switchport mode E
FC_switch_A_1# switchport trunk mode on
FC_switch_A_1# switchport trunk allowed vsan 10
FC_switch_A_1# switchport trunk allowed vsan add 20
```

```

FC_switch_A_1# channel-group 1
fc1/36 added to port-channel 1 and disabled

FC_switch_B_1# conf t
FC_switch_B_1# shutdown
FC_switch_B_1# switchport rate-mode dedicated
FC_switch_B_1# switchport speed 8000
FC_switch_B_1# switchport fcrxbbcredit 60
FC_switch_B_1# switchport mode E
FC_switch_B_1# switchport trunk mode on
FC_switch_B_1# switchport trunk allowed vsan 10
FC_switch_B_1# switchport trunk allowed vsan add 20
FC_switch_B_1# channel-group 1
fc1/36 added to port-channel 1 and disabled

```

13. Issue the following command on both switches to restart the ports:

```
no shutdown
```

14. Repeat the previous steps for the other ISL ports in the fabric.
15. Add native vsan to port-channel interface on both switches in the same fabric:

```
interface port-channel number

switchport trunk allowed vsan add native_san_id
```

16. Verify configuration of the port-channel:

```
show interface port-channel number
```

The port channel should have the following attributes:

- The port-channel is `trunking`.
- Admin port mode is `E`, trunk mode is `on`.
- Speed shows the cumulative value of all the ISL link speeds.
For example, two ISL ports operating at 4 Gbps should show a speed of 8 Gbps.
- Trunk vsans (admin allowed and active) shows all the allowed VSANs.
- Trunk vsans (up) shows all the allowed VSANs.
- The member list shows all the ISL ports that were added to the port-channel.
- The port VSAN number should be the same as the VSAN that contains the ISLs (usually native vsan 1).

Example

```

FC_switch_A_1(config-if)# show int port-channel 1
port-channel 1 is trunking
Hardware is Fibre Channel
Port WWN is 24:01:54:7f:ee:e2:8d:a0

```

```

Admin port mode is E, trunk mode is on
snmp link state traps are enabled
Port mode is TE
Port vsan is 1
Speed is 8 Gbps
Trunk vsans (admin allowed and active) (1,10,20)
Trunk vsans (up) (1,10,20)
Trunk vsans (isolated) ( )
Trunk vsans (initializing) ( )
5 minutes input rate 1154832 bits/sec,144354 bytes/sec, 170
frames/sec
5 minutes output rate 1299152 bits/sec,162394 bytes/sec, 183
frames/sec
535724861 frames input,1069616011292 bytes
0 discards,0 errors
0 invalid CRC/FCS,0 unknown class
0 too long,0 too short
572290295 frames output,1144869385204 bytes
0 discards,0 errors
5 input OLS,11 LRR,2 NOS,0 loop inits
14 output OLS,5 LRR, 0 NOS, 0 loop inits
Member[1] : fcl/36
Member[2] : fcl/40
Interface last changed at Thu Oct 16 11:48:00 2014

```

17. Exit interface configuration on both switches:

```
end
```

18. Copy the updated configuration to the startup configuration on both fabrics:

```
copy running-config startup-config
```

Example

```

FC_switch_A_1(config-if)# end
FC_switch_A_1# copy running-config startup-config

FC_switch_B_1(config-if)# end
FC_switch_B_1# copy running-config startup-config

```

19. Repeat the previous steps on the second switch fabric.

Related concepts

Recommended port assignments for FC switches on page 39

Configuring zoning on a Cisco FC switch

You must assign the switch ports to separate zones to isolate storage (HBA) and controller (FC-VI) traffic.

About this task

These steps must be performed on both FC switch fabrics.

Steps

1. Clear the existing zones and zoneset if present.

- a. Determine which zones and zonesets are active:

```
show zoneset active
```

Example

```
FC_switch_A_1# show zoneset active
FC_switch_B_1# show zoneset active
```

- b. Disable the active zonesets identified in the previous step:

```
no zoneset activate name zoneset_name vsan vsan_id
```

Example

The following example shows two zonesets being disabled:

- ZoneSet_A on FC_switch_A_1 in VSAN 10
- ZoneSet_B on FC_switch_B_1 in VSAN 20

```
FC_switch_A_1# no zoneset activate name ZoneSet_A vsan 10
FC_switch_B_1# no zoneset activate name ZoneSet_B vsan 20
```

- c. After all zonesets are deactivated, clear the zone database:

```
clear zone database zone-name
```

Example

```
FC_switch_A_1# clear zone database 10
FC_switch_A_1# copy running-config startup-config

FC_switch_B_1# clear zone database 20
FC_switch_B_1# copy running-config startup-config
```

2. Obtain the switch worldwide name (WWN):

```
show wwn switch
```

3. Configure the basic zone settings:

- a. Set the default zoning policy to deny:

```
no system default zone default-zone permit
```

- b. Enable the full zone distribution:

```
system default zone distribute full
```

- c. Set the default zoning policy for each VSAN:

```
no zone default-zone permit vsanid
```

- d. Set the default full zone distribution for each VSAN:

```
zoneset distribute full vsanid
```

Example

```
FC_switch_A_1# conf t
FC_switch_A_1(config)# no system default zone default-zone permit
FC_switch_A_1(config)# system default zone distribute full
FC_switch_A_1(config)# no zone default-zone permit 10
FC_switch_A_1(config)# no zone default-zone permit 20
FC_switch_A_1(config)# zoneset distribute full vsan 10
FC_switch_A_1(config)# zoneset distribute full vsan 20
FC_switch_A_1(config)# end
FC_switch_A_1# copy running-config startup-config

FC_switch_B_1# conf t
FC_switch_B_1(config)# no system default zone default-zone permit
FC_switch_B_1(config)# system default zone distribute full
FC_switch_B_1(config)# no zone default-zone permit 10
FC_switch_B_1(config)# no zone default-zone permit 20
FC_switch_B_1(config)# zoneset distribute full vsan 10
FC_switch_B_1(config)# zoneset distribute full vsan 20
FC_switch_B_1(config)# end
FC_switch_B_1# copy running-config startup-config
```

4. Create storage zones and add the storage ports to it.

These steps only need to be performed on one switch in each fabric.

Each storage zone contains the HBA initiator ports from all controllers and one single port connecting an FC-to-SAS bridge. Each zone has 9 members.

- a. Create the storage zones by entering the following command:

```
zone name STOR_zone-name vsan vsanid
```

- b. Add storage ports to the zone by entering the following command:

```
member STOR_zone-name
```

- c. Activate the zone set by entering the following command:

```
zoneset activate name STOR_zonenameesetname vsan vsanid
```

Example

```
FC_switch_A_1# conf t
FC_switch_A_1(config)# zone name STOR_Zone_1_20_25 vsan 20
FC_switch_A_1(config-zone)# member interface fc1/5 swwn
20:00:00:05:9b:24:cb:78
FC_switch_A_1(config-zone)# member interface fc1/9 swwn
20:00:00:05:9b:24:cb:78
FC_switch_A_1(config-zone)# member interface fc1/17 swwn
20:00:00:05:9b:24:cb:78
FC_switch_A_1(config-zone)# member interface fc1/21 swwn
20:00:00:05:9b:24:cb:78
FC_switch_A_1(config-zone)# member interface fc1/5 swwn
20:00:00:05:9b:24:12:99
FC_switch_A_1(config-zone)# member interface fc1/9 swwn
20:00:00:05:9b:24:12:99
FC_switch_A_1(config-zone)# member interface fc1/17 swwn
20:00:00:05:9b:24:12:99
FC_switch_A_1(config-zone)# member interface fc1/21 swwn
20:00:00:05:9b:24:12:99
FC_switch_A_1(config-zone)# member interface fc1/25 swwn
20:00:00:05:9b:24:cb:78
FC_switch_A_1(config-zone)# end
FC_switch_A_1# copy running-config startup-config
```

5. Create an FCVI zone set and add the FCVI ports to it:

These steps only need to be performed on one switch in the fabric.

- a. Create the FCVI zone set by entering the following command:

```
zoneset name FCVI_zonesetname vsan vsanid
```

- b. Add FCVI zones to the zone set by entering the following command:

```
member FCVI_zonename
```

- c. Activate the zone set by entering the following command:

```
zoneset activate name FCVI_zonesetname vsan vsanid
```

Example

```
FC_switch_A_1# conf t
FC_switch_A_1(config)# zoneset name FCVI_Zoneset_1_20 vsan 20
FC_switch_A_1(config-zoneset)# member FCVI_Zone_1_20_25
FC_switch_A_1(config-zoneset)# member FCVI_Zone_1_20_29
...
FC_switch_A_1(config-zoneset)# exit
FC_switch_A_1(config)# zoneset activate name FCVI_ZoneSet_1_20 vsan 20
FC_switch_A_1(config)# exit
FC_switch_A_1# copy running-config startup-config
```

6. Verify the zoning:

```
show zone
```

7. Repeat the previous steps on the second FC switch fabric.

Ensuring the FC switch configuration is saved

You must make sure the FC switch configuration is saved to the startup config on all switches.

Step

1. Issue the following command on both FC switch fabrics:

```
copy running-config startup-config
```

Example

```
FC_switch_A_1# copy running-config startup-config
FC_switch_B_1# copy running-config startup-config
```

Configuring hardware for sharing a Brocade 6510 FC fabric during transition

If your 7-Mode fabric MetroCluster configuration uses Brocade 6510 switches, you can share the existing switch fabrics with the new clustered MetroCluster configuration. Shared switch fabrics means the new MetroCluster configuration does not require a new, separate switch fabric. This temporary configuration is only supported with the Brocade 6510 switch for transition purposes.

Before you begin

- The 7-Mode fabric MetroCluster must be using Brocade 6510 switches.
If the MetroCluster configuration is currently not using Brocade 6510 switches, the switches must be upgraded to Brocade 6510 prior to using this procedure.
- The 7-Mode fabric MetroCluster configuration must be using SAS storage shelves only.
If the existing configuration includes FC storage shelves (such as the DS14mk4 FC), FC switch fabric sharing is not supported.
- All 48 ports of the Brocade 6510 switches must be licensed.
- The SFPs on the switch ports used by the new, clustered MetroCluster configuration must support 16-Gbps rates.
The existing 7-Mode fabric MetroCluster can remain connected to ports using 8-Gbps or 16-Gbps SFPs.
- On each of the four Brocade 6510 switches, ports 24 through 45 must be available to connect the ports of the new MetroCluster components.
- You should ensure that the existing Inter-Switch Links (ISLs) are on ports 46 and 47.
- The Brocade 6510 switches must be running a FOS firmware version that is supported on both the 7-Mode fabric MetroCluster and clustered Data ONTAP MetroCluster configuration.

After you finish

After sharing the fabric and completing the MetroCluster configuration, you can transition data from the 7-Mode fabric MetroCluster configuration.

After transitioning the data, you can remove the 7-Mode fabric MetroCluster cabling and, if desired, move the clustered Data ONTAP MetroCluster cabling to the lower-numbered ports previously used for the 7-Mode MetroCluster cabling. The ports are shown in [Reviewing FC switch port assignments for a four node MetroCluster](#) on page 39. You must adjust the zoning for the rearranged ports.

Steps

1. [Reviewing Brocade license requirements](#) on page 97

2. [Racking the hardware components](#) on page 97
3. [Cable the new MetroCluster controllers to the existing FC fabrics](#) on page 98
4. [Configuring switch fabrics sharing between the 7-Mode and clustered MetroCluster configuration](#) on page 100

Related information

[7-Mode Transition Tool 2.0 Data and Configuration Transition Guide](#)

Reviewing Brocade license requirements

You need certain licenses for the switches in a MetroCluster configuration. You must install these licenses on all four switches.

The MetroCluster configuration has the following Brocade license requirements:

- Trunking license for systems using more than one ISL, as recommended.
- Extended Fabric license (for ISL distances over 6 km)
- Enterprise license for sites with more than one ISL and an ISL distance greater than 6 km
The Enterprise license includes Brocade Network Advisor and all licenses except for additional port licenses.

You can verify that the licenses are installed by using the `licenseshow` command. If you do not have these licenses, contact your sales representative before proceeding.

Racking the hardware components

If you have not received the equipment already installed in cabinets, you must rack the components.

About this task

This task must be performed on both MetroCluster sites.

Steps

1. Plan out the positioning of the MetroCluster components.
The rack space will depend on the platform model of the storage controllers, the switch types, and the number of disk shelf stacks in your configuration.
2. Properly ground yourself.
3. Install the storage controllers in the rack or cabinet.

[Installation and Setup Instructions FAS8040/FAS8060 Systems](#)

Installation and setup Instructions FAS80xx Systems with I/O Expansion Modules

Installation and Setup Instructions FAS8020 systems

Installation and Setup Instructions 62xx Systems

Installation and Setup Instructions 32xx Systems

4. Install the FC switches in the rack or cabinet.
5. Install the disk shelves, power them on, and set the shelf IDs.

SAS Disk Shelves Installation and Service Guide for DS4243, DS2246, DS4486, and DS4246

- You must power-cycle each disk shelf.
- Shelf IDs must be unique for each SAS disk shelf within the entire MetroCluster configuration (including both sites).

6. Install each FC-to-SAS bridge:

- a. Secure the “L” brackets on the front of the bridge to the front of the rack (flush-mount) with the four screws.

The openings in the bridge “L” brackets are compliant with rack standard ETA-310-X for 19-inch (482.6 mm) racks.

For more information and an illustration of the installation, see the *ATTO FibreBridge 6500N Installation and Operation Manual*.

- b. Connect each bridge to a power source that provides a proper ground.
- c. Power on each bridge.

Note: For maximum resiliency, ensure that bridges attached to the same stack of disk shelves are connected to different power sources.

The bridge Ready LED might take up to 30 seconds to illuminate, indicating that the bridge has completed its power-on self test sequence.

Cable the new MetroCluster controllers to the existing FC fabrics

On each controller module in the clustered Data ONTAP MetroCluster configuration, the FC-VI adapter and HBAs must be cabled to specific ports on the existing FC switches.

Steps

1. Cable the FC-VI and HBA ports according to the following table:

Site A		Site B	
Connect this Site A component and port...	To this port on FC_switch_A_1.. ..	Connect this Site B component and port...	To this port on FC_switch_B_1...
controller_A_1 FC-VI port 1	32	controller_B_1 FC-VI port 1	32
controller_A_1 HBA port 1	33	controller_B_1 HBA port 1	33
controller_A_1 HBA port 2	34	controller_B_1 HBA port 2	34
controller_A_2 FC-VI port 1	35	controller_B_2 FC-VI port 1	35
controller_A_2 HBA 1	36	controller_B_2 HBA 1	36
controller_A_2 HBA 2	37	controller_B_2 HBA 2	37

2. Cable each FC-SAS bridge in the first switch fabric to the FC switches.

The number of bridges varies depending on the number of SAS storage stacks.

Site A		Site B	
Cable this site A bridge...	To this port on FC_switch_A_1.. ..	Cable this Site B bridge...	To this port on FC_switch_B_1...
bridge_A_1_38	38	bridge_B_1_38	38
bridge_A_1_39	39	bridge_B_1_39	39
bridge_A_1_40	40	bridge_B_1_40	40
bridge_A_1_41	41	bridge_B_1_41	41
bridge_A_1_42	42	bridge_B_1_42	42
bridge_A_1_43	43	bridge_B_1_43	43
bridge_A_1_44	44	bridge_B_1_44	44
bridge_A_1_45	45	bridge_B_1_45	45

3. Cable each bridge in the second switch fabric to the FC switches.

The number of bridges varies depending on the number of SAS storage stacks.

Site A		Site B	
Cable this site A bridge...	To this port on FC_switch_A_2. ..	Cable this Site B bridge...	To this port on FC_switch_B_2...
bridge_A_2_38	38	bridge_B_2_38	38
bridge_A_2_39	39	bridge_B_2_39	39
bridge_A_2_40	40	bridge_B_2_40	40
bridge_A_2_41	41	bridge_B_2_41	41
bridge_A_2_42	42	bridge_B_2_42	42
bridge_A_2_43	43	bridge_B_2_43	43
bridge_A_2_44	44	bridge_B_2_44	44
bridge_A_2_45	45	bridge_B_2_45	45

Configuring switch fabrics sharing between the 7-Mode and clustered MetroCluster configuration

To share switch fabrics between the existing 7-Mode fabric MetroCluster and the new MetroCluster configuration, you must set up specific zoning and other settings that are different than an unshared configuration.

About this task

This task must be performed on both switch fabrics, one at a time.

Disabling one of the switch fabrics

You must disable one of the switch fabrics so you can modify its configuration. After you complete the configuration and reenble the switch fabric, you will repeat the process on the other fabric.

Before you begin

You must have run the `fmc_dc` utility on the existing 7-Mode fabric MetroCluster configuration and resolved any issues prior to beginning the configuration process.

About this task

To ensure continued operation of the MetroCluster configuration, you must not disable the second fabric while the first fabric is disabled.

Steps

1. Disable each of the switches in the fabric:

switchCfgPersistentDisable

If this command is not available, use the `switchDisable` command.

Example

The following example shows the command issued on FC_switch_A_1:

```
FC_switch_A_1:admin> switchCfgPersistentDisable
```

The following example shows the command issued on FC_switch_B_1:

```
FC_switch_B_1:admin> switchCfgPersistentDisable
```

2. Ensure that the 7-Mode MetroCluster configuration is functioning correctly using the redundant fabric:
 - a. Confirm that controller failover is healthy:

cf status

Example

```
node_A> cf status
Controller Failover enabled, node_A is up.
VIA Interconnect is up (link 0 down, link 1 up).
```

- b. Confirm that disks are visible:

storage show disk -p

Example

```
node_A> storage show disk -p
```

PRIMARY	PORT	SECONDARY	PORT	SHELF	BAY
Brocade-6510-2K0GG:5.126L27	B			1	0
Brocade-6510-2K0GG:5.126L28	B			1	1
Brocade-6510-2K0GG:5.126L29	B			1	2
Brocade-6510-2K0GG:5.126L30	B			1	3
Brocade-6510-2K0GG:5.126L31	B			1	4
.					
.					
.					

- c. Confirm that the aggregates are healthy:

```
aggr status
```

Example

```
node_A> aggr status
      Aggr State      Status      Options
aggr0 online      raid_dp, aggr      root, nosnap=on
                    mirrored
                    64-bit
```

Deleting TI zoning and configuring IOD settings

You must delete the existing TI zoning and reconfigure in-order-delivery (IOD) settings on the switch fabric.

Steps

1. Identify the TI zones that are configured on the fabric:

```
zone --show
```

Example

The following example shows the zone FCVI_TI_FAB_2.

```
Brocade-6510:admin> zone --show
Defined TI zone configuration:
TI Zone Name:    FCVI_TI_FAB_2
Port List:      1,0; 1,3; 2,0; 2,3
configured Status: Activated / Failover-Disabled
Enabled Status: Activated / Failover-Disabled
```

2. Delete the TI zones:

```
zone --delete zone-name
```

Example

The following example shows the deletion of zone FCVI_TI_FAB_2.

```
Brocade-6510:admin> zone --delete FCVI_TI_FAB_2
```

3. Confirm that the zones have been deleted:

```
zone --show
```

Example

The output should be similar to the following:

```
Brocade-6510:admin> zone --show

Defined TI zone configuration:
no TI zone configuration defined
```

4. Save the configuration:

```
cfgsave
```

5. Enable in-order-delivery:

```
iodset
```

6. Select Advanced Performance Tuning (APT) policy 1, the Port Based Routing Policy:

```
aptpolicy 1
```

7. Disable Dynamic Load Sharing (DLS):

```
dlsreset
```

8. Verify the IOD settings using the followign commands:

```
iodshow
```

```
aptpolicy
```

```
dlsshow
```

Example

The output should be similar to the following:

```
Brocade-6510:admin> iodshow

IOD is set

Brocade-6510:admin> aptpolicy
Current Policy: 1

3 : Default Policy
1: Port Based Routing Policy
2: Device Based Routing Policy (FICON support only)
3: Exchange Based Routing Policy
Brocade-6510:admin> dlsshow

DLS is not set
```

Ensuring ISLs are in the same port group and configuring zoning

You must make sure that the Inter-Switch Links (ISLs) are in the same port group and configure zoning for the MetroCluster configurations to successfully share the switch fabrics.

Steps

1. If the ISLs are not in the same port group, move one of the ISL ports to the same port group as the other one.

You can use any available port except 32 through 45, which are used by the new MetroCluster configuration. The recommended ISL ports are 46 and 47.

2. Follow the steps in [Configuring zoning on a Brocade FC switch](#) on page 65 to enable trunking and the QoS zone.

The port numbers when sharing fabrics are different than those shown in the section. When sharing, use ports 46 and 47 for the ISL ports. If you moved your ISL ports, you need to use the procedure in [Configuring the E-ports \(ISL ports\) on a Brocade FC switch](#) on page 58 to configure the ports.

3. Follow the steps in [Configuring the non-E ports on the Brocade switch](#) on page 64 to configure the non-E ports.
4. Do not delete the zones or zone sets that already exist in the backend switches (for the 7-Mode fabric MetroCluster) except the Traffic Isolation (TI) zones in Step 3.
5. Follow the steps in [Configuring the E-ports \(ISL ports\) on a Brocade FC switch](#) on page 58 to add the zones required by the new MetroCluster to the existing zone sets.

Example

The following example shows the commands and system output for creating the zones:

```
Brocade-6510-2K0GG:admin> zonecreate "QOSH2_FCVI_1", "2,32; 2,35;
1,32; 1,35"

Brocade-6510-2K0GG:admin> zonecreate "STOR_A_2_47", "2,33; 2,34;
2,36; 2,37; 1,33; 1,34; 1,36; 1,37; 1,47"

Brocade-6510-2K0GG:admin> zonecreate "STOR_B_2_47", "2,33; 2,34;
2,36; 2,37; 1,33; 1,34; 1,36; 1,37; 2,47"

Brocade-6510-2K0GG:admin> cfgadd config_1_FAB2, "QOSH2_FCVI_1;
STOR_A_2_47; STOR_B_2_47"

Brocade-6510-2K0GG:admin> cfgenable "config_1_FAB2"
You are about to enable a new zoning configuration.
This action will replace the old zoning configuration with the
current configuration selected. If the update includes changes
```



```

to one or more traffic isolation zones, the update may result in
localized disruption to traffic on ports associated with
the traffic isolation zone changes
Do you want to enable 'config_1_FAB2' configuration (yes, y, no, n):
[no] yes

Brocade-6510-2K0GG:admin> cfigsave
You are about to save the Defined zoning configuration. This
action will only save the changes on Defined configuration.
Do you want to save the Defined zoning configuration only? (yes, y,
no, n): [no] yes
Nothing changed: nothing to save, returning ...
Brocade-6510-2K0GG:admin>

```

Reenabling the switch fabric and verify operation

You must enable the FC switch fabric and make sure that the switches and devices are operating correctly.

Steps

1. Enable the switches:

switchCfgPersistentEnable

If this command is not available, the switch should be in the enabled state after the `fastBoot` command is issued.

Example

The following example shows the command issued on FC_switch_A_1:

```
FC_switch_A_1:admin> switchCfgPersistentEnable
```

The following example shows the command issued on FC_switch_B_1:

```
FC_switch_B_1:admin> switchCfgPersistentEnable
```

2. Verify that the switches are online and all devices are properly logged in:

switchShow

Example

The following example shows the command issued on FC_switch_A_1:

```
FC_switch_A_1:admin> switchShow
```

The following example shows the command issued on FC_switch_B_1:

```
FC_switch_B_1:admin> switchShow
```

3. Run the `fmc_dc` utility to ensure that the 7-Mode fabric MetroCluster is functioning correctly.
You can ignore errors related to Traffic Isolation (TI) zoning and trunking.
4. Repeat the tasks for the second switch fabric.

Configuring the MetroCluster software in Data ONTAP

Each node in the MetroCluster must be set up in Data ONTAP, including the node-level configurations and the configuration of the nodes into two sites. Finally, the MetroCluster relationship is implemented between the two sites. The steps for systems with native disk shelves are slightly different than those for systems with array LUNs.

For new systems, you do not need to configure the Data ONTAP software except to change the pre-configured IP addresses. If your system is new, you can proceed to [Verifying the MetroCluster configuration](#) on page 137.

Steps

1. [Gathering required information and reviewing the workflow](#) on page 107
2. [Similarities and differences between regular cluster and MetroCluster configurations](#) on page 113
3. [Verifying disk assignment in Maintenance mode](#) on page 114
4. [Verifying the HA state of components is mcc in Maintenance mode](#) on page 118
5. [Running System Setup to configure the nodes and clusters](#) on page 119
6. [Configuring the clusters into a MetroCluster configuration](#) on page 122
7. [Checking for MetroCluster configuration errors with Config Advisor](#) on page 137
8. [Verifying local HA operation](#) on page 138
9. [Verifying switchover, healing, and switchback](#) on page 140
10. [Protecting configuration backup files](#) on page 140

Gathering required information and reviewing the workflow

You need to gather the required IP addresses for the controller and review the software installation workflow before you begin the configuration process.

Worksheet for IP network information for site A

You must obtain IP addresses and other network information for the first MetroCluster site (site A) from your network administrator before you configure the system.

Site A switch information

When you cable the system, you need a host name and management IP address for each cluster switch:

Cluster switch	Host name	IP address	Network mask	Default gateway
Interconnect 1 Not required if using two-node switchless cluster.				
Interconnect 2 Not required if using two-node switchless cluster.				
Management 1				
Management 2				

Site A cluster creation information

When you first create the cluster, you need the following information:

Type of information	Your values
Cluster name Example used in this guide: site_A	
DNS domain	
DNS name servers	
Location	
Administrator password	

Site A node information

For each node in the cluster, you need a management IP address, a network mask, and a default gateway:

Node	Port	IP address	Network mask	Default gateway
Node 1 Example used in this guide: controller_A_1				

Node	Port	IP address	Network mask	Default gateway
Node 2 Example used in this guide: controller_A_2				

Site A LIFs and ports for cluster peering

For each node in the cluster, you need the IP address of an intercluster LIF, a network mask, and a default gateway. The intercluster LIFs are used to peer the clusters:

Node	Port	IP address of intercluster LIF	Network mask	Default gateway
Node 1				
Node 2				

Site A time server information

You must synchronize the time, which will require one or more NTP time servers:

Node	Host name	IP address	Network mask	Default gateway
NTP server 1				
NTP server 2				

Site A AutoSupport information

You must configure AutoSupport on each node, which will require the following information:

Type of information		Your values
From email address		
Mail hosts	IP addresses or names	
Transport protocol	HTTP, HTTPS, or SMTP	
	Proxy server	
Recipient email addresses or distribution lists	Full-length messages	
	Concise messages	
	Partners	

Site A service processor information

You must enable access to the service processor of each node for troubleshooting and maintenance, which requires the following network information for each node:

Node	IP address	Network mask	Default gateway
Node 1			
Node 2			

Worksheet for IP network information for site B

You must obtain IP addresses and other network information for the second MetroCluster site (site B) from your network administrator before you configure the system.

Site B cluster switch information (if not using two-node switchless cluster configuration)

When you cable the system, you need a host name and management IP address for each cluster switch:

Cluster switch	Host name	IP address	Network mask	Default gateway
Interconnect 1 Not required if using two-node switchless cluster.				
Interconnect 2 Not required if using two-node switchless cluster.				
Management 1				
Management 2				

Site B cluster creation information

When you first create the cluster, you need the following information:

Type of information	Your values
Cluster name Example used in this guide: site_B	

Type of information	Your values
DNS domain	
DNS name servers	
Location	
Administrator password	

Site B node information

For each node in the cluster, you need a management IP address, a network mask, and a default gateway:

Node	Port	IP address	Network mask	Default gateway
Node 1 Example used in this guide: controller_B_1				
Node 2 Example used in this guide: controller_B_2				

Site B LIFs and ports for cluster peering

For each node in the cluster, you need the IP address of an intercluster LIF, a network mask, and a default gateway. The intercluster LIFs are used to peer the clusters.

Node	Port	IP address of intercluster LIF	Network mask	Default gateway
Node 1				
Node 2				

Site B time server information

You must synchronize the time, which will require one or more NTP time servers:

Node	Host name	IP address	Network mask	Default gateway
NTP server 1				
NTP server 2				

Site B autoSupport information

You must configure AutoSupport on each node, which will require the following information:

Type of information		Your values
From email address		
Mail hosts	IP addresses or names	
Transport protocol	HTTP, HTTPS, or SMTP	
	Proxy server	
Recipient email addresses or distribution lists	Full-length messages	
	Concise messages	
	Partners	

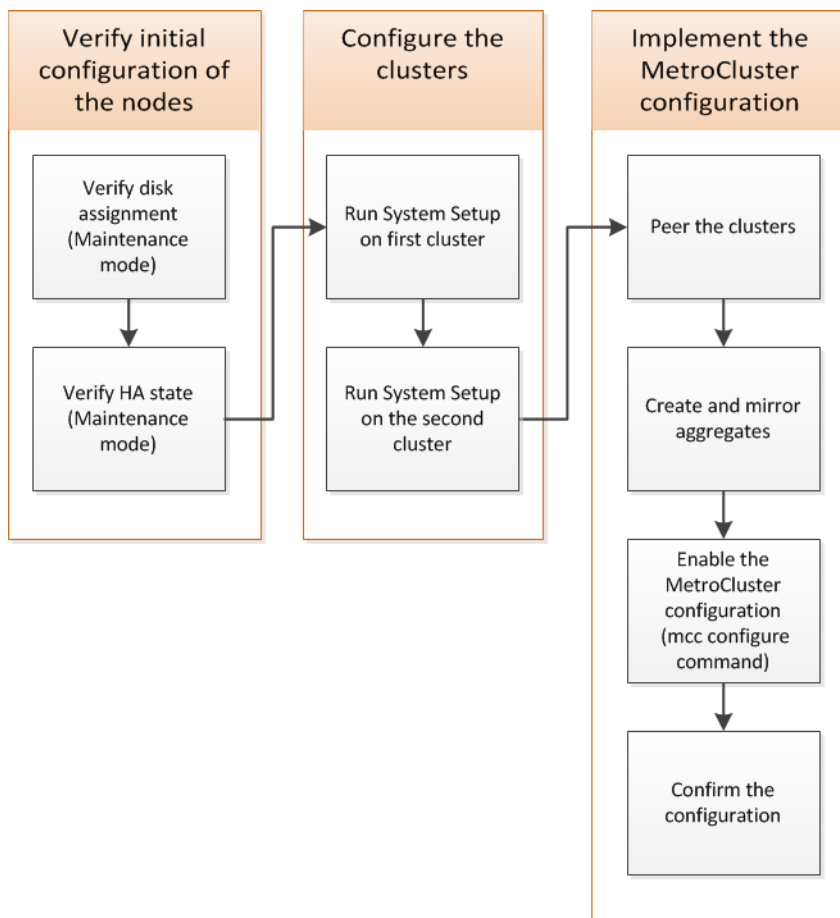
Site B service processor information

You must enable access to the service processor of each node for troubleshooting and maintenance, which requires the following network information for each node:

Node	IP address	Network mask	Default gateway
Node 1 (controller_B_1)			
Node 2 (controller_B_2)			

Software configuration workflow

Configuring the MetroCluster in Data ONTAP involves configuring each node, clustering the nodes into two clusters, one at each site, and then peering the clusters and issuing the command to implement the MetroCluster operations.



Similarities and differences between regular cluster and MetroCluster configurations

The configuration of the nodes in each cluster in a MetroCluster configuration is similar to that of nodes in a standard cluster.

The MetroCluster configuration is built on two standard clusters. Physically, the configuration must be symmetrical, with each node having the same hardware configuration, and all the MetroCluster

components must be cabled and configured. But the basic software configuration for nodes in a MetroCluster configuration is the same as that for nodes in a standard cluster.

Configuration step	Standard cluster configuration	MetroCluster configuration
Configure management, cluster, and data LIFs on each node	Same in both types of clusters	
Configure root aggregate		
Configure nodes in the cluster as HA pairs		
Set up cluster on one node in the cluster		
Join the other node to the cluster		
Create a mirrored root aggregate	Optional	Required
Create a mirrored data aggregate on each node	Optional	Required
Peer the clusters	Optional	Required
Enable the MetroCluster configuration	Does not apply	Required

Verifying disk assignment in Maintenance mode

Before fully booting the system to Data ONTAP, you can optionally boot to Maintenance mode and verify the disk assignment on the nodes. The disks should be assigned to create a fully symmetric active-active configuration, where each node and each pool have an equal number of disks assigned to them.

About this task

New MetroCluster systems have disk assignment completed prior to shipment.

The following table shows example pool assignments for a MetroCluster configuration. Disks are assigned to pools on a per-shelf basis.

Disk shelf (<i>example name</i>)...	At site...	Belongs to...	And is assigned to that node's...
Disk shelf 1 (shelf_A_1_1)	Site A	Node A 1	Pool 0
Disk shelf 2 (shelf_A_1_3)			
Disk shelf 3 (shelf_B_1_1)		Node B 1	Pool 1
Disk shelf 4 (shelf_B_1_3)			
Disk shelf 5 (shelf_A_2_1)		Node A 2	Pool 0
Disk shelf 6 (shelf_A_2_3)			
Disk shelf 7 (shelf_B_2_1)		Node B 2	Pool 1
Disk shelf 8 (shelf_B_2_3)			
Disk shelf 9 (shelf_B_1_2)	Site B	Node B 1	Pool 0
Disk shelf 10 (shelf_B_1_4)			
Disk shelf 11 (shelf_A_1_2)		Node A 1	Pool 1
Disk shelf 12 (shelf_A_1_4)			
Disk shelf 13 (shelf_B_2_2)		Node B 2	Pool 0
Disk shelf 14 (shelf_B_2_4)			
Disk shelf 15 (shelf_A_2_2)		Node A 2	Pool 1
Disk shelf 16 (shelf_A_2_4)			

Steps

1. Issue the following command to confirm the shelf assignments:

```
disk show -v
```

2. If necessary, 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.
3. Issue the following command to show the disk shelf IDs and bays for each disk:

```
storage show disk -x
```

Assigning disk ownership shelf by shelf

If the MetroCluster nodes do not have the disks correctly assigned, you must assign disks to each of the nodes in the MetroCluster configuration to create a configuration in which each node has the same number of disks in its local and remote disk pools. Each node can have a different number of disks or shelves.

Before you begin

The storage controllers must be in Maintenance mode.

About this task

This task is only required if disks were not correctly assigned when received from the factory.

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.

Note: Pool 0 always contains the disks that are found 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. If you have not done so, boot each system into Maintenance mode.
2. Assign the disk shelves to the nodes located at the first site (site A):

Disk shelves at the same site as the node are assigned to pool 0 and disk shelves located at the partner site are assigned to pool 1.

You should assign an equal number of shelves to each pool.

- a. On the first node, systematically assign the local disk shelves to pool 0 and the remote disk shelves to pool 1:

```
disk assign -shelf local-switch-name:shelf-name.port -p pool
```

Example

If storage controller Controller_A_1 has four shelves, you issue the following commands:

```
*> disk assign -shelf FC_switch_A_1:1-4.shelf1 -p 0
*> disk assign -shelf FC_switch_A_1:1-4.shelf2 -p 0

*> disk assign -shelf FC_switch_B_1:1-4.shelf1 -p 1
*> disk assign -shelf FC_switch_B_1:1-4.shelf2 -p 1
```

- b. Repeat the process for the second node at the local site, systematically assigning the local disk shelves to pool 0 and the remote disk shelves to pool 1:

```
disk assign -shelf local-switch-name:shelf-name.port -p pool
```

Example

If storage controller Controller_A_1 has four shelves, you issue the following commands:

```
*> disk assign -shelf FC_switch_A_1:1-4.shelf3 -p 0
*> disk assign -shelf FC_switch_B_1:1-4.shelf4 -p 1

*> disk assign -shelf FC_switch_A_1:1-4.shelf3 -p 0
*> disk assign -shelf FC_switch_B_1:1-4.shelf4 -p 1
```

3. Assign the disk shelves to the nodes located at the second site (site B):

Disk shelves at the same site as the node are assigned to pool 0 and disk shelves located at the partner site are assigned to pool 1.

You should assign an equal number of shelves to each pool.

- a. On the first node at the remote site, systematically assign its local disk shelves to pool 0 and its remote disk shelves to pool 1:

```
disk assign -shelf local-switch-namesshelf-name -p pool
```

Example

If storage controller Controller_B_1 has four shelves, you issue the following commands:

```
*> disk assign -shelf FC_switch_B_1:1-5.shelf1 -p 0
*> disk assign -shelf FC_switch_B_1:1-5.shelf2 -p 0

*> disk assign -shelf FC_switch_A_1:1-5.shelf1 -p 1
*> disk assign -shelf FC_switch_A_1:1-5.shelf2 -p 1
```

- b. Repeat the process for the second node at the remote site, systematically assigning its local disk shelves to pool 0 and its remote disk shelves to pool 1:

```
disk assign -shelf shelf-name -p pool
```

Example

If storage controller Controller_B_2 has four shelves, you issue the following commands:

```
*> disk assign -shelf FC_switch_B_1:1-5.shelf3 -p 0
*> disk assign -shelf FC_switch_B_1:1-5.shelf4 -p 0

*> disk assign -shelf FC_switch_A_1:1-5.shelf3 -p 1
*> disk assign -shelf FC_switch_A_1:1-5.shelf4 -p 1
```

4. Confirm the shelf assignments:

```
storage show shelf
```

5. Exit Maintenance mode:

```
halt
```

6. Display the boot menu:

```
boot_ontap menu
```

7. On each node, select option 4 to initialize all disks.

Verifying the HA state of components is mcc in Maintenance mode

In a MetroCluster configuration, the HA state of the controller and chassis components must be set to `mcc` so they boot up properly. For systems received from the factory, this value is pre-configured and you do not need to verify it.

Before you begin

The system must be in Maintenance mode.

About this task

This task is not required on systems received from the factory.

This task must be performed on each node.

Steps

1. In Maintenance mode, display the HA state of the controller module and chassis:

```
ha-config show
```

The HA state should be `mcc` for all components.

2. If the displayed system state of the controller is not `mcc`, set the HA state for the controller module to `mcc`:

```
ha-config modify controller mcc
```
3. If the displayed system state of the chassis is not `mcc`, set the HA state for the chassis to `mcc`:

```
ha-config modify chassis mcc
```
4. Boot the node to Data ONTAP:

```
boot_ontap
```
5. Repeat these steps on each node in the MetroCluster configuration.

Running System Setup to configure the nodes and clusters

After you boot up each node, you are prompted to run the System Setup tool to perform basic node and cluster configuration. After configuring the cluster, you return to the Data ONTAP CLI to create aggregates and create the MetroCluster configuration.

Before you begin

You must have cabled the MetroCluster configuration.

You must not have configured the Service Processor prior to performing this task.

About this task

New MetroCluster systems are pre-configured; you do not need to perform these steps. However, you should configure autosupport.

This task must be performed on both clusters in the MetroCluster configuration.

This procedure uses the System Setup tool. If desired, you can use the CLI cluster setup wizard.

[Clustered Data ONTAP 8.3 Software Setup Guide](#)

Steps

1. If you have not already done so, power up each node and let them boot up.
 If the system is in Maintenance mode, issue the `halt` command to exit Maintenance mode, and then issue the following command from the `LOADER` prompt:

```
boot_ontap
```

Example

The output should be similar to the following:

```
Welcome to node setup

You can enter the following commands at any time:
  "help" or "?" - if you want to have a question clarified,
  "back" - if you want to change previously answered questions, and
  "exit" or "quit" - if you want to quit the setup wizard.
                  Any changes you made before quitting will be saved.

To accept a default or omit a question, do not enter a value.
.
.
.
```

2. Enable autosupport by following the directions provided by the system.
3. Respond to the prompts to configure the node management interface.

Example

The prompts are similar to the following:

```
Enter the node management interface port: [e0M]:
Enter the node management interface IP address: 10.228.160.229
Enter the node management interface netmask: 225.225.252.0
Enter the node management interface default gateway: 10.228.160.1
```

4. Confirm that nodes are configured in high-availability mode:

```
storage failover show
```

If not, issue the following command on each node and reboot the node:

```
storage failover modify -mode ha -node localhost
```

5. Confirm that you have four ports configured as cluster interconnects:

```
network port show
```

Example

```
cluster_A::> network port show
```

Node	Port	IPspace	Broadcast	Domain	Link	MTU	Speed (Mbps) Admin/Oper
node_A_1							
	e0a	Cluster	Cluster		up	1500	auto/1000
	e0b	Cluster	Cluster		up	1500	auto/1000
	e0c	Default	Default		up	1500	auto/1000
	e0d	Default	Default		up	1500	auto/1000
	e0e	Default	Default		up	1500	auto/1000
	e0f	Default	Default		up	1500	auto/1000
	e0g	Default	Default		up	1500	auto/1000
node_A_2							
	e0a	Cluster	Cluster		up	1500	auto/1000
	e0b	Cluster	Cluster		up	1500	auto/1000

e0c	Default	Default	up	1500	auto/1000
e0d	Default	Default	up	1500	auto/1000
e0e	Default	Default	up	1500	auto/1000
e0f	Default	Default	up	1500	auto/1000
e0g	Default	Default	up	1500	auto/1000
14 entries were displayed.					

6. If you are creating a two-node switchless cluster (a cluster without cluster interconnect switches), enable the switchless-cluster networking mode:
 - a. Issue the following command at either node's prompt to change to the advanced privilege level:


```
set -privilege advanced
```

You can respond **y** when prompted to continue into advanced mode. The advanced mode prompt appears (*>).
 - b. Enable switchless-cluster mode:


```
network options switchless-cluster modify -enabled true
```
 - c. Return to the admin privilege level:


```
set -privilege admin
```
7. Launch the System Setup tool as directed by the information that appears on the system console after the initial bootup.
8. Use the System Setup tool to configure each node and create the cluster, but do not create aggregates.

Note: You will create mirrored aggregates in later tasks.

After you finish

Return to the Data ONTAP command line interface and complete the MetroCluster configuration by performing the tasks that follow.

Configuring the clusters into a MetroCluster configuration

You must mirror the root aggregates, create a mirrored data aggregate, and then issue the command to implement the MetroCluster operations.

Peering the clusters

The clusters in the MetroCluster configuration must be in a peer relationship so they can communicate with each other and perform the data mirroring essential to MetroCluster disaster recovery.

Steps

- 1. [Configuring intercluster LIFs to use dedicated intercluster ports](#) on page 122
- 2. [Configuring intercluster LIFs to share data ports](#) on page 126
- 3. [Creating the cluster peer relationship](#) on page 128

Configuring intercluster LIFs to use dedicated intercluster ports

Configuring intercluster LIFs to use dedicated data ports allows greater bandwidth than using shared data ports on your intercluster networks for cluster peer relationships.

About this task

Creating intercluster LIFs that use dedicated ports involves creating a failover group for the dedicated ports and assigning LIFs to those ports. In this procedure, a two-node cluster exists in which each node has two data ports that you have added, e0e and e0f. These ports are ones you will dedicate for intercluster replication and currently are in the default IPspace. These ports will be grouped together as targets for the intercluster LIFs you are configuring. You must configure intercluster LIFs on the peer cluster before you can create cluster peer relationships. In your own environment, you would replace the ports, networks, IP addresses, subnet masks, and subnets with those specific to your environment.

Steps

- 1. List the ports in the cluster by using `network port show` command.

Example

```
cluster01::> network port show
```

Node	Port	IPspace	Broadcast Domain	Link	MTU	Speed (Mbps) Admin/Oper

cluster01-01	e0a	Cluster	Cluster	up	1500	auto/1000
	e0b	Cluster	Cluster	up	1500	auto/1000
	e0c	Default	Default	up	1500	auto/1000

	e0d	Default	Default	up	1500	auto/1000
	e0e	Default	Default	up	1500	auto/1000
	e0f	Default	Default	up	1500	auto/1000
cluster01-02						
	e0a	Cluster	Cluster	up	1500	auto/1000
	e0b	Cluster	Cluster	up	1500	auto/1000
	e0c	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000
	e0e	Default	Default	up	1500	auto/1000
	e0f	Default	Default	up	1500	auto/1000

- Determine whether any of the LIFs are using ports that are dedicated for replication by using the `network interface show` command.

Example

Ports e0e and e0f do not appear in the following output; therefore, they do not have any LIFs located on them:

```
cluster01::> network interface show -fields home-port,curr-port
vserver lif             home-port curr-port
-----
Cluster cluster01-01_clus1 e0a      e0a
Cluster cluster01-01_clus2 e0b      e0b
Cluster cluster01-02_clus1 e0a      e0a
Cluster cluster01-02_clus2 e0b      e0b
cluster01
  cluster_mgmt          e0c      e0c
cluster01
  cluster01-01_mgmt1    e0c      e0c
cluster01
  cluster01-02_mgmt1    e0c      e0c
```

- Group the ports that you will use for the intercluster LIFs by using the `network interface failover-groups create` command.

Example

```
cluster01::> network interface failover-groups create -vserver cluster01
-failover-group intercluster01 -targets cluster01-01:e0e,cluster01-01:e0f,
cluster01-02:e0e,cluster01-02:e0f
```

- Display the failover-group that you created by using the `network interface failover-groups show` command.

Example

```
cluster01::> network interface failover-groups show
Vserver      Group      Failover
-----
Cluster
  Cluster
    cluster01-01:e0a, cluster01-01:e0b,
    cluster01-02:e0a, cluster01-02:e0b
cluster01
  Default
    cluster01-01:e0c, cluster01-01:e0d,
```

	cluster01-02:e0c, cluster01-02:e0d, cluster01-01:e0e, cluster01-01:e0f cluster01-02:e0e, cluster01-02:e0f
intercluster01	cluster01-01:e0e, cluster01-01:e0f cluster01-02:e0e, cluster01-02:e0f

5. Create an intercluster LIF on the admin SVM cluster01 by using the `network interface create` command.

Example

This example uses the LIF naming convention `adminSVMname_icl#` for the intercluster LIF:

```
cluster01::> network interface create -vserver cluster01 -lif cluster01_icl01 -role
intercluster -home-node cluster01-01 -home-port e0e
-address 192.168.1.201 -netmask 255.255.255.0 -failover-group intercluster01

cluster01::> network interface create -vserver cluster01 -lif cluster01_icl02 -role
intercluster -home-node cluster01-02 -home-port e0e
-address 192.168.1.202 -netmask 255.255.255.0 -failover-group intercluster01
```

6. Verify that the intercluster LIFs were created properly by using the `network interface show` command.

Example

```
cluster01::> network interface show
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
Cluster						
	cluster01-01_clus_1	up/up	192.168.0.xxx/24	cluster01-01	e0a	true
	cluster01-01_clus_2	up/up	192.168.0.xxx/24	cluster01-01	e0b	true
	cluster01-02_clus_1	up/up	192.168.0.xxx/24	cluster01-01	e0a	true
	cluster01-02_clus_2	up/up	192.168.0.xxx/24	cluster01-01	e0b	true
cluster01	cluster_mgmt	up/up	192.168.0.xxx/24	cluster01-01	e0c	true
	cluster01_icl01	up/up	192.168.1.201/24	cluster01-01	e0e	true
	cluster01_icl02	up/up	192.168.1.202/24	cluster01-02	e0e	true
	cluster01-01_mgmt1	up/up	192.168.0.xxx/24	cluster01-01	e0c	true
	cluster01-02_mgmt1	up/up	192.168.0.xxx/24	cluster01-02	e0c	true

7. Verify that the intercluster LIFs are configured for redundancy by using the `network interface show` command with the `-role intercluster` and `-failover` parameters.

Example

The LIFs in this example are assigned the e0e home port on each node. If the e0e port fails, the LIF can fail over to the e0f port.

```
cluster01::> network interface show -role intercluster -failover
```

Vserver	Logical Interface	Home Node:Port	Failover Policy	Failover Group
cluster01-01	cluster01-01_icl01	cluster01-01:e0e	local-only	intercluster01
		Failover Targets:	cluster01-01:e0e, cluster01-01:e0f	
	cluster01-01_icl02	cluster01-02:e0e	local-only	intercluster01
		Failover Targets:	cluster01-02:e0e, cluster01-02:e0f	

8. Display the routes in the cluster by using the `network route show` command to determine whether intercluster routes are available or you must create them.

Creating a route is required only if the intercluster addresses in both clusters are not on the same subnet and a specific route is needed for communication between the clusters.

Example

In this example, no intercluster routes are available:

```
cluster01::> network route show
```

Vserver	Destination	Gateway	Metric
Cluster			
	0.0.0.0/0	192.168.0.1	20
cluster01	0.0.0.0/0	192.168.0.1	10

9. If communication between intercluster LIFs in different clusters requires routing, create an intercluster route by using the `network route create` command.

Example

In this example, 192.168.1.1 is the gateway address for the 192.168.1.0/24 network. If the destination is specified as 0.0.0.0/0, then it becomes a default route for the intercluster network.

```
cluster01::> network route create -vserver cluster01
-destination 0.0.0.0/0 -gateway 192.168.1.1 -metric 40
```

10. Verify that you created the routes correctly by using the `network route show` command.

Example

```
cluster01::> network route show
```

Vserver	Destination	Gateway	Metric
Cluster			
	0.0.0.0/0	192.168.0.1	20
cluster01	0.0.0.0/0	192.168.0.1	10
	0.0.0.0/0	192.168.1.1	40

11. Repeat these steps to configure intercluster networking in the peer cluster.

12. Verify that the ports have access to the proper subnets, VLANs, and so on.

Dedicating ports for replication in one cluster does not require dedicating ports in all clusters; one cluster might use dedicated ports, while the other cluster shares data ports for intercluster replication.

Configuring intercluster LIFs to share data ports

Configuring intercluster LIFs to share data ports enables you to use existing data ports to create intercluster networks for cluster peer relationships. Sharing data ports reduces the number of ports you might need for intercluster networking.

About this task

Creating intercluster LIFs that share data ports involves assigning LIFs to existing data ports. In this procedure, a two-node cluster exists in which each node has two data ports, e0c and e0d, and these data ports are in the default IPspace. These are the two data ports that are shared for intercluster replication. You must configure intercluster LIFs on the peer cluster before you can create cluster peer relationships. In your own environment, you replace the ports, networks, IP addresses, subnet masks, and subnets with those specific to your environment.

Steps

1. List the ports in the cluster by using the `network port show` command:

Example

```
cluster01::> network port show
```

Node	Port	IPspace	Broadcast	Domain	Link	MTU	Speed (Mbps) Admin/Oper

cluster01-01							
	e0a	Cluster	Cluster		up	1500	auto/1000
	e0b	Cluster	Cluster		up	1500	auto/1000
	e0c	Default	Default		up	1500	auto/1000
	e0d	Default	Default		up	1500	auto/1000
cluster01-02							
	e0a	Cluster	Cluster		up	1500	auto/1000
	e0b	Cluster	Cluster		up	1500	auto/1000
	e0c	Default	Default		up	1500	auto/1000
	e0d	Default	Default		up	1500	auto/1000

2. Create an intercluster LIF on the admin SVM cluster01 by using the `network interface create` command.

Example

This example uses the LIF naming convention `adminSVMname_icl#` for the intercluster LIF:

```
cluster01::> network interface create -vserver cluster01 -lif cluster01_icl01 -role
intercluster
-home-node cluster01-01 -home-port e0c -address 192.168.1.201 -netmask 255.255.255.0
```

```
cluster01::> network interface create -vserver cluster01 -lif cluster01_icl02 -role
intercluster
-home-node cluster01-02 -home-port e0c -address 192.168.1.202 -netmask 255.255.255.0
```

3. Verify that the intercluster LIFs were created properly by using the `network interface show` command with the `-role intercluster` parameter:

Example

```
cluster01::> network interface show -role intercluster
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
cluster01	cluster01_icl01	up/up	192.168.1.201/24	cluster01-01	e0c	true
	cluster01_icl02	up/up	192.168.1.202/24	cluster01-02	e0c	true

4. Verify that the intercluster LIFs are configured to be redundant by using the `network interface show` command with the `-role intercluster` and `-failover` parameters.

Example

The LIFs in this example are assigned the e0c port on each node. If the e0c port fails, the LIF can fail over to the e0d port.

```
cluster01::> network interface show -role intercluster -failover
```

Vserver	Logical Interface	Home Node:Port	Failover Policy	Failover Group
cluster01	cluster01_icl01	cluster01-01:e0c	local-only	192.168.1.201/24
		Failover Targets:	cluster01-01:e0c, cluster01-01:e0d	
	cluster01_icl02	cluster01-02:e0c	local-only	192.168.1.201/24
		Failover Targets:	cluster01-02:e0c, cluster01-02:e0d	

5. Display the routes in the cluster by using the `network route show` command to determine whether intercluster routes are available or you must create them.

Creating a route is required only if the intercluster addresses in both clusters are not on the same subnet and a specific route is needed for communication between the clusters.

Example

In this example, no intercluster routes are available:

```
cluster01::> network route show
```

Vserver	Destination	Gateway	Metric
Cluster			
	0.0.0.0/0	192.168.0.1	20
cluster01			
	0.0.0.0/0	192.168.0.1	10

6. If communication between intercluster LIFs in different clusters requires routing, create an intercluster route by using the `network route create` command.

Example

In this example, 192.168.1.1 is the gateway address for the 192.168.1.0/24 network. If the destination is specified as 0.0.0.0/0, then it becomes a default route for the intercluster network.

```
cluster01::> network route create -vserver cluster01
-destination 0.0.0.0/0 -gateway 192.168.1.1 -metric 40
```

7. Verify that you created the routes correctly by using the `network route show` command.

Example

```
cluster01::> network route show
```

Vserver	Destination	Gateway	Metric
Cluster			
	0.0.0.0/0	192.168.0.1	20
cluster01			
	0.0.0.0/0	192.168.0.1	10
	0.0.0.0/0	192.168.1.1	40

8. Repeat these steps on the cluster to which you want to connect.

Creating the cluster peer relationship

You create the cluster peer relationship using a set of intercluster logical interfaces to make information about one cluster available to the other cluster for use in cluster peering applications.

Before you begin

- Intercluster LIFs should be created on all of the nodes of both clusters you want to peer
- You should ensure that the intercluster LIFs of the clusters can route to each other.
- If there are different administrators for each cluster, the passphrase used to authenticate the cluster peer relationship should be agreed upon.

Steps

1. Create the cluster peer relationship on each cluster by using the `cluster peer create` command.

The passphrase that you use is not displayed as you type it.

Example

In the following example, cluster01 is peered with a remote cluster named cluster02. Cluster01 is a two-node cluster that has one intercluster LIF per node. The IP addresses of the intercluster LIFs created in cluster01 are 192.168.2.201 and 192.168.2.202. Similarly, cluster02 is a two-node cluster that has one intercluster LIF per node. The IP addresses of the intercluster LIFs created in cluster02 are 192.168.2.203 and 192.168.2.204. These IP addresses are used to create the cluster peer relationship.

```
cluster01::> cluster peer create -peer-addr
192.168.2.203,192.168.2.204
Please type the passphrase:
Please type the passphrase again:
```

```
cluster02::> cluster peer create -peer-addr
192.168.2.201,192.168.2.202
Please type the passphrase:
Please type the passphrase again:
```

If DNS is configured to resolve host names for the intercluster IP addresses, you can use host names in the `-peer-addr` option. It is not likely that intercluster IP addresses frequently change; however, using host names allows intercluster IP addresses to change without having to modify the cluster peer relationship.

2. Display the cluster peer relationship by using the `cluster peer show` command with the `-instance` parameter.

Displaying the cluster peer relationship verifies that the relationship was established successfully.

Example

```
cluster01::> cluster peer show -instance
Peer Cluster Name: cluster02
Remote Intercluster Addresses: 192.168.2.203,192.168.2.204
Availability: Available
Remote Cluster Name: cluster02
Active IP Addresses: 192.168.2.203,192.168.2.204
Cluster Serial Number: 1-80-000013
```

3. Preview the health of the nodes in the peer cluster by using the `cluster peer health show` command.

Previewing the health checks the connectivity and status of the nodes on the peer cluster.

Example

```
cluster01::> cluster peer health show
```

Node	cluster-Name	Node-Name	RDB-Health	Cluster-Health	Avail...

cluster01-01	cluster02	cluster02-01			
	Data: interface_reachable				
	ICMP: interface_reachable	true	true		true
		cluster02-02			
	Data: interface_reachable				
	ICMP: interface_reachable	true	true		true
cluster01-02	cluster02	cluster02-01			
	Data: interface_reachable				
	ICMP: interface_reachable	true	true		true
		cluster02-02			
	Data: interface_reachable				
	ICMP: interface_reachable	true	true		true

Mirroring the root aggregates

You must mirror the root aggregates to ensure data protection.

Steps

- 1. Mirror the root aggregate:

```
storage aggregate mirror aggr_ID
```

Example

The following command mirrors the root aggregate for controller_A_1:

```
controller_A_1::> storage aggregate mirror aggr0_controller_A_1
```

This creates an aggregate with a local plex located at the local MetroCluster site and a remote plex located at the remote MetroCluster site.

- 2. Repeat the previous steps for each node in the MetroCluster configuration.

Creating a mirrored data aggregate on each node

You must create a mirrored data aggregate on each node in the DR group.

Before you begin

- You should know what drives or array LUNs will be used in the new aggregate.
- If you have multiple drive types in your system (heterogeneous storage), you should understand how you can ensure that the correct drive type is selected.

About this task

- Drives and array LUNs are owned by a specific node; when you create an aggregate, all drives in that aggregate must be owned by the same node, which becomes the home node for that aggregate.
- Aggregate names should conform to the naming scheme you determined when you planned your MetroCluster configuration.
- The *Clustered Data ONTAP Data Protection Guide* contains more information about mirroring aggregates.

Steps

1. Display a list of available spares:

```
storage disk show -spare -owner node_name
```

2. Create the aggregate by using the `storage aggregate create -mirror true` command.

If you are logged in to the cluster on the cluster management interface, you can create an aggregate on any node in the cluster. To ensure that the aggregate is created on a specific node, use the `-node` parameter or specify drives that are owned by that node.

You can specify the following options:

- Aggregate's home node (that is, the node that owns the aggregate in normal operation)
- List of specific drives or array LUNs that are to be added to the aggregate
- Number of drives to include
- Checksum style to use for the aggregate
- Type of drives to use
- Size of drives to use
- Drive speed to use
- RAID type for RAID groups on the aggregate
- Maximum number of drives or array LUNs that can be included in a RAID group
- Whether drives with different RPM are allowed

For more information about these options, see the `storage aggregate create` man page.

Example

The following command creates a mirrored aggregate with 10 disks:

```
controller_A_1::> storage aggregate create aggr1_controller_A_1 -
diskcount 10 -node controller_A_1 -mirror true
[Job 15] Job is queued: Create aggr1_controller_A_1.
[Job 15] The job is starting.
[Job 15] Job succeeded: DONE
```

3. Verify the RAID group and drives of your new aggregate:

```
storage aggregate show-status -aggregate aggregate-name
```

Implementing the MetroCluster configuration

You must run the `metrocluster configure` command to start data protection in the MetroCluster configuration.

Before you begin

- There must be at least two nonroot mirrored data aggregates on each cluster, and all aggregates must be mirrored.

You can verify this with the `storage aggregate show` command.

- The ha-config state of the controllers and chassis must be `mcc`.
This state is pre-configured on systems that have shipped from the factory.

About this task

You issue the `metrocluster configure` command once, on any of the nodes, to enable the MetroCluster configuration. It does not need to be issued on each of the sites or nodes. It does not matter which node or site it is issued on.

Steps

1. Enter the `metrocluster configure` command in the following format: `metrocluster configure -node-name node-name`.

Example

The following command enables MetroCluster configuration on all nodes in the DR group that contains `controller_A_1`:

```
controller_A_1::*> metrocluster configure -node-name controller_A_1
[Job 121] Job succeeded: Configure is successful.
```

2. Enter the following command to check the networking status on site A:

```
network port show
```

Example

```
cluster_A::> network port show
```

Node	Port	IPspace	Broadcast Domain	Link	MTU	Speed (Mbps) Admin/Oper

controller_A_1						
	e0a	Cluster	Cluster	up	1500	auto/1000
	e0b	Cluster	Cluster	up	1500	auto/1000
	e0c	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000
	e0e	Default	Default	up	1500	auto/1000
	e0f	Default	Default	up	1500	auto/1000
	e0g	Default	Default	up	1500	auto/1000
controller_A_2						
	e0a	Cluster	Cluster	up	1500	auto/1000
	e0b	Cluster	Cluster	up	1500	auto/1000
	e0c	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000
	e0e	Default	Default	up	1500	auto/1000
	e0f	Default	Default	up	1500	auto/1000
	e0g	Default	Default	up	1500	auto/1000

14 entries were displayed.

3. Confirm the MetroCluster configuration from both sites in the MetroCluster configuration:

- a. Confirm the configuration from site A;

```
metrocluster show
```

Example

```
cluster_A::> metrocluster show
```

Cluster	Configuration	State	Mode

Local: cluster_A	configured		normal
Remote: cluster_B	configured		normal

- b. Confirm the configuration from site B;

```
metrocluster show
```

Example

```
cluster_B::> metrocluster show
```

Cluster	Configuration	State	Mode

Local: cluster_B	configured		normal
Remote: cluster_A	configured		normal

Configuring MetroCluster components for health monitoring

You must perform some special configuration steps before monitoring the components in the MetroCluster configuration.

Related information

Clustered Data ONTAP 8.3 System Administration Guide for Cluster Administrators

Configuring the MetroCluster FC switches for health monitoring

You must perform some special configuration steps to monitor the FC switches in the MetroCluster configuration.

Steps

1. Issue the following command on each MetroCluster node:

```
storage switch add -switch-ipaddress ipaddress
```

This command must be repeated on all four switches in the MetroCluster configuration.

Example

The following example shows the command to add a switch with IP 10.10.10.10:

```
controller_A_1::> storage switch add -switch-ipaddress 10.10.10.10
```

2. Verify that all switches are properly configured:

```
storage switch show
```

It may take up to 15 minutes to reflect all data due to the 15-minute polling interval.

Example

The following example shows the command given to verify the MetroCluster's FC switches are configured:

```
controller_A_1::> storage switch show
Fabric      Switch Name      Vendor  Model      Switch WWN      Status
-----
1000000533a9e7a6 brcd6505-fcs40  Brocade Brocade6505 1000000533a9e7a6 OK
1000000533a9e7a6 brcd6505-fcs42  Brocade Brocade6505 1000000533d3660a OK
1000000533ed94d1 brcd6510-fcs44  Brocade Brocade6510 1000000533eda031 OK
1000000533ed94d1 brcd6510-fcs45  Brocade Brocade6510 1000000533ed94d1 OK
4 entries were displayed.

controller_A_1::>
```

If the switch's worldwide name (WWN) is shown, the Data ONTAP health monitor is able to contact and monitor the FC switch.

Configuring FC-to-SAS bridges for health monitoring

You must perform some special configuration steps to monitor the FC-to-SAS bridges in the MetroCluster configuration.

Steps

1. Configure the FC-to-SAS bridges for monitoring by issuing the following command for each bridge on each storage controller:

```
storage bridge add -address ipaddress
```

This command must be repeated for all FC-to-SAS bridges in the MetroCluster configuration.

Example

The following example shows the command you must use to add an FC-to-SAS bridge with an IP address of 10.10.20.10:

```
controller_A_1::> storage bridge add -address 10.10.20.10
```

2. Verify that all FC-to-SAS bridges are properly configured:

```
storage bridge show
```

It might take as long as 15 minutes to reflect all data due to the polling interval.

Example

The following example shows that the FC-to-SAS bridges are configured:

```
controller_A_1::> storage bridge show
```

Bridge	Symbolic Name	Vendor	Model	Bridge WWN	Monitored	Status
ATTO_1	atto6500n-1	Atto	FibreBridge 6500N	WW11	true	ok
ATTO_2	atto6500n-2	Atto	FibreBridge 6500N	WW11	true	ok
ATTO_3	atto6500n-3	Atto	FibreBridge 6500N	WW11	true	ok
ATTO_4	atto6500n-4	Atto	FibreBridge 6500N	WW11	true	ok

```
controller_A_1::>
```

If the FC-to-SAS bridge's worldwide name (WWN) is shown, the Data ONTAP health monitor is able to contact and monitor the bridge.

Checking the MetroCluster configuration

You can check that the components and relationships in the MetroCluster configuration are working correctly. You should do a check after initial configuration and after making any changes to the

MetroCluster configuration. After you run the `metrocluster check run` command, you then display the results of the check with the `metrocluster check show` command.

About this task

If the `metrocluster check run` command is issued twice within a short time, on either or both clusters, a conflict can occur and the command might not collect all data. Subsequent `metrocluster check show` commands will not show the expected output.

Steps

1. Check the configuration:

```
metrocluster check run
```

Example

The following example shows the output for a healthy MetroCluster configuration:

```
controller_A_1::> metrocluster check run

Last Checked On: 9/24/2014 17:10:33

Component          Result
-----
nodes               ok
lifs                ok
config-replication ok
aggregates          ok
4 entries were displayed.

Command completed. Use the "metrocluster check show -instance"
command or sub-commands in "metrocluster check" directory for
detailed results.
To check if the nodes are ready to do a switchover or switchback
operation, run "metrocluster switchover -simulate" or "metrocluster
switchback
-simulate", respectively.
```

2. Display more detailed results from the most recent `metrocluster check run` command:

```
metrocluster check aggregate show
```

```
metrocluster check config-replication show
```

```
metrocluster check lif show
```

```
metrocluster check node show
```

The `metrocluster check show` commands show the results of the most recent `metrocluster check run` command. You should always run the `metrocluster check`

run command prior to using the `metrocluster check show` commands to ensure that the information displayed is current.

Example

The following example shows the `metrocluster check aggregate show` output for a healthy MetroCluster configuration:

```
controller_A_1:> metrocluster check aggregate show
Last Checked On: 8/5/2014 00:42:58
```

Node	Aggregate	Check	Result
controller_A_1	aggr0_controller_A_1_0	mirroring-status	ok
		disk-pool-allocation	ok
	controller_A_1_aggr1	mirroring-status	ok
		disk-pool-allocation	ok
controller_A_2	aggr0_controller_A_2	mirroring-status	ok
		disk-pool-allocation	ok
	controller_A_2_aggr1	mirroring-status	ok
		disk-pool-allocation	ok
controller_B_1	aggr0_controller_B_1_0	mirroring-status	ok
		disk-pool-allocation	ok
	controller_B_1_aggr1	mirroring-status	ok
		disk-pool-allocation	ok
controller_B_2	aggr0_controller_B_2	mirroring-status	ok
		disk-pool-allocation	ok
	controller_B_2_aggr1	mirroring-status	ok
		disk-pool-allocation	ok

16 entries were displayed.

Related information

[Clustered Data ONTAP 8.3 Physical Storage Management Guide](#)

[Clustered Data ONTAP 8.3 Network Management Guide](#)

[Clustered Data ONTAP 8.3 Data Protection Guide](#)

Checking for MetroCluster configuration errors with Config Advisor

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. You can deploy it 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. Go to [NetApp Downloads: Config Advisor](#).
2. After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

Verifying local HA operation

You should verify the operation of the local HA pairs in the MetroCluster configuration.

About this task

The examples in this task use standard naming conventions:

- cluster_A
 - controller_A_1
 - controller_A_2
- cluster_B
 - controller_B_1
 - controller_B_2

Steps

1. On cluster_A, perform a failover and giveback in both directions.
 - a. Confirm that storage failover is enabled:

```
storage failover show
```

Example

The output should indicate that takeover is possible for both nodes:

```
cluster_A::> storage failover show
Node           Partner           Takeover
Possible State Description
-----
controller_A_1 controller_A_2 true      Connected to controller_A_2
```

```
controller_A_2 controller_A_1 true      Connected to controller_A_1
2 entries were displayed.
```

- b. Takeover controller_A_2 from controller_A_1:

```
storage failover takeover controller_A_2
```

You can use the `storage failover show-takeover` command to monitor the progress of the takeover operation.

- c. Confirm that the takeover is complete:

```
storage failover show
```

Example

The output should indicate that controller_A_1 is in takeover state, meaning that it has taken over its HA partner:

```
cluster_A::> storage failover show
Node           Partner           Takeover
Possible State Description
-----
controller_A_1 controller_A_2 false      In takeover
controller_A_2 controller_A_1 -          Unknown
2 entries were displayed.
```

- d. Give back controller_A_2:

```
storage failover giveback controller_A_2
```

You can use the `storage failover show-giveback` command to monitor the progress of the giveback operation.

- e. Confirm that storage failover has returned to a normal state:

```
storage failover show
```

Example

The output should indicate that takeover is possible for both nodes:

```
cluster_A::> storage failover show
Node           Partner           Takeover
Possible State Description
-----
controller_A_1 controller_A_2 true        Connected to controller_A_2
```

```
controller_A_2 controller_A_1 true      Connected to controller_A_1
2 entries were displayed.
```

- f. Repeat the previous substeps, this time taking over controller_A_1 from controller_A_2.
2. Repeat the previous step on cluster_B.

Related information

[Clustered Data ONTAP 8.3 High-Availability Configuration Guide](#)

Verifying switchover, healing, and switchback

You should verify the switchover, healing, and switchback operations of the MetroCluster configuration.

Step

1. Use the procedures for negotiated switchover, healing, and switchback in the *[Clustered Data ONTAP 8.3 MetroCluster Management and Disaster Recovery Guide](#)*.

Protecting configuration backup files

You can provide additional protection for the cluster configuration backup files by specifying a remote URL (either HTTP or FTP) where the configuration backup files will be uploaded in addition to the default locations in the local cluster.

Step

1. Set the URL of the remote destination for the configuration backup files:

```
system configuration backup settings modify URL-of-destination
```

Related information

[Clustered Data ONTAP 8.3 System Administration Guide for Cluster Administrators](#)

Planning and installing a MetroCluster configuration with array LUNs

If you are using array LUNs in your MetroCluster configuration, you need to plan the installation and follow the specific procedures for such a configuration. Systems support either a mix of array LUNs and NetApp disk shelves or array LUNs only.

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. Installing a MetroCluster configuration involves connecting and configuring a number of devices, which might be done by different people. Therefore, the plan also helps you communicate with other people involved in the installation.

Requirements for a MetroCluster configuration with array LUNs

There are some unique requirements for setting up a MetroCluster configuration with array LUNs.

Requirements for Data ONTAP systems using array LUNs in a MetroCluster configuration

- The platforms must be identified as supported for MetroCluster configurations.
[NetApp Interoperability Matrix Tool](#)
- All the Data ONTAP systems in a MetroCluster configuration must be of the same model.
- Ensure that FC-VI adapters are installed into the appropriate slots for each Data ONTAP system depending on the platform model.
[NetApp Hardware Universe](#)
- Sharing of multiple FC initiator ports with a single target port is *not* supported in a MetroCluster configuration. Similarly, sharing of multiple target ports with a single FC initiator port is also *not* supported.

Requirements for storage arrays

- The storage arrays must be identified as supported for MetroCluster configurations.
[NetApp Interoperability Matrix Tool](#)
- The storage arrays in the MetroCluster configuration must be symmetric, which means the following:

- The two storage arrays must be from the same supported vendor family and have the same firmware version installed.
[FlexArray Virtualization Implementation Guide for NetApp E-Series Storage](#)
[FlexArray Virtualization Implementation Guide for Third-Party Storage](#)
- You must have two sets of array LUNs—one set for the aggregate on the local storage array and another set of LUNs at the remote storage array for the mirror of the aggregate.
The array LUNs must be of the same size for mirroring the aggregate.
- Disk types (for example, SATA, SSD, or SAS) used for mirrored storage must be the same on both storage arrays.
- The parameters for configuring storage arrays, such as raid type and tiering, should be the same across both the sites.

Requirements for FC switches

- The switches and switch firmware must be identified as supported for MetroCluster configurations.
[NetApp Interoperability Matrix Tool](#)
- 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 and path failures.
- Data ONTAP supports using up to four ISLs, depending on the configuration.

Note: For more information about basic switch configuration, ISL settings, and FC-VI configurations, see *[Configuring the Cisco or Brocade FC switches manually](#)* on page 51.

Zoning requirements

- A single-initiator to single-target zoning scheme must be followed for MetroCluster configurations.
Single-initiator to single-target zoning limits each zone to a single FC initiator port and a single target port.
- FC-VI ports must be zoned end-to-end across the fabric.
- Sharing of multiple initiator ports with a single target port is not supported. Similarly, sharing of multiple target ports with a single initiator port is also not supported.

SyncMirror requirements

- SyncMirror is required for a MetroCluster configuration.
- Two separate storage arrays are required for the mirrored storage.

- 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).

Clustered Data ONTAP 8.3 Data Protection Guide

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 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. Setting up the storage array to present array LUNs to Data ONTAP and configuring the parameters required for a storage array to work with Data ONTAP.
 - *FlexArray Virtualization Implementation Guide for Third-Party Storage*
 - *FlexArray Virtualization Implementation Guide for NetApp E-Series Storage*
 - *FlexArray Virtualization Installation Requirements and Reference Guide*
2. Installing the FC-VI adapter on each Data ONTAP system, if it is not installed already.
3. Connecting the Data ONTAP systems at both the sites to the fabric.
4. Connecting the ISLs between the switches.
5. Connecting the storage array to the fabric.
6. Connecting the Data ONTAP systems to the cluster interconnect switches.
7. Configuring zoning.
8. Setting the HA state of the controller and chassis components in all the Data ONTAP systems to mcc.

Verifying the HA state of components is mcc in Maintenance mode on page 118
9. Assigning array LUNs to specific Data ONTAP systems.
10. Installing clustered Data ONTAP.
 - *FlexArray Virtualization Installation Requirements and Reference Guide*
 - *Clustered Data ONTAP 8.3 Software Setup Guide*
11. Setting up the Data ONTAP systems for the MetroCluster configuration.
12. Configuring various Data ONTAP features.

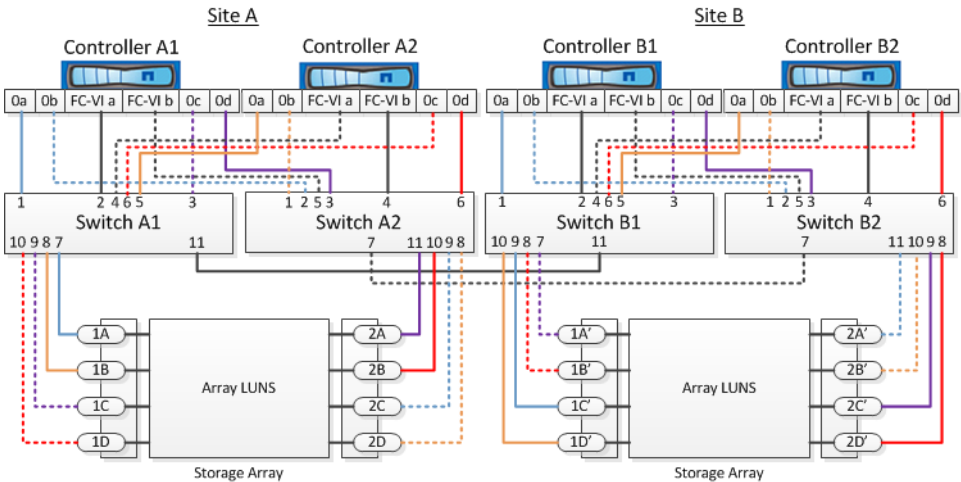
NetApp Documentation: Data ONTAP 8 (current releases)

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.

About this task

The following *example* illustration shows the connections between devices in a MetroCluster configuration that uses array LUNs.



Steps

1. Connect the FC-VI ports a and b from each controller to the alternate FC switch.

Perform the controller-to-switch connections at both Site A and Site B.

The following table lists the connections between the FC-VI ports and the FC switches for the configuration shown in the example MetroCluster illustration:

FC-VI Ports	FC Switch Ports
Site A:	
Controller A1: FC-VI Port a	Switch A1, Port 2
Controller A1: FC-VI Port b	Switch A2, Port 5
Controller A2: FC-VI Port a	Switch A1, Port 4

FC-VI Ports	FC Switch Ports
Controller A2: FC-VI Port b	Switch A2, Port 4
Site B:	
Controller B1: FC-VI Port a	Switch B1, Port 2
Controller B1: FC-VI Port b	Switch B2, Port 5
Controller B2: FC-VI Port a	Switch B1, Port 4
Controller B2: FC-VI Port b	Switch B2, Port 4

- Form ISL connections from each FC switch at Site A to a corresponding switch at Site B.

In the example MetroCluster configuration, FC switch A1 at Site A is connected to FC switch B1 at Site B to form a switch fabric. Similarly, switches A2 and B2 form another switch fabric.

- Connect the FC initiator ports from each controller to the FC switches.

For every controller at a site, ensure that at least two initiator ports are connected to the corresponding switch and the other two ports are connected to the alternate switch.

The following table lists the interconnections between the FC initiator ports and the switches in the example MetroCluster configuration:

FC Initiator Ports	FC Switch Ports
Site A:	
Controller A1: Port 0a	Switch A1, Port 1
Controller A1: Port 0b	Switch A2, Port 2
Controller A1: Port 0c	Switch A1, Port 3
Controller A1: Port 0d	Switch A2, Port 3
Controller A2: Port 0a	Switch A1: Port 5
Controller A2: Port 0b	Switch A2, Port 1
Controller A2: Port 0c	Switch A1, Port 6
Controller A2: Port 0d	Switch A2, Port 6
Site B:	
Controller B1: Port 0a	Switch B1, Port 1
Controller B1: Port 0b	Switch B2, Port 2
Controller B1: Port 0c	Switch B1, Port 3
Controller B1: Port 0d	Switch B2, Port 3

FC Initiator Ports	FC Switch Ports
Controller B2: Port 0a	Switch B1, Port 5
Controller B2: Port 0b	Switch B2, Port 1
Controller B2: Port 0c	Switch B1, Port 6
Controller B2: Port 0d	Switch B2, Port 6

4. Connect the ports of the storage arrays to the FC switch ports.

Perform the storage array-to-switch connections at Site A and Site B.

The following table lists the connections between the storage array ports and the FC switches for the example MetroCluster configuration:

Storage Array Ports	FC Switch Ports
Site A:	
Storage array 1: Port 1A	Switch A1, Port 7
Storage array 1: Port 1B	Switch A1, Port 8
Storage array 1: Port 1C	Switch A1, Port 9
Storage array 1: Port 1D	Switch A1, Port 10
Storage array 1: Port 2A	Switch A2, Port 11
Storage array 1: Port 2B	Switch A2, Port 10
Storage array 1: Port 2C	Switch A2, Port 9
Storage array 1: Port 2D	Switch A2, Port 8
Site B:	
Storage array 2: Port 1A'	Switch B1, Port 7
Storage array 2: Port 1B'	Switch B1, Port 8
Storage array 2: Port 1C'	Switch B1, Port 9
Storage array 2: Port 1D'	Switch B1, Port 10
Storage array 2: Port 2A'	Switch B2, Port 11
Storage array 2: Port 2B'	Switch B2, Port 10
Storage array 2: Port 2C'	Switch B2, Port 9
Storage array 2: Port 2D'	Switch B2, Port 8

5. Configure the switch zoning according to the zone plan.

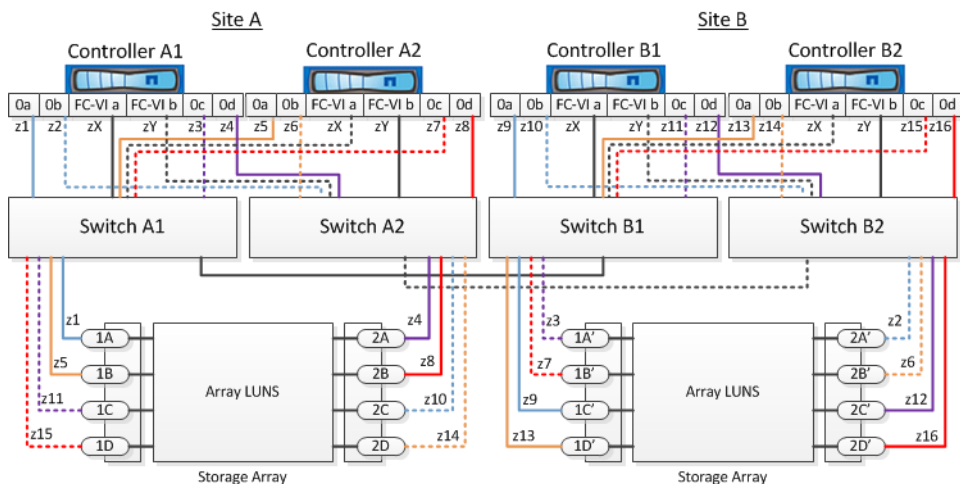
Switch zoning for a MetroCluster configuration with array LUNs

Switch zoning defines paths between connected nodes. Configuring the zoning enables you to define which array LUNs can be viewed by a specific Data ONTAP system.

When using switch zoning for a MetroCluster configuration with array LUNs, you must ensure that it meets these requirements:

- A single-initiator to single-target zoning scheme must be followed for MetroCluster configurations.
Single-initiator to single-target zoning limits each zone to a single FC initiator port and a single target port.
- FC-VI ports must be zoned end-to-end across the fabric.
- Sharing of multiple initiator ports with a single target port is not supported. Similarly, sharing of multiple target ports with a single initiator port is also not supported.

You can use the following *example* as a reference when determining zoning for a MetroCluster configuration with array LUNs. The example shows single-initiator zoning for a MetroCluster configuration. The lines in the following example represent zones rather than connections; each line is labeled with its zone number:



In the example illustration, 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 Data ONTAP system has two paths to array LUNs. The ports on the storage array are redundant, and are configured as follows:

- Storage array at Site A:
 - Ports 1A and 2A are a redundant port pair.

- Ports 1B and 2B are a redundant port pair.
- Storage array at Site B:
 - Ports 1C' and 2C' are a redundant port pair.
 - Ports 1D' and 2D' are a redundant port pair.

The redundant port pairs on each storage array form alternate paths. Therefore, both the ports of the port pairs can access the LUNs on the respective storage arrays.

The following table shows the zones for this example:

Zone	Data ONTAP controller and FC initiator port	Storage array port
Switch A1		
z1	Controller A1: Port 0a	Port 1A
z2	Controller A1: Port 0b	Port 2A'
z3	Controller A1: Port 0c	Port 1A'
z4	Controller A1: Port 0d	Port 2A
Switch A2		
z5	Controller A2: Port 0a	Port 1B
z6	Controller A2: Port 0b	Port 2B'
z7	Controller A2: Port 0c	Port 1B'
z8	Controller A2: Port 0d	Port 2B
Switch B1		
z9	Controller A3: Port 0a	Port 1C'
z10	Controller A3: Port 0b	Port 2C
z11	Controller A3: Port 0c	Port 1C
z12	Controller A3: Port 0d	Port 2C'
Switch B2		
z13	Controller A4: Port 0a	Port 1D'
z14	Controller A4: Port 0b	Port 2D
z15	Controller A4: Port 0c	Port 1D
z16	Controller A4: Port 0d	Port 2D'

The following table shows the zones for the FC-VI connections at Site A and Site B:

Zone	Data ONTAP controller and FC initiator port	Switch
Site A		
zX	Controller A1: Port FC-VI a	Switch A1
zY	Controller A1: Port FC-VI b	Switch A2
zX	Controller A2: Port FC-VI a	Switch A1
zY	Controller A2: Port FC-VI b	Switch A2
Site B		
zX	Controller B1: Port FC-VI a	Switch B1
zY	Controller B1: Port FC-VI b	Switch B2
zX	Controller B2: Port FC-VI a	Switch B1
zY	Controller B2: Port FC-VI b	Switch B2

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 must also configure any required Data ONTAP feature.

Steps

1. Set the HA state of the controller and chassis components in all the Data ONTAP systems to mcc.
[Verifying the HA state of components is mcc in Maintenance mode](#) on page 118
2. Assigning array LUNs to specific Data ONTAP systems.
3. Install clustered Data ONTAP on the Data ONTAP systems connected to array LUNs.
 - [FlexArray Virtualization Installation Requirements and Reference Guide](#)
 - [Clustered Data ONTAP 8.3 Software Setup Guide](#)
4. Set up a peering relationship between the clusters.
[Peering the clusters](#) on page 122
5. Mirror the root aggregate for each Data ONTAP node in the MetroCluster configuration.

[Mirroring the root aggregates](#) on page 130

6. Create a mirrored data aggregate on each Data ONTAP node in the MetroCluster configuration.

[Creating a mirrored data aggregate on each node](#) on page 130

7. Implement the MetroCluster configuration by using the `metrocluster configure` command.

[Implementing the MetroCluster configuration](#) on page 132

8. Verify if the components and the mirroring relationships in the MetroCluster configuration are working correctly.

[Checking the MetroCluster configuration](#) on page 135

9. Configure other Data ONTAP features as desired.

[NetApp Documentation: Data ONTAP 8 \(current releases\)](#)

Implementing a MetroCluster configuration with both disks and array LUNs

A MetroCluster configuration can consist of both disks and array LUNs if the Data ONTAP systems used in the configuration have the capability of attaching to storage arrays.

[NetApp Interoperability Matrix Tool](#)

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.

Considerations for implementing a MetroCluster configuration with disks and array LUNs

When planning your MetroCluster configuration for use with disks and array LUNs, you must consider various factors, such as the order of setting up access to storage, root aggregate location, and the usage of FC initiator ports, switches, and FC-to-SAS bridges.

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 (disks or array LUNs) first. You must 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
Location of the root aggregate	<ul style="list-style-type: none"> If you are setting up a <i>new</i> MetroCluster deployment with both disks and array LUNs, you must create the root aggregate on native disks. When doing this, ensure that <i>at least one</i> disk shelf (with 24 disk drives) is set up at each of the sites. If you are adding native disks to an <i>existing</i> MetroCluster configuration that uses array LUNs, the root aggregate can remain on an array LUN.
Using switches and FC-to-SAS bridges	<p>FC-to-SAS bridges are required to connect the Data ONTAP systems with the disk shelves through the switches.</p> <p>You must use the same switches to connect to the storage arrays and the FC-to-SAS bridges.</p>
Using FC initiator ports	<p>The FC initiator ports used to connect to an FC-to-SAS bridge must be different from the initiator ports used to connect to the switches, which connect to the storage arrays.</p> <p>A minimum of eight FC initiator ports is required to connect a Data ONTAP system to both disks and array LUNs.</p>

Related concepts

[Switch zoning for a MetroCluster configuration with array LUNs](#) on page 147

Related tasks

[Configuring the Cisco or Brocade FC switches manually](#) on page 51

[Installing FC-to-SAS bridges and SAS disk shelves](#) on page 41

Related information

[NetApp Hardware Universe](#)

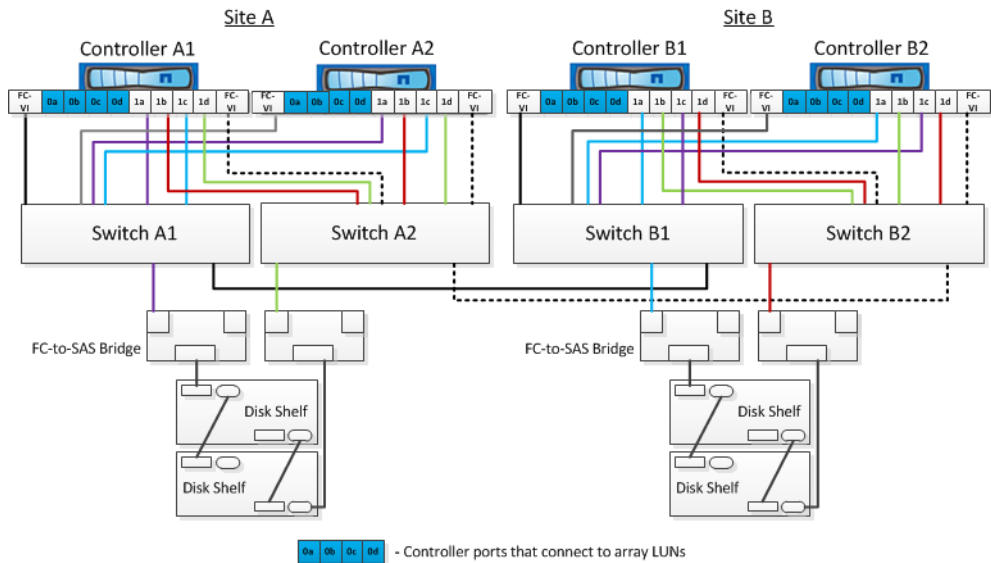
Example of a MetroCluster configuration with disks and array LUNs

In a MetroCluster configuration with native disks and array LUNs, you need FC-to-SAS bridges to connect the Data ONTAP systems with the disk shelves through the FC switches. You can connect array LUNs through the FC switches to the Data ONTAP systems.

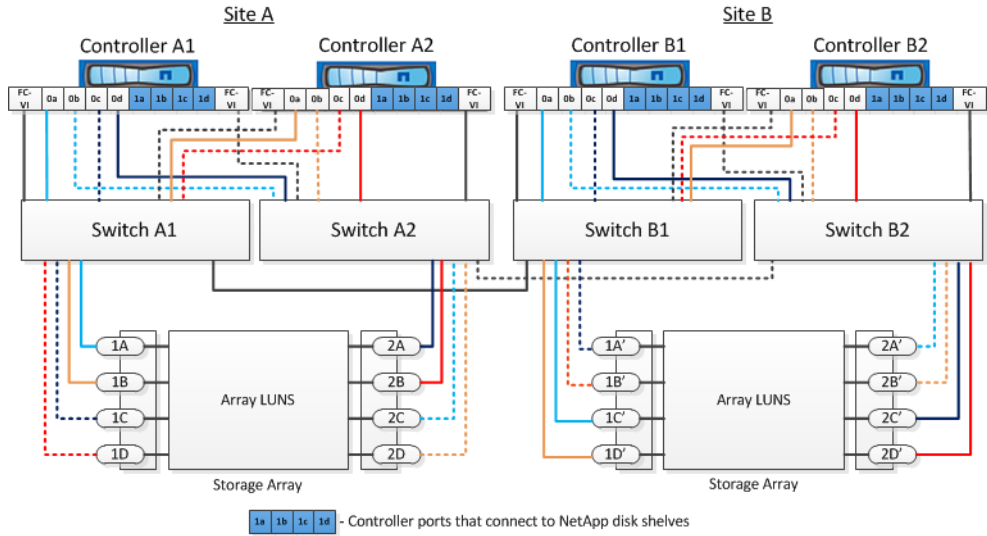
A minimum of eight FC initiator ports is required for a Data ONTAP system to connect to both native disks and array LUNs.

The following diagrams represent *examples* of a MetroCluster configuration with disks and array LUNs. They both represent the same MetroCluster configuration; the representations for disks and array LUNs are separated only for simplification.

In the following configuration diagram that displays the connectivity between Data ONTAP systems and disks, the initiator ports 1a through 1d are used for connectivity with disks through the FC-to-SAS bridges.



In the following configuration diagram that displays the connectivity between Data ONTAP systems and array LUNs, the initiator ports 0a through 0d are used for connectivity with array LUNs because ports 1a through 1d are used for connectivity with disks.



Using the OnCommand management tools for further configuration and monitoring

The OnCommand management tools can be used for GUI management of the clusters and monitoring of the configuration.

Each node has OnCommand System Manager pre-installed. To load System Manager, enter the cluster management LIF address as the URL in a web browser that has connectivity to the node.

You can also use OnCommand Unified Manager and OnCommand Performance Manager to monitor the MetroCluster configuration.

Related information

[NetApp Documentation: OnCommand Unified Manager Core Package \(current releases\)](#)

[NetApp Documentation: OnCommand System Manager \(current releases\)](#)

Requirements and limitations when using Data ONTAP in a MetroCluster configuration

When using Data ONTAP in a MetroCluster configuration, you should be aware of certain requirements and limitations when configuring Data ONTAP features.

- Both sites should be licensed for the same site-licensed features.
- All nodes should be licensed for the same node-locked features.
- Infinite Volumes are not supported in a MetroCluster configuration.

Job schedules in a MetroCluster configuration

Because job schedules are not replicated between the MetroCluster clusters, if you create, modify, or delete a job schedule, you must perform the same operation on the partner cluster. This ensures that job schedules on both clusters are identical.

Related information

[Clustered Data ONTAP 8.3 System Administration Guide for Cluster Administrators](#)

Cluster peering from the MetroCluster sites to a third cluster

Because peering configuration is not replicated, if you peer one of the clusters in the MetroCluster configuration to a third cluster outside of that configuration, you must also configure the peering on the partner MetroCluster cluster. This ensures peering can be maintained in the event of a switchover.

The non-MetroCluster cluster must be running Data ONTAP 8.3 or peering will be lost in the event of a switchover.

Volume creation on a root aggregate

The system will not allow the creation of new volumes on the root aggregate (aggregates with an HA policy of `CFO`) of a node in a MetroCluster configuration.

Networking and LIF creation guidelines for MetroCluster configurations

You should be aware of how LIFs are created and replicated within the MetroCluster configuration, and you must also know about the requirement for consistency so you can make proper decisions when configuring your network .

IPspace configuration

IPspace names must match between the two sites.

IPspace objects must be manually replicated to the partner cluster. Any SVMs created and assigned to an IPspace before the IPspace is replicated will not be replicated to the partner cluster.

IPv6 configuration

If IPv6 is configured on one site, it must be configured on the other site.

LIF creation

You can confirm the successful creation of a LIF in a MetroCluster configuration by running the `metrocluster check lif show`. If there are issues, you can use the `metrocluster check lif repair-placement` command.

Duplicate LIFs

You should not create duplicate LIFs (multiple LIFs with the same IP address) within the same IPspace.

Intercluster LIFs

intercluster LIFs are limited to the default IPspace, owned by the admin SVM.

Replication of LIFs to the partner cluster

When you create a LIF on a cluster in a MetroCluster configuration, that LIF is replicated on the partner cluster. The system must meet the following conditions to place the replicated LIF on the partner cluster:

1. DR partner availability

The system attempts to place the replicated LIF on the DR partner of the node on which it was created.

2. Connectivity

- For IP or iSCSI LIFs, the system places them on a reachable subnet.
- For FCP LIFs, the system attempts to place them on a reachable FC fabric.

3. Port attributes

The system attempts to place the LIF on a port with the desired VLAN, adapter type, and speed attributes.

An EMS message is displayed if the LIF replication fails.

You can also check the failure by using the `metrocluster check lif show` command. Failures can be corrected by running the `metrocluster check lif repair-placement` command for any LIF that fails to find a correct port. You should resolve any LIF failures as soon as possible to ensure LIF operation in the event of a MetroCluster switchover operation.

Note: Even if the source Storage Virtual Machine (SVM) is down, LIF placement may proceed normally if there is a LIF belonging to a different SVM in a port with the same ipspace and network in the destination.

Placement of replicated LIFs when the DR partner node is down

When a LIF is created on a node whose DR partner has been taken over, the replicated LIF is placed on the DR auxiliary partner node. After a subsequent giveback, the LIFs are not automatically moved to the DR partner. This could lead to LIFs being concentrated on a single node in the partner cluster. In the event of a MetroCluster switchover operation, subsequent attempts to map LUNs belonging to the SVM will fail.

You should run the `metrocluster check lif show` after a takeover or giveback to ensure correct LIF placement. If errors exist, you can run the `metrocluster check lif repair-placement` command.

Related information

[Clustered Data ONTAP 8.3 Network Management Guide](#)

Volume or FlexClone command VLDB errors

If a volume or FlexClone volume command (such as `volume create` or `volume delete`) fails and the error message indicates that the failure is due to a VLDB error, you should manually retry the job.

If the retry fails with an error that indicates a duplicate volume name, there is a stale entry in the internal volume database. Please call customer support for assistance in removing the stale entry.

Removing the entry helps ensure that configuration inconsistencies do not develop between the two MetroCluster clusters.

Modifying volumes to set NVFAIL in case of switchover

You can modify a volume so that, in event of a MetroCluster switchover, the NVFAIL flag is set on the volume. The NVFAIL flag causes the volume to be fenced off from any modification. This is required for volumes that need to be handled as if committed writes to the volume were lost after the switchover.

Step

1. Enable MetroCluster to trigger NVFAIL on switchover by setting the `vol -dr-force-nvfail` parameter to `on`:

```
vol modify -vserver vservice-name -volume volume-name -dr-force-nvfail on
.
```

Monitoring and protecting database validity by using NVFAIL

The `-nvfail` parameter of the `volume modify` command enables Data ONTAP to detect nonvolatile RAM (NVRAM) inconsistencies when the system is booting or after a switchover operation. It also warns you and protects the system against data access and modification until the volume can be manually recovered.

If Data ONTAP finds any problems, database or file system instances stop responding or shut down, and Data ONTAP sends error messages to the console to alert you to check the state of the database or file system. You can enable NVFAIL to warn database administrators of NVRAM inconsistencies among clustered nodes that can compromise database validity. After a system crash or switchover operation, NFS clients cannot access data from any of the nodes until the NVFAIL state is cleared. CIFS clients are unaffected.

How NVFAIL protects database files

The NVFAIL state is set in two cases, either when Data ONTAP detects NVRAM errors when booting up or when a MetroCluster switchover operation occurs. If no errors are detected at startup, the file service is started normally. However, if NVRAM errors are detected or the `force-fail` option was set and then there was a switchover, Data ONTAP stops database instances from responding.

When you enable the NVFAIL option, one of the following processes takes place during bootstrap.

If...	Then...
Data ONTAP detects no NVRAM errors	File service starts normally.
Data ONTAP detects NVRAM errors	<ul style="list-style-type: none"> Data ONTAP returns a stale file handle (ESTALE) error to NFS clients trying to access the database, causing the application to stop responding, crash, or shut down. Data ONTAP then sends an error message to the system console and log file. When the application restarts, files are available to CIFS clients, even if you have not verified that they are valid. For NFS clients, files remain inaccessible until you reset the <code>in-nvfailed-state</code> option on the affected volume.
Data ONTAP detects NVRAM errors on a volume that contains LUNs	LUNs in that volume are brought offline. Then the <code>in-nvfailed-state</code> option on the volume must be cleared, and the <code>NVFAIL</code> attribute on the LUNs must be cleared by bringing each LUN in the affected volume online. You can perform the steps to check the integrity of the LUNs and recover the LUN from Snapshot copy or backup as necessary. After all the LUNs in the volume are recovered, the <code>in-nvfailed-state</code> option on the affected volume is cleared.

Commands for monitoring data loss events

If you enable the `NVFAIL` option, you receive notification when a system crash caused by NVRAM inconsistencies or a MetroCluster switchover occurs.

By default, the `NVFAIL` parameter is not enabled.

If you want to...	Use this command...
Create a new volume with <code>NVFAIL</code> enabled	<code>volume create -nvfail on</code>
Enable <code>NVFAIL</code> on an existing volume	<code>volume modify</code> Note: You set the <code>-nvfail</code> option to <code>on</code> to enable <code>NVFAIL</code> on the created volume.

If you want to...	Use this command...
Display whether NVFAIL is currently enabled for a specified volume	<pre>volume show</pre> <p>Note: You set the <code>-fields</code> parameter to <code>nvfail</code> to display the NVFAIL attribute for a specified volume.</p>

See the man page for each command for more information.

Accessing volumes in NVFAIL state after a switchover

After a switchover, you must clear the NVFAIL state by resetting the `-in-nvfailed-state` parameter of the `volume modify` command to remove the restriction of clients to access data.

Before you begin

The database or file system must not be running or trying to access the affected volume.

About this task

Setting `-in-nvfailed-state` parameter requires advanced-level privilege.

Step

1. Recover the volume by using the `volume modify` command with the `-in-nvfailed-state` parameter set to **false**.

After you finish

For instructions about examining database file validity, see the documentation for your specific database software.

If your database uses LUNs, review the steps to make the LUNs accessible to the host after an NVRAM failure.

Recovering LUNs in NVFAIL states after switchover

After a switchover, the host no longer has access to data on the LUNs that are in NVFAIL states. You must perform a number of actions before the database has access to the LUNs.

Before you begin

The database must not be running.

Steps

1. Clear the NVFAIL state on the affect volume that hosts the LUNs by resetting the `-in-nvfailed-state` parameter of the `volume modify` command.
2. Bring the affected LUNs online.
3. Examine the LUNs for any data inconsistencies and resolve them.
This might involve host-based recovery or recovery done on the storage controller using SnapRestore.
4. Bring the database application online after recovering the LUNs.

Glossary of MetroCluster terms

aggregate

A grouping of physical storage resources (disks or array LUNs) that provides storage to volumes associated with the aggregate. Aggregates provide the ability to control the RAID configuration for all associated volumes.

data SVM

Formerly known as data Vserver. In clustered Data ONTAP, a Storage Virtual Machine (SVM) that facilitates data access from the cluster; the hardware and storage resources of the cluster are dynamically shared by data SVMs within a cluster.

admin SVM

Formerly known as admin Vserver. In clustered Data ONTAP, a Storage Virtual Machine (SVM) that has overall administrative access to all objects in the cluster, including all objects owned by other SVMs, but does not provide data access to clients or hosts.

inter-switch link (ISL)

A connection between two switches using the E-port.

destination

The storage to which source data is backed up, mirrored, or migrated.

disaster recovery (DR) group

The four nodes in a MetroCluster configuration that synchronously replicate each others' configuration and data.

disaster recovery (DR) partner

A node's partner at the remote MetroCluster site. The node mirrors its DR partner's NVRAM or NVMEM partition.

disaster recovery auxiliary (DR auxiliary) partner

The HA partner of a node's DR partner. The DR auxiliary partner mirrors a node's NVRAM or NVMEM partition in the event of an HA takeover after a MetroCluster switchover operation.

HA pair

- In Data ONTAP 8.x, a pair of nodes whose controllers are configured to serve data for each other if one of the two nodes stops functioning.
Depending on the system model, both controllers can be in a single chassis, or one controller can be in one chassis and the other controller can be in a separate chassis.
- In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an *active/active configuration*.

HA partner

A node's partner within the local HA pair. The node mirrors its HA partner's NVRAM or NVMEM cache.

high availability (HA)

In Data ONTAP 8.x, the recovery capability provided by a pair of nodes (storage systems), called an *HA pair*, that are configured to serve data for each other if one of the two nodes stops functioning. In the Data ONTAP 7.3 and 7.2 release families, this functionality is referred to as an *active/active configuration*.

healing

The two required MetroCluster operations that prepare the storage located at the DR site for switchback. The first heal operation resynchronizes the mirrored plexes. The second heal operation returns ownership of root aggregates to the DR nodes.

LIF (logical interface)

A logical network interface, representing a network access point to a node. LIFs currently correspond to IP addresses, but could be implemented by any interconnect. A LIF is generally bound to a physical network port; that is, an Ethernet port. LIFs can fail over to other physical ports (potentially on other nodes) based on policies interpreted by the LIF manager.

NVRAM

nonvolatile random-access memory.

NVRAM cache

Nonvolatile RAM in a storage system, used for logging incoming write data and NFS requests. Improves system performance and prevents loss of data in case of a storage system or power failure.

NVRAM mirror

A synchronously updated copy of the contents of the storage system NVRAM (nonvolatile random access memory) contents kept on the partner storage system.

node

- In Data ONTAP, one of the systems in a cluster or an HA pair.
To distinguish between the two nodes in an HA pair, one node is sometimes called the *local node* and the other node is sometimes called the *partner node* or *remote node*.
- In Protection Manager and Provisioning Manager, the set of storage containers (storage systems, aggregates, volumes, or qtrees) that are assigned to a dataset and designated either primary data (primary node), secondary data (secondary node), or tertiary data (tertiary node).
A dataset node refers to any of the nodes configured for a dataset.
A backup node refers to either a secondary or tertiary node that is the destination of a backup or mirror operation.
A disaster recovery node refers to the dataset node that is the destination of a failover operation.

remote storage

The storage that is accessible to the local node, but is at the location of the remote node.

root volume

A special volume on each Data ONTAP system. The root volume contains system files and configuration information, and can also contain data. It is required for the system to be able to boot and to function properly. Core dump files, which are important for troubleshooting, are written to the root volume if there is enough space.

switchback

The MetroCluster operation that restores service back to one of the MetroCluster sites.

switchover

The MetroCluster operation that transfers service from one of the MetroCluster sites.

- A *negotiated* switchover is planned in advance and cleanly shuts down components of the target MetroCluster site.
- A *forced* switchover immediately transfers service; the shut down of the target site might not be clean.

Copyright information

Copyright © 1994–2015 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.

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

Trademark information

NetApp, the NetApp logo, Go Further, Faster, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

Cisco and the Cisco logo are trademarks of Cisco in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

How to send comments about documentation and receive update notification

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 to 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., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

- 7-Mode fabric MetroCluster
 - sharing FC switch fabric [25](#)
- 7-Mode MetroCluster configurations
 - differences with clustered MetroCluster configurations [8](#)

A

- addresses
 - gathering required network information [107](#), [110](#)
- aggregates
 - mirrored data, creating on each node of a MetroCluster configuration [130](#)
- architecture
 - of MetroCluster configurations [9](#)
- array LUNs
 - in a MetroCluster configuration [14](#)
 - planning for a MetroCluster configuration with [141](#)
- array LUNs and disks
 - considerations for MetroCluster configuration with [150](#)
 - example MetroCluster configuration [151](#)
- assigning disk shelves
 - MetroCluster configuration [116](#)
- ATTO FibreBridge
 - See FC-to-SAS bridges
- authentication policy
 - setting [70](#)

B

- bridges
 - installing FC-to-SAS [41](#)
- Brocade
 - license requirements [52](#), [97](#)
- Brocade 6510s
 - sharing during transition [96](#)
- Brocade FC switch configuration
 - configuring ISL ports [58](#), [64](#)
 - configuring zoning [65](#)
 - setting basic settings [56](#)
 - setting to default values [53](#)
- Brocade switches
 - enabling ISL encryption [72](#)
 - setting authentication policy [70](#)

- setting ISL encryption [69](#)
- setting payload for [69](#)

C

- cabling inter-switch link (ISL) [36](#)
- checking
 - MetroCluster configuration operation [135](#)
- checklists
 - hardware setup, for factory configured clusters [18](#)
 - software setup, for factory-configured MetroCluster [20](#)
- Cisco 9148 switches
 - manually enabling ports in [78](#)
- Cisco FC switch configuration
 - calculating buffer-to-buffer credits [88](#)
 - community string [75](#)
 - configuring ISL ports [81](#), [88](#)
 - configuring port channels [88](#)
 - configuring VSANs [83](#)
 - configuring zoning [92](#)
 - setting basic settings [75](#)
 - setting to default values [74](#)
- Cisco FC switches
 - configuration requirements [73](#)
 - configuring with configuration files [50](#)
- Cisco switch configuration
 - saving [95](#)
- Cisco switches
 - manually enabling ports in MDS 9148 [78](#)
 - port licensing [76](#)
- cluster configurations
 - hardware and software [113](#)
 - regular and MetroCluster [113](#)
- cluster interconnects
 - cabling MetroCluster configurations [37](#)
- cluster peering
 - illustration [14](#)
 - MetroCluster [14](#)
 - MetroCluster configuration [122](#)
 - to an outside cluster [155](#)
- cluster peering connections
 - cabling in MetroCluster configurations [37](#)
- cluster peers
 - creating relationships between [128](#)
- clustered MetroCluster configurations

- differences with 7-Mode MetroCluster configurations [8](#)
 - clusters
 - example names [15](#)
 - commands
 - metrocluster configure [132](#)
 - volume [159](#)
 - comments
 - how to send feedback about documentation [167](#)
 - community string
 - FC switches [75](#)
 - Health Monitors [75](#)
 - components
 - preconfigured when new [17](#)
 - racking [30, 97](#)
 - Config Advisor
 - checking for common configuration errors [137](#)
 - downloading and running [137](#)
 - configuration backup files
 - for MetroCluster clusters
 - setting remote destinations for preservation [140](#)
 - configuration files
 - configuring Brocade FC switches [50](#)
 - configuring Cisco FC switches with [50](#)
 - configurations
 - implementing MetroCluster configurations [132](#)
 - MetroCluster
 - overview of implementing with array LUNs [143](#)
 - considerations
 - MetroCluster configuration with disks and array LUNs [150](#)
 - controller ports
 - checking connectivity with partner site [37](#)
 - controllers
 - racking [30, 97](#)
 - verifying and setting HA state [118](#)
- ## D
- data aggregates
 - mirrored, creating on each node of a MetroCluster configuration [130](#)
 - data field size
 - setting [69](#)
 - Data ONTAP health monitoring
 - for FC-to-SAS bridges [135](#)
 - data ports
 - cabling [39](#)
 - configuring intercluster LIFs to share [126](#)
 - database files
 - how NVFAIL protects [158](#)
 - databases
 - accessing after a switchover [160](#)
 - introduction to using NVFAIL to monitor and protect validity of [158](#)
 - destinations
 - specifying URL for configuration backup [140](#)
 - disabling
 - virtual fabric in a Brocade switch [69](#)
 - disaster recovery group
 - MetroCluster configurations [9](#)
 - disk assignment
 - verifying in a MetroCluster configuration [114](#)
 - disk shelves
 - assigning in a MetroCluster configuration [116](#)
 - racking [30, 97](#)
 - disks and array LUNs
 - considerations for MetroCluster configuration with [150](#)
 - documentation
 - how to receive automatic notification of changes to [167](#)
 - how to send feedback about [167](#)
 - where to find MetroCluster documentation [6](#)
- ## E
- E-ports
 - configuring [58](#)
 - events
 - monitoring data loss [159](#)
 - example names
 - MetroCluster components [15](#)
- ## F
- fabric-attached MetroCluster configurations
 - array LUN requirements [141](#)
 - failover and giveback
 - verifying in a MetroCluster configuration [138](#)
 - FC switch
 - verifying [105](#)
 - FC switch configuration
 - configuring ISL ports
 - Cisco [81](#)
 - configuring ports
 - Cisco [79](#)
 - configuring VSANs
 - Cisco [83](#)

- downloading configuration files, Brocade [50](#)
 - setting basic settings
 - Brocade [56](#)
 - Cisco [75](#)
 - community string [75](#)
 - setting to default values
 - Brocade [53](#)
 - Cisco [74](#)
 - switch names
 - setting for Brocade switches [53](#)
 - worksheet [26](#)
- FC switch configurations
 - for Brocade switch fabrics [51](#)
 - recommended port assignments [39](#)
 - requirements [51](#)
- FC switch fabric
 - reenabling [105](#)
- FC switch fabrics
 - redundant configuration in the MetroCluster architecture [13](#)
 - sharing during 7-Mode transition [96](#)
- FC switches
 - configuration requirements for Cisco [73](#)
 - configuring Cisco, with configuration files [50](#)
 - configuring for health monitoring [134](#)
 - example names [15](#)
 - introduction to manually configuring Cisco and Brocade [51](#)
 - racking [30, 97](#)
- FC-to-SAS bridge ports
 - cabling during 7-Mode transition [98](#)
- FC-to-SAS bridges
 - configuring for health monitoring [135](#)
 - example names [15](#)
 - in the MetroCluster architecture [12](#)
 - installing [41](#)
 - meeting preinstallation requirements [42](#)
- FC-VI ports
 - cabling [31](#)
 - cabling during 7-Mode transition [98](#)
- feedback
 - how to send comments about documentation [167](#)
- FibreBridge
 - See* FC-to-SAS bridges
- files
 - how NVFAIL protects database [158](#)
- FlexClone command errors
 - in MetroCluster configurations [157](#)

H

- HA interconnects
 - cabling [38](#)
- HA pair operation
 - verifying in a MetroCluster configuration [138](#)
- HA pairs
 - illustration of local MetroCluster [12](#)
 - MetroCluster configurations
 - illustration of local HA pairs [12](#)
- HA state
 - verifying and setting controller [118](#)
- hardware
 - setup checklist for factory configured clusters [18](#)
- hardware components
 - racking [30, 97](#)
- HBA ports
 - cabling during 7-Mode transition [98](#)
- healing
 - verifying in a MetroCluster configuration [140](#)
- health monitoring
 - configuring FC switches for health monitoring [134](#)
 - for FC-to-SAS bridges [135](#)
- Health Monitors
 - community string [75](#)
- host names
 - gathering required network information [107, 110](#)

I

- illustration
 - FC switch fabrics in the MetroCluster architecture [13](#)
 - FC-to-SAS bridges in the MetroCluster architecture [12](#)
- Infinite Volumes
 - in a MetroCluster configuration [155](#)
- information
 - how to send feedback about improving documentation [167](#)
- initial configuration
 - performing with the System Setup tool [119](#)
- installation
 - for systems sharing an FC switch fabric [25](#)
 - for systems with array LUNs [25](#)
 - for systems with native disks [25](#)
 - preparations for installing FC-to-SAS bridges [42](#)
- intercluster LIFs
 - configuring to share data ports [126](#)
 - configuring to use dedicated intercluster ports [122](#)

- intercluster networks
 - configuring intercluster LIFs for [122, 126](#)
- intercluster ports
 - configuring intercluster LIFs to use dedicated [122](#)
- IOD settings
 - deleting TI zoning and configuring [102](#)
- IP addresses
 - gathering required network information [107, 110](#)
- ISL encryption
 - enabling in a MetroCluster configuration [72](#)
 - enabling in switches [72](#)
 - setting on Brocade switches [69](#)
- ISL port group
 - sharing a switch fabric [104](#)
- ISL ports
 - alternate name for [58](#)
 - configuring [58](#)
 - configuring on Brocade switches [64](#)
 - configuring on Cisco switches [88](#)

J

- jobs
 - scheduling in MetroCluster configurations [155](#)

L

- licensing
 - in a MetroCluster configuration [155](#)
- LIF creation
 - in a MetroCluster configuration [156](#)
- LIF replication
 - in a MetroCluster configuration [156](#)
- LIFs
 - configuring to use dedicated intercluster ports [122](#)
- LIFs, intercluster
 - configuring to share data ports [126](#)
- local HA pairs
 - illustration of MetroCluster [12](#)
 - verifying operation of in a MetroCluster configuration [138](#)
- LUNs
 - recovering after NVRAM failures [160](#)
- LUNs (array)
 - MetroCluster configuration implementation overview [143](#)

M

- management ports

- cabling [39](#)
- MetroCluster
 - architecture of [9](#)
 - disaster recovery group [9](#)
 - illustration of
 - MetroCluster illustration [9](#)
 - sharing existing switch fabric [96](#)
- MetroCluster architecture
 - cluster peering [14](#)
 - FC switch fabrics [13](#)
 - FC-to-SAS bridges [12](#)
 - illustration
 - cluster peering in the MetroCluster architecture [14](#)
- MetroCluster components
 - cabling cluster interconnects [37](#)
 - cabling Inter-Switch Links (ISLs) [36](#)
 - cabling the HA interconnect [38](#)
 - racking disk shelves [30, 97](#)
 - racking FC switches [30, 97](#)
 - racking storage controllers [30, 97](#)
- MetroCluster configuration networking
 - IPv6 configuration [156](#)
 - LIF creation [156](#)
 - LIF replication [156](#)
- MetroCluster configurationd
 - setting the authentication policy in Brocade switch [70](#)
- MetroCluster configurations
 - array LUN requirements and restrictions [141](#)
 - array LUNs
 - connecting devices in [144](#)
 - implementation overview [143](#)
 - requirements for zoning with [147](#)
 - setting up ONTAP after connecting devices [149](#)
 - array LUNs and disks [151](#)
 - cabling FC-VI adapters [31](#)
 - cabling Inter-Switch Links (ISLs) [36](#)
 - cabling MetroCluster components
 - cabling FC-VI adapters
 - cabling HBA adapters [31](#)
 - checking operation [135](#)
 - creating mirrored data aggregates on each node of [130](#)
 - disks and array LUNs [150](#)
 - enabling ISL encryption on Brocade switches [72](#)
 - example configuration [151](#)
 - implementing [132](#)
 - setting ISL encryption on Brocade switch [69](#)
 - storage arrays [150](#)

172 | MetroCluster Installation and Configuration Guide

- metrocluster configure command
 - creating MetroCluster relationships [132](#)
- MetroCluster hardware installation
 - workflow [28](#)

N

- native disk shelves
 - in a MetroCluster configuration [14](#)
- network information
 - gathering required [110](#)
 - gathering required network information [107](#)
- node configuration after MetroCluster setup
 - additional configuration with OnCommand System Manager [154](#)
- nodes
 - creating mirrored data aggregates on each MetroCluster [130](#)

- NVFAIL
 - description of [158](#)
 - how it protects database files [158](#)
 - modifying volumes to set NVFAIL in case of switchover [158](#)

- NVRAM failures
 - recovering LUNs after [160](#)

O

- OnCommand Performance Manager
 - monitoring with [154](#)
- OnCommand System Manager
 - node configuration with [154](#)
- OnCommand Unified Manager
 - monitoring with [154](#)

P

- partner cluster
 - cabling cluster peering connections [37](#)
- payload
 - setting [69](#)
- peer relationships
 - creating cluster [128](#)
- peering clusters
 - MetroCluster configuration [122](#)
- planning
 - gathering required network information [107](#), [110](#)
- port channels
 - configuring on Cisco switches [88](#)
- port configuration

- on FC switch
 - Cisco [79](#)
- port licensing
 - Cisco switches [76](#)
- port numbers
 - Brocade FC switches [39](#)
 - Cisco FC switch [39](#)
- ports
 - configuring intercluster LIFs to use dedicated intercluster [122](#)
 - manually enabling in Cisco MDS 9148 [78](#)
- ports, data
 - configuring intercluster LIFs to share [126](#)
- preconfiguration
 - MetroCluster components [17](#)

Q

- QoS policies
 - VSANs on Cisco switches [83](#)

R

- relationships
 - creating cluster peer [128](#)
 - creating MetroCluster relationships [132](#)
- required hardware
 - MetroCluster [15](#)
- requirements
 - gathering network information [107](#), [110](#)
- root aggregates
 - mirroring [130](#)
 - volume creation on [156](#)

S

- SAS disk shelves
 - installing [41](#)
- setting
 - authentication policy in Brocade switch in a MetroCluster configuration [70](#)
- setup
 - hardware, checklist for factory configured clusters [18](#)
 - software, checklist for factory configured MetroCluster [20](#)
- software
 - settings already enabled [17](#)
 - setup checklist for factory configured clusters [20](#)
- storage controllers

- example names [15](#)
- suggestions
 - how to send feedback about documentation [167](#)
- switch configuration, Cisco
 - saving [95](#)
- switch fabrics
 - configuring switch ports [58](#)
 - disabling before modifying configuration [100](#)
 - enabling sharing for a port group [104](#)
 - Fibre Channel [13](#)
 - reenabling [100](#)
- switch parameters
 - setting [69](#)
- switch zoning
 - requirements for MetroCluster configuration with array LUNs [147](#)
- switchback
 - verifying in a MetroCluster configuration [140](#)
- switches
 - configuration requirements for Cisco FC [73](#)
 - licensing requirements
 - Brocade [52](#), [97](#)
 - Cisco [74](#)
 - requirements in MetroCluster configurations with array LUNs [141](#)
- switchover
 - accessing the database after [160](#)
 - verifying in a MetroCluster configuration [140](#)
- SyncMirror
 - requirements for a MetroCluster configuration [141](#)
- System Setup tool
 - performing basic software configuration [119](#)
 - performing cluster configuration [119](#)

T

- third-party storage
 - MetroCluster configurations with, requirements and restrictions [141](#)
- TI zoning
 - deleting and configuring IOD settings [102](#)
- transition
 - 7-Mode to clustered Data ONTAP [17](#)
 - sharing FC switch fabric [17](#)
- twitter
 - how to receive automatic notification of documentation changes [167](#)

U

- utilities
 - checking for common configuration errors with Config Advisor [137](#)
 - downloading and running Config Advisor [137](#)

V

- verification
 - booting to Maintenance mode [114](#)
 - performing before booting to Data ONTAP [114](#)
 - verifying disk assignment [114](#)
- verifying
 - MetroCluster configuration operation [135](#)
- Virtual fabric
 - disabling [69](#)
- VLDB errors
 - in MetroCluster configurations [157](#)
- volume command errors
 - in MetroCluster configurations [157](#)
- volume creation
 - in a MetroCluster configuration [156](#)
- volumes
 - commands [159](#)
 - recovering after a switchover [160](#)
- VSANs
 - Cisco [83](#)
- Vservers
 - See* SVMs

W

- workflow
 - MetroCluster hardware installation [28](#)
 - MetroCluster software configuration [113](#)
- worksheet
 - FC switch configuration [26](#)
 - for site configuration [107](#), [110](#)

Z

- zoning
 - sharing a switch fabric [104](#)
- zoning configuration
 - configuring on a Brocade FC switch [65](#)
 - configuring on a Cisco FC switch [92](#)