

Install a fabric-attached MetroCluster

ONTAP MetroCluster

NetApp August 29, 2025

This PDF was generated from https://docs.netapp.com/us-en/ontap-metrocluster/install-fc/index.html on August 29, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Install a fabric-attached MetroCluster	1
Overview	1
Prepare for the MetroCluster installation	1
Differences among the ONTAP MetroCluster configurations	1
Cluster peering	2
Considerations for ISLs	5
Considerations for using TDM/WDM equipment with fabric-attached MetroCluster configurations	6
Requirements for using a Brocade DCX 8510-8 switch	7
Considerations when using unmirrored aggregates	8
Firewall usage at MetroCluster sites	9
Cable a fabric-attached MetroCluster configuration	10
Cabling a fabric-attached MetroCluster configuration	10
Parts of a fabric MetroCluster configuration	10
Required MetroCluster FC components and naming conventions	17
Configuration worksheets for FC switches and FC-to-SAS bridges	21
Install and cable MetroCluster components	21
Configure the FC switches	57
Install FC-to-SAS bridges and SAS disk shelves	182
Configuring the MetroCluster software in ONTAP	197
Gathering required information	198
Similarities and differences between standard cluster and MetroCluster configurations	204
Verifying and configuring the HA state of components in Maintenance mode	205
Restoring system defaults and configuring the HBA type on a controller module	206
Configuring FC-VI ports on a X1132A-R6 quad-port card on FAS8020 systems	208
Verifying disk assignment in Maintenance mode in an eight-node or a four-node configuration	210
Verifying disk assignment in Maintenance mode in a two-node configuration	217
Setting up ONTAP	218
Configuring the clusters into a MetroCluster configuration	224
Checking for MetroCluster configuration errors with Config Advisor	254
Verifying local HA operation	255
Verifying switchover, healing, and switchback	257
Protecting configuration backup files	257
Considerations for using virtual IP and Border Gateway Protocol with a MetroCluster configuration	257
ONTAP limitations	258
Guidelines for using this Layer 3 solution with a MetroCluster configuration	259
Testing the MetroCluster configuration.	259
Verifying negotiated switchover	260
Verifying healing and manual switchback	261
Loss of a single FC-to-SAS bridge	264
Verifying operation after power line disruption	266
Verifying operation after a switch fabric failure	267
Verifying operation after loss of a single storage shelf.	269
Remove MetroCluster configurations	279

How to use the Active IQ Unified Manager and ONTAP System Manager for further configuration and	
monitoring	0
Use the Active IQ Unified Manager and ONTAP System Manager for further configuration and	
monitoring	0
Synchronize the system time using NTP	0
Considerations when using ONTAP in a MetroCluster configuration	1
FlexCache support in a MetroCluster configuration 28	2
FabricPool support in MetroCluster configurations 28	2
FlexGroup support in MetroCluster configurations 28	3
Consistency group support in MetroCluster configurations	3
Job schedules in a MetroCluster configuration	3
Cluster peering from the MetroCluster site to a third cluster	4
LDAP client configuration replication in a MetroCluster configuration	4
Networking and LIF creation guidelines for MetroCluster configurations.	4
SVM disaster recovery in a MetroCluster configuration	9
Output for the "storage aggregate plex show" command is indeterminate after a MetroCluster	
switchover	2
Modifying volumes to set the NVFAIL flag in case of switchover	2
Where to find additional information	3
MetroCluster and miscellaneous information 29	3

Install a fabric-attached MetroCluster

Overview

To install your fabric-attached MetroCluster configuration, you must perform a number of procedures in the correct order.

- Prepare for the installation and understand all requirements.
- Cable the components
- Configure the software
- Test the configuration

Prepare for the MetroCluster installation

Differences among the ONTAP MetroCluster configurations

The various MetroCluster configurations have key differences in the required components.

In all configurations, each of the two MetroCluster sites are configured as an ONTAP cluster. In a two-node MetroCluster configuration, each node is configured as a single-node cluster.

Feature	IP configurations	Fabric attached	configurations	Stretch configurations	
		Four- or eight- node	Two-node	Two-node bridge-attached	Two-node direct-attached
Number of controllers	Four or eight ¹	Four or eight	Two	Two	Two
Uses an FC switch storage fabric	No	Yes	Yes	No	No
Uses an IP switch storage fabric	Yes	No	No	No	No
Uses FC-to-SAS bridges	No	Yes	Yes	Yes	No
Uses direct- attached SAS storage	Yes (local attached only)	No	No	No	Yes

Supports ADP	Yes (beginning with ONTAP 9.4)	No	No	No	No
Supports local HA	Yes	Yes	No	No	No
Supports ONTAP automatic unplanned switchover (AUSO)	No	Yes	Yes	Yes	Yes
Supports unmirrored aggregates	Yes (beginning with ONTAP 9.8)	Yes	Yes	Yes	Yes
Supports ONTAP Mediator	Yes (beginning with ONTAP 9.7)	No	No	No	No
Supports MetroCluster Tiebreaker	Yes (not in combination with ONTAP Mediator)	Yes	Yes	Yes	Yes
Supports All SAN Arrays	Yes	Yes	Yes	Yes	Yes

Notes

1. Review the following considerations for eight-node MetroCluster IP configurations:

- Eight-node configurations are supported beginning with ONTAP 9.9.1.
- Only NetApp-validated MetroCluster switches (ordered from NetApp) are supported.
- Configurations using IP-routed (layer 3) backend connections are not supported.

Support for All SAN Array systems in MetroCluster configurations

Some of the All SAN Arrays (ASAs) are supported in MetroCluster configurations. In the MetroCluster documentation, the information for AFF models applies to the corresponding ASA system. For example, all cabling and other information for the AFF A400 system also applies to the ASA AFF A400 system.

Supported platform configurations are listed in the NetApp Hardware Universe.

Cluster peering

Each MetroCluster site is configured as a peer to its partner site. You must be familiar with the prerequisites and guidelines for configuring the peering relationships. This is important when deciding on whether to use shared or dedicated ports for those relationships.

Related information

Cluster and SVM peering express configuration

Prerequisites for cluster peering

Before you set up cluster peering, you should confirm that connectivity between port, IP address, subnet, firewall, and cluster-naming requirements are met.

Connectivity requirements

Every intercluster LIF on the local cluster must be able to communicate with every intercluster LIF on the remote cluster.

Although it is not required, it is typically simpler to configure the IP addresses used for intercluster LIFs in the same subnet. The IP addresses can reside in the same subnet as data LIFs, or in a different subnet. The subnet used in each cluster must meet the following requirements:

• The subnet must have enough IP addresses available to allocate to one intercluster LIF per node.

For example, in a four-node cluster, the subnet used for intercluster communication must have four available IP addresses.

Each node must have an intercluster LIF with an IP address on the intercluster network.

Intercluster LIFs can have an IPv4 address or an IPv6 address.



ONTAP 9 enables you to migrate your peering networks from IPv4 to IPv6 by optionally allowing both protocols to be present simultaneously on the intercluster LIFs. In earlier releases, all intercluster relationships for an entire cluster were either IPv4 or IPv6. This meant that changing protocols was a potentially disruptive event.

Port requirements

You can use dedicated ports for intercluster communication, or share ports used by the data network. Ports must meet the following requirements:

• All ports used to communicate with a given remote cluster must be in the same IPspace.

You can use multiple IPspaces to peer with multiple clusters. Pair-wise full-mesh connectivity is required only within an IPspace.

• The broadcast domain used for intercluster communication must include at least two ports per node so that intercluster communication can fail over from one port to another port.

Ports added to a broadcast domain can be physical network ports, VLANs, or interface groups (ifgrps).

- All ports must be cabled.
- All ports must be in a healthy state.
- The MTU settings of the ports must be consistent.

Firewall requirements

Firewalls and the intercluster firewall policy must allow the following protocols:

- ICMP service
- TCP to the IP addresses of all the intercluster LIFs over the ports 10000, 11104, and 11105
- Bidirectional HTTPS between the intercluster LIFs

The default intercluster firewall policy allows access through the HTTPS protocol and from all IP addresses (0.0.0.0/0). You can modify or replace the policy if necessary.

Considerations when using dedicated ports

When determining whether using a dedicated port for intercluster replication is the correct intercluster network solution, you should consider configurations and requirements such as LAN type, available WAN bandwidth, replication interval, change rate, and number of ports.

Consider the following aspects of your network to determine whether using a dedicated port is the best intercluster network solution:

- If the amount of available WAN bandwidth is similar to that of the LAN ports, and the replication interval is such that replication occurs while regular client activity exists, then you should dedicate Ethernet ports for intercluster replication to avoid contention between replication and the data protocols.
- If the network utilization generated by the data protocols (CIFS, NFS, and iSCSI) is such that the network utilization is above 50 percent, then dedicate ports for replication to allow for nondegraded performance if a node failover occurs.
- When physical 10 GbE or faster ports are used for data and replication, you can create VLAN ports for replication and dedicate the logical ports for intercluster replication.

The bandwidth of the port is shared between all VLANs and the base port.

• Consider the data change rate and replication interval and whether the amount of data, that must be replicated on each interval, requires enough bandwidth. This might cause contention with data protocols if sharing data ports.

Considerations when sharing data ports

When determining whether sharing a data port for intercluster replication is the correct intercluster network solution, you should consider configurations and requirements such as LAN type, available WAN bandwidth, replication interval, change rate, and number of ports.

Consider the following aspects of your network to determine whether sharing data ports is the best intercluster connectivity solution:

• For a high-speed network, such as a 40-Gigabit Ethernet (40-GbE) network, a sufficient amount of local LAN bandwidth might be available to perform replication on the same 40-GbE ports that are used for data access.

In many cases, the available WAN bandwidth is far less than the 10 GbE LAN bandwidth.

- All nodes in the cluster might have to replicate data and share the available WAN bandwidth, making data port sharing more acceptable.
- Sharing ports for data and replication eliminates the extra port counts required to dedicate ports for replication.
- The maximum transmission unit (MTU) size of the replication network will be the same size as that used on the data network.

- Consider the data change rate and replication interval and whether the amount of data, that must be replicated on each interval, requires enough bandwidth. This might cause contention with data protocols if sharing data ports.
- When data ports for intercluster replication are shared, the intercluster LIFs can be migrated to any other intercluster-capable port on the same node to control the specific data port that is used for replication.

Considerations for ISLs

You must determine how many ISLs you need for each FC switch fabric in the MetroCluster configuration. Beginning with ONTAP 9.2, in some cases, instead of dedicating FC switches and ISLs to each individual MetroCluster configuration, you can share the same four switches.

ISL sharing considerations (ONTAP 9.2)

Beginning with ONTAP 9.2, you can use ISL sharing in the following cases:

- One two-node and one four-node MetroCluster configurations
- Two separate four-node MetroCluster configurations
- Two separate two-node MetroCluster configurations
- Two DR groups within one eight-node MetroCluster configuration

The number of ISLs required between the shared switches depends on the bandwidth of the platform models connected to the shared switches.

Consider the following aspects of your configuration when determining how many ISLs you need.

- Non-MetroCluster devices should not be connected to any of the FC switches that provide the back-end MetroCluster connectivity.
- ISL sharing is supported on all switches except the Cisco 9250i and Cisco 9148 switches.
- All nodes must be running ONTAP 9.2 or later.
- The FC switch cabling for ISL sharing is the same as for the eight-node MetroCluster cabling.
- The RCF files for ISL sharing are same as for the eight-node MetroCluster cabling.
- You should verify that all hardware and software versions are supported.

NetApp Hardware Universe

- The speed and number of ISLs must be sized to support the client load on both MetroCluster systems.
- The back-end ISLs and the back-end components must be dedicated to the MetroCluster configuration only.
- The ISL must use one of the supported speeds: 4 Gbps, 8 Gbps, 16 Gbps, or 32 Gbps.
- The ISLs on one fabric should all be the same speed and length.
- The ISLs on one fabric should all have the same topology. For example, they should all be direct links, or if your system uses WDM, then they should all use WDM.

Platform-specific ISL considerations

The number of recommended ISLs is platform-model specific. The following table shows the ISL requirements for each fabric by platform model. It assumes that each ISL has a 16-Gbps capacity.

Platform model	Recommended number of ISLs per four-node DR group (per switch fabric)
AFF A900 and FAS9500	Eight
AFF A700	Six
FAS9000	Six
8080	Four
All others	Two

If the switch fabric is supporting eight nodes (either part of a single, eight-node MetroCluster configuration, or two four-node configurations that are sharing ISLs), the recommended total number of ISLs for the fabric is the sum of that required for each four-node DR group. For example:

- If DR group 1 includes four AFF A700 systems, it requires six ISLs.
- If DR group 2 includes four FAS8200 systems, it requires two ISLs.
- The total number of recommended ISLs for the switch fabric is eight.

Considerations for using TDM/WDM equipment with fabric-attached MetroCluster configurations

The Hardware Universe tool provides some notes about the requirements that Time Division Multiplexing (TDM) or Wavelength Division Multiplexing (WDM) equipment must meet to work with a fabric-attached MetroCluster configuration. These notes also include information about various configurations, which can help you to determine when to use in-order delivery (IOD) of frames or out-of-order delivery (OOD) of frames.

An example of such requirements is that the TDM/WDM equipment must support the link aggregation (trunking) feature with routing policies. The order of delivery (IOD or OOD) of frames is maintained within a switch, and is determined by the routing policy that is in effect.

NetApp Hardware Universe

The following table provides the routing policies for configurations containing Brocade switches and Cisco switches:

Switches	Configuring MetroCluster	Configuring MetroCluster		
	configurations for IOD	configurations for OOD		

Brocade	 AptPolicy must be set to 1 	AptPolicy must be set to 3
	DLS must be set to off	DLS must be set to on
	 IOD must be set to on 	IOD must be set to off
Cisco	Policies for the FCVI-designated VSAN:	Not applicable
	 Load balancing policy: srcid and dstid 	
	 IOD must be set to on 	
	Policies for the storage-designated VSAN:	
	 Load balancing policy: srcid, dstid, and oxid 	
	 VSAN must not have the in- order-guarantee option set 	

When to use IOD

It is best to use IOD if it is supported by the links. The following configurations support IOD:

- A single ISL
- The ISL and the link (and the link equipment, such as TDM/WDM, if used) supports configuration for IOD.
- A single trunk, and the ISLs and the links (and the link equipment, such as TDM/WDM, if used) support configuration for IOD.

When to use OOD

- You can use OOD for all configurations that do not support IOD.
- You can use OOD for configurations that do not support the trunking feature.

Using encryption devices

When using dedicated encryption devices on the ISL or encryption on WDM devices in the MetroCluster configuration, you must meet the following requirements:

• The external encryption devices or WDM equipment has been self certified by the vendor with the FC switch in question.

The self certification should cover the operating mode (such as trunking and encryption).

• The added latency due to encryption should be no more than 10 microseconds.

Requirements for using a Brocade DCX 8510-8 switch

As you prepare for the MetroCluster installation, you should understand the MetroCluster hardware architecture and required components.

- The DCX 8510-8 switches used in MetroCluster configurations must be purchased from NetApp.
- For scalability, you should leave one port-chunk between MetroCluster configurations if cabling only two MetroClusters in 4x48-port modules. This enables you to expand port usage in MetroCluster configurations without recabling.
- Each Brocade DCX 8510-8 switch in the MetroCluster configuration must be correctly configured for the ISL ports and storage connections. For port usage, see the following section: Port assignments for FC switches.
- ISLs cannot be shared and each MetroCluster requires two ISLs for each fabric.
- The DCX 8510-8 switch used for backend MetroCluster connectivity should not be used for any other connectivity.

Non-MetroCluster devices should not be connected to these switches and non-MetroCluster traffic should not flow through DCX 8510-8 switches.

• One line card can either be connected to ONTAP MetroClusters or ONTAP 7-Mode MetroClusters.



RCF files are not available for this switch.

The following are the requirements for using two Brocade DCX 8510-8 switches:

- You must have one DCX 8510-8 switch at each site.
- You must use a minimum of two 48-port blades that contain 16Gb SFPs in each switch.

The following are the requirements for using four DCX 8510-8 switches at each site in a MetroCluster configuration:

- You must have two DCX 8510-8 switches at each site.
- You must use at least one 48-port blade for each DCX 8510-8 switch.
- Each blade is configured as a virtual switch using virtual fabrics.

The following NetApp products are not supported by Brocade DCX 8510-8 switches:

- · Config Advisor
- Fabric Health Monitor
- MyAutoSupport (system risks might show false positives)
- Active IQ Unified Manager (formerly OnCommand Unified Manager)



Ensure that all the components needed for this configuration are in the NetApp Interoperability Matrix Tool. Read the notes section in the Interoperability Matrix Tool for information on supported configurations.

Considerations when using unmirrored aggregates

Considerations when using unmirrored aggregates

If your configuration includes unmirrored aggregates, you must be aware of potential access issues that follow switchover operations.

Considerations for unmirrored aggregates when doing maintenance requiring power shutdown

If you are performing a negotiated switchover for maintenance reasons requiring site-wide power shutdown, you should first manually take offline any unmirrored aggregates owned by the disaster site.

If you do not take any unmirrored aggregates offline, nodes at the surviving site might go down due to multidisk panics. This could occur if switched over unmirrored aggregates go offline, or are missing, because of the loss of connectivity to storage at the disaster site. This is the result of a power shutdown or a loss of ISLs.

Considerations for unmirrored aggregates and hierarchical namespaces

If you are using hierarchical namespaces, you should configure the junction path so that all of the volumes in that path are either on mirrored aggregates only or on unmirrored aggregates only. Configuring a mix of unmirrored and mirrored aggregates in the junction path might prevent access to the unmirrored aggregates after the switchover operation.

Considerations for unmirrored aggregates and CRS metadata volume and data SVM root volumes

The configuration replication service (CRS) metadata volume and data SVM root volumes must be on a mirrored aggregate. You cannot move these volumes to an unmirrored aggregate. If they are on an unmirrored aggregate, negotiated switchover and switchback operations are vetoed. The MetroCluster check command provides a warning if this is the case.

Considerations for unmirrored aggregates and SVMs

SVMs should be configured on mirrored aggregates only, or on unmirrored aggregates only. Configuring a mix of unmirrored and mirrored aggregates can result in a switchover operation that exceeds 120 seconds and result in a data outage if the unmirrored aggregates do not come online.

Considerations for unmirrored aggregates and SAN

In ONTAP versions prior to 9.9.1, a LUN should not be located on an unmirrored aggregate. Configuring a LUN on an unmirrored aggregate can result in a switchover operation that exceeds 120 seconds and a data outage.

Firewall usage at MetroCluster sites

Considerations for firewall usage at MetroCluster sites

If you are using a firewall at a MetroCluster site, you must ensure access for required ports.

The following table shows TCP/UDP port usage in an external firewall positioned between two MetroCluster sites.

Traffic type	Port/services
Cluster peering	11104 / TCP
	11105 / TCP
ONTAP System Manager	443 / TCP

MetroCluster IP intercluster LIFs	65200 / TCP
	10006 / TCP and UDP
Hardware assist	4444 / TCP

Cable a fabric-attached MetroCluster configuration

Cabling a fabric-attached MetroCluster configuration

The MetroCluster components must be physically installed, cabled, and configured at both geographic sites.



Parts of a fabric MetroCluster configuration

Parts of a fabric MetroCluster configuration

As you plan your MetroCluster configuration, you should understand the hardware components and how they interconnect.

Disaster Recovery (DR) groups

A fabric MetroCluster configuration consists of one or two DR groups, depending on the number of nodes in the MetroCluster configuration. Each DR group consists of four nodes.

• An eight-node MetroCluster configuration consists of two DR groups.

• A four-node MetroCluster configuration consists of one DR group.

The following illustration shows the organization of nodes in an eight-node MetroCluster configuration:



The following illustration shows the organization of nodes in a four-node MetroCluster configuration:



Key hardware elements

A 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 FC 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 the cluster configuration, which includes storage virtual machine (SVM) configuration. The configuration of all of the SVMs on one cluster is mirrored to the partner cluster.

Eight-node fabric MetroCluster configuration

An eight-node configuration consists of two clusters, one at each geographically separated site. cluster_A is located at the first MetroCluster site. cluster_B is located at the second MetroCluster site. Each site has one SAS storage stack. 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 is not shown.

An eight-node 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 that an external interconnect is not required.

· Ethernet connections from the controllers to the customer-provided network that is used for cluster peering

SVM configuration is replicated over the cluster peering network.

• A cluster interconnect between each controller in the local cluster

Four-node fabric MetroCluster configuration

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



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



Two-node fabric MetroCluster configuration

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



A two-node configuration consists of two clusters, one at each geographically separated site. cluster_A is located at the first MetroCluster site. cluster_B is located at the second MetroCluster site. Each site has one SAS storage stack. Additional storage stacks are supported, but only one is shown at each site.



In a two-node configuration, the nodes are not configured as an HA pair.

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



A two-node configuration includes the following connections:

- FC connections between the FC-VI adapter on each controller module
- FC connections from each controller module's HBAs to the FC-to-SAS bridge for each SAS shelf stack
- SAS connections between each SAS shelf and from the top and bottom of each stack to an FC-to-SAS bridge
- Ethernet connections from the controllers to the customer-provided network that is used for cluster peering

SVM configuration is replicated over the cluster peering network.

Illustration of the local HA pairs in a MetroCluster configuration

In eight-node or four-node MetroCluster configurations, each site consists of storage controllers configured as one or two HA pairs. 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

Illustration of redundant FC-to-SAS bridges

Redundant FC switch fabrics

Illustration of the cluster peering network

ONTAP concepts

Illustration of redundant FC-to-SAS bridges

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



Related information

Illustration of the local HA pairs in a MetroCluster configuration

Redundant FC switch fabrics

Illustration of the cluster peering network

Redundant FC switch fabrics

Each switch fabric includes inter-switch links (ISLs) that connect the sites. Data is replicated from site-to-site over the ISL. Each switch fabric must be on different physical paths for redundancy.



Related information

Illustration of the local HA pairs in a MetroCluster configuration

Illustration of redundant FC-to-SAS bridges

Illustration of the cluster peering network

Illustration of the cluster peering network

The two clusters in the MetroCluster configuration are peered through a customerprovided 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 the SVM configuration is carried out over this network through the Configuration Replication Service.



Related information

Illustration of the local HA pairs in a MetroCluster configuration

Illustration of redundant FC-to-SAS bridges

Redundant FC switch fabrics

Cluster and SVM peering express configuration

Considerations for configuring cluster peering

Cabling the cluster peering connections

Peering the clusters

Required MetroCluster FC components and naming conventions

When planning your MetroCluster FC configuration, you must understand the required and supported hardware and software components. For convenience and clarity, you should also understand the naming conventions used for components in examples throughout the documentation. For example, one site is referred to as Site A and the other site is referred to as Site B.

Supported software and hardware

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

NetApp Hardware Universe

When using AFF systems, all controller modules in the MetroCluster configuration must be configured as AFF systems.



Long-wave SFPs are not supported in the MetroCluster storage switches. For a table of supported SPFs, see the MetroCluster Technical Report.

Hardware redundancy in the MetroCluster FC configuration

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

Requirement for two ONTAP clusters

The fabric-attached MetroCluster FC configuration requires two ONTAP clusters, one at each MetroCluster site.

Naming must be unique within the MetroCluster configuration.

Example names:

- Site A: cluster_A
- Site B: cluster_B

Requirement for four FC switches

The fabric-attached MetroCluster FC configuration requires four FC switches (supported Brocade or Cisco models).

The four switches form two switch storage fabrics that provide the ISL between each of the clusters in the MetroCluster FC configuration.

Naming must be unique within the MetroCluster configuration.

Requirement for two, four, or eight controller modules

The fabric-attached MetroCluster FC configuration requires two, four, or eight controller modules.

In a four or eight-node MetroCluster configuration, the controller modules at each site form one or two HA pairs. Each controller module has a DR partner at the other site.

The controller modules must meet the following requirements:

- Naming must be unique within the MetroCluster configuration.
- All controller modules in the MetroCluster configuration must be running the same version of ONTAP.
- All controller modules in a DR group must be of the same model.

However, in configurations with two DR groups, each DR group can consist of different controller module models.

• All controller modules in a DR group must use the same FC-VI configuration.

Some controller modules support two options for FC-VI connectivity:

- · Onboard FC-VI ports
- An FC-VI card in slot 1 A mix of one controller module using onboard FC-VI ports and another using an add-on FC-VI card is not supported. For example, if one node uses onboard FC-VI configuration, then all other nodes in the DR group must use onboard FC-VI configuration as well.

Example names:

- Site A: controller_A_1
- Site B: controller_B_1

Requirement for four cluster interconnect switches

The fabric-attached MetroCluster FC configuration requires four cluster interconnect switches (if you are not using two-node switchless clusters)

These switches provide cluster communication among the controller modules in each cluster. The switches are not required if the controller modules at each site are configured as a two-node switchless cluster.

Requirement for FC-to-SAS bridges

The fabric-attached MetroCluster FC configuration requires one pair of FC-to-SAS bridges for each stack group of SAS shelves.



FibreBridge 6500N bridges are not supported in configurations running ONTAP 9.8 and later.

- FibreBridge 7600N or 7500N bridges support up to four SAS stacks.
- Each stack can use different models of IOM.

A mix of IOM12 modules and IOM3 modules is not supported within the same storage stack. A mix of IOM12 modules and IOM6 modules is supported within the same storage stack if your system is running a supported version of ONTAP.

Supported IOM modules depend on the version of ONTAP you are running.

• Naming must be unique within the MetroCluster configuration.

The suggested names used as examples in this documentation identify the controller module and stack that the bridge connects to, as shown below.

Pool and drive requirements (minimum supported)

Eight SAS disk shelves are recommended (four shelves at each site) to allow disk ownership on a per-shelf basis.

The MetroCluster configuration requires the minimum configuration at each site:

• Each node has at least one local pool and one remote pool at the site.

For example, in a four-node MetroCluster configuration with two nodes at each site, four pools are required at each site.

• At least seven drives in each pool.

In a four-node MetroCluster configuration with a single mirrored data aggregate per node, the minimum configuration requires 24 disks at the site.

In a minimum supported configuration, each pool has the following drive layout:

- Three root drives
- Three data drives
- One spare drive

In a minimum supported configuration, at least one shelf is needed per site.

MetroCluster configurations support RAID-DP and RAID4.

Drive location considerations for partially populated shelves

For correct auto-assignment of drives when using shelves that are half populated (12 drives in a 24-drive shelf), drives should be located in slots 0-5 and 18-23.

In a configuration with a partially populated shelf, the drives must be evenly distributed in the four quadrants of the shelf.

Mixing IOM12 and IOM 6 modules in a stack

Your version of ONTAP must support shelf mixing. Refer to the Interoperability Matrix Tool (IMT) to see if your version of ONTAP supports shelf mixing. IMT

For further details on shelf mixing see: Hot-adding shelves with IOM12 modules to a stack of shelves with IOM6 modules

Bridge naming conventions

The bridges use the following example naming:

```
bridge_site_stack grouplocation in pair
```

This portion of the name	Identifies the	Possible values
site	Site on which the bridge pair physically resides.	A or B
stack group	Number of the stack group to which the bridge pair connects. FibreBridge 7600N or 7500N bridges support up to four stacks in the stack group. The stack group can contain no more than 10 storage shelves.	1, 2, etc.
location in pair	Bridge within the bridge pair.A pair of bridges connect to a specific stack group.	a or b

Example bridge names for one stack group on each site:

- bridge_A_1a
- bridge_A_1b
- bridge_B_1a
- bridge_B_1b

Configuration worksheets for FC switches and FC-to-SAS bridges

Before beginning to configure the MetroCluster sites, you can use the following worksheets to record your site information:

Site A worksheet

Site B worksheet

Install and cable MetroCluster components

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 depends on the platform model of the controller modules, the switch types, and the number of disk shelf stacks in your configuration.

- 2. Properly ground yourself.
- 3. Install the controller modules in the rack or cabinet.

ONTAP Hardware Systems Documentation

- 4. Install the FC switches in the rack or cabinet.
- 5. Install the disk shelves, power them on, and then set the shelf IDs.
 - You must power-cycle each disk shelf.
 - Shelf IDs must be unique for each SAS disk shelf within each MetroCluster DR group (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.

The ATTO FibreBridge Installation and Operation Manual for your bridge model contains more information and an illustration of the installation.



For adequate port space access and FRU serviceability, you must leave 1U space below the bridge pair and cover this space with a tool-less blanking panel.

b. Connect each bridge to a power source that provides a proper ground.

c. Power on each bridge.



For maximum resiliency, bridges that are attached to the same stack of disk shelves must be 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 new controller module's FC-VI and HBA ports to the FC switches

The FC-VI ports and HBAs (host bus adapters) must be cabled to the site FC switches on each controller module in the MetroCluster configuration.

Steps

1. Cable the FC-VI ports and HBA ports, using the table for your configuration and switch model.

Port assignments for FC switches

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 must be done for both switch fabrics.

Steps

1. Connect the FC switches at each site to all ISLs, using the cabling in the table that corresponds to your configuration and switch model.

Port assignments for FC switches

Related information

Considerations for ISLs

Port assignments for FC switches

Port assignments for MetroCluster FC switches

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

You can reconfigure ports that are not used for attaching initiator ports, FC-VI ports, or ISLs to act as storage ports. However, if you are using the supported RCFs, you must change the zoning accordingly.

If you use the supported RCFs, ISL ports might not connect to the same ports shown and might need to be reconfigured manually.

If you configured your switches using the port assignments for ONTAP 9, you can continue to use the older assignments. However, new configurations running ONTAP 9.1 or later should use the port assignments shown here.

Overall cabling guidelines

You should be aware of the following guidelines when using the cabling tables:

- 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 cabling is the same for each FC switch in the switch fabric.
- You can order AFF A300 and FAS8200 storage systems with one of two options for FC-VI connectivity:
 - Onboard ports 0e and 0f configured in FC-VI mode.
 - Ports 1a and 1b on an FC-VI card in slot 1.
- AFF A700 and FAS9000 storage systems require four FC-VI ports. The following tables show cabling for the FC switches with four FC-VI ports on each controller except for the Cisco 9250i switch.

For other storage systems, use the cabling shown in the tables but ignore the cabling for FC-VI ports c and d.

You can leave those ports empty.

- AFF A400 and FAS8300 storage systems use ports 2a and 2b for FC-VI connectivity.
- If you have two MetroCluster configurations sharing ISLs, use the same port assignments as that for an eight-node MetroCluster cabling.
- The number of ISLs you cable can vary depending on site's requirements.
- See the section on ISL considerations.

Considerations for ISLs

AFF A900 and FAS9500 cabling guidelines

• AFF A900 or FAS9500 storage systems require eight FC-VI ports. If you are using an AFF A900 or FAS9500, you need to use the eight port configuration. If the configuration includes the other storage system models, use the cabling shown in the tables but ignore the cabling for unneeded FC-VI ports.

Port assignments for systems using two initiator ports

You can configure FAS8200 and AFF A300 systems using a single initiator port for each fabric and two initiator ports for each controller.

You can follow the cabling for the FibreBridge 7600N bridge using only one FC port (FC1 or FC2). Instead of using four initiators, connect only two initiators and leave the other two that are connected to the switch port empty.

If zoning is performed manually, then follow the zoning used for a FibreBridge 7600N bridge using one FC port (FC1 or FC2). In this scenario, one initiator port rather than two is added to each zone member per fabric.

You can change the zoning or perform an upgrade from a FibreBridge 6500N to a FibreBridge 7500N using the procedure in Hot-swapping a FibreBridge 6500N bridge with a FibreBridge 7600N or 7500N bridge.

The following table shows port assignments for Brocade FC switches when using a single initiator port for each fabric and two initiator ports for each controller.

Configurations using FibreBridge 7500N or 7600N using one FC port (FC1 or FC2) only					
MetroCluster 1 or DR Group 1					
Component	Port	6505, 6510, 6520, 7840, G620, G630, G610, G710, G720, G730 and DCX 8510-8			
		Connects to FC switch	Connects to switch port		
controller_x_1	FC-VI port a	1	0		
	FC-VI port b	2	0		
	FC-VI port c	1	1		
	FC-VI port d	2	1		
	HBA port a	1	2		
	HBA port b	2	2		
	HBA port c	-	-		
	HBA port d	-	-		
Stack 1	bridge_x_1a	1	8		
	bridge_x_1b	2	8		
Stack y	bridge_x_ya	1	11		
	bridge_x_yb	2	11		

Brocade port usage for controllers in a MetroCluster FC configuration

Learn about the port assignments required to cable Brocade FC switches to your controllers.

The following tables show the maximum supported configuration, with four controller modules per DR group. For smaller configurations, ignore the rows for the additional controller modules. Note that eight ISLs are supported only on the Brocade 6510, Brocade DCX 8510-8, G620, G630, G620-1, G630-1, G720, and G730 switches.

Review the following information before using these tables:

• Port usage for the Brocade 6505, G610, and G710 switches in an eight-node MetroCluster configuration is not shown. Due to the limited number of ports, port assignments must be made on a site-by-site basis

depending on the controller module model and the number of ISLs and bridge pairs in use.

- The Brocade DCX 8510-8 switch can use the same port layout as the 6510 switch *or* the 7840 switch.
- Brocade 6520, 7810, and 7840 switches aren't supported on systems that use eight FC-VI ports (AFF A900 and FAS9500 systems).
- Brocade 7810 switches only support one DR group.

MetroCluster 1 or DR group 1

The following table shows the supported controller configurations in MetroCluster 1 or DR group 1 on Brocade switches.

Compon ent	Port	Connect s to FC switch	6505, G610, G710 port	6510, DCX 8510-8 port	6520 port	7810 port	7840 port	G620, G620-1, G630, G630-1 port	G720, G730 port
controller _x_1	FC-VI port a	1	0	0	0	0	0	0	0
	FC-VI port b	2	0	0	0	0	0	0	0
	FC-VI port c	1	1	1	1	1	1	1	1
	FC-VI port d	2	1	1	1	1	1	1	1
	FC-VI-2 port a	1	16	20	n/a	n/a	n/a	16	2
	FC-VI-2 port b	2	16	20	n/a	n/a	n/a	16	2
	FC-VI-2 port c	1	17	21	n/a	n/a	n/a	17	3
	FC-VI-2 port d	2	17	21	n/a	n/a	n/a	17	3
	HBA port a	1	2	2	2	2	2	2	8
	HBA port b	2	2	2	2	2	2	2	8
	HBA port c	1	3	3	3	3	3	3	9
	HBA port d	2	3	3	3	3	3	3	9

Compon ent	Port	Connect s to FC switch	6505, G610, G710 port	6510, DCX 8510-8 port	6520 port	7810 port	7840 port	G620, G620-1, G630, G630-1 port	G720, G730 port
controller _x_2	FC-VI port a	1	4	4	4	4	4	4	4
	FC-VI port b	2	4	4	4	4	4	4	4
	FC-VI port c	1	5	5	5	5	5	5	5
	FC-VI port d	2	5	5	5	5	5	5	5
	FC-VI-2 port a	1	18	22	n/a	n/a	n/a	20	6
	FC-VI-2 port b	2	18	22	n/a	n/a	n/a	20	6
	FC-VI-2 port c	1	19	23	n/a	n/a	n/a	21	7
	FC-VI-2 port d	2	19	23	n/a	n/a	n/a	21	7
	HBA port a	1	6	6	6	6	6	6	12
	HBA port b	2	6	6	6	6	6	6	12
	HBA port c	1	7	7	7	7	7	7	13
	HBA port d	2	7	7	7	7	7	7	13

MetroCluster 2 or DR group 2

The following table shows the supported controller configurations in MetroCluster 2 or DR group 2 on Brocade switches.

Compon ent	Port	Connect s to FC switch	6505, G610, G710 port	6510, DCX 8510-8 port	6520 port	7810 port	7840 port	G620, G620-1, G630, G630-1 port	G720, G730 port
controller _x_3	FC-VI port a	1	n/a	24	48	n/a	12	18	18
	FC-VI port b	2	n/a	24	48	n/a	12	18	18
	FC-VI port c	1	n/a	25	49	n/a	13	19	19
	FC-VI port d	2	n/a	25	49	n/a	13	19	19
	FC-VI-2 port a	1	n/a	36	n/a	n/a	n/a	36	24
	FC-VI-2 port b	2	n/a	36	n/a	n/a	n/a	36	24
	FC-VI-2 port c	1	n/a	37	n/a	n/a	n/a	37	25
	FC-VI-2 port d	2	n/a	37	n/a	n/a	n/a	37	25
	HBA port a	1	n/a	26	50	n/a	14	24	26
	HBA port b	2	n/a	26	50	n/a	14	24	26
	HBA port c	1	n/a	27	51	n/a	15	25	27
	HBA port d	2	n/a	27	51	n/a	15	25	27

Compon ent	Port	Connect s to FC switch	6505, G610, G710 port	6510, DCX 8510-8 port	6520 port	7810 port	7840 port	G620, G620-1, G630, G630-1 port	G720, G730 port
controller _x_4	FC-VI port a	1	n/a	28	52	n/a	16	22	22
	FC-VI port b	2	n/a	28	52	n/a	16	22	22
	FC-VI port c	1	n/a	29	53	n/a	17	23	23
	FC-VI port d	2	n/a	29	53	n/a	17	23	23
	FC-VI-2 port a	1	n/a	38	n/a	n/a	n/a	38	28
	FC-VI-2 port b	2	n/a	38	n/a	n/a	n/a	38	28
	FC-VI-2 port c	1	n/a	39	n/a	n/a	n/a	39	29
	FC-VI-2 port d	2	n/a	39	n/a	n/a	n/a	39	29
	HBA port a	1	n/a	30	54	n/a	18	28	30
	HBA port b	2	n/a	30	54	n/a	18	28	30
	HBA port c	1	n/a	31	55	n/a	19	29	31
	HBA port d	2	n/a	31	55	n/a	19	29	31

MetroCluster 3 or DR group 3

The following table shows the supported controller configurations in MetroCluster 3 or DR group 3 on Brocade switches.

Component	Port	Connects to FC switch	G630, G630-1 port	G730 port
controller_x_5	FC-VI port a	1	48	48
	FC-VI port b	2	48	48
	FC-VI port c	1	49	49
	FC-VI port d	2	49	49
	FC-VI-2 port a	1	64	50
	FC-VI-2 port b	2	64	50
	FC-VI-2 port c	1	65	51
	FC-VI-2 port d	2	65	51
	HBA port a	1	50	56
	HBA port b	2	50	56
	HBA port c	1	51	57
	HBA port d	2	51	57

Component	Port	Connects to FC switch	G630, G630-1 port	G730 port
controller_x_6	FC-VI port a	1	52	52
	FC-VI port b	2	52	52
	FC-VI port c	1	53	53
	FC-VI port d	2	53	53
	FC-VI-2 port a	1	68	54
	FC-VI-2 port b	2	68	54
	FC-VI-2 port c	1	69	55
	FC-VI-2 port d	2	69	55
	HBA port a	1	54	60
	HBA port b	2	54	60
	HBA port c	1	55	61
	HBA port d	2	55	61

MetroCluster 4 or DR group 4

The following table shows the supported controller configurations in MetroCluster 4 or DR group 4 on Brocade switches.

Component	Port	Connects to FC switch	G630, G630-1 port	G730 port
controller_x_7	FC-VI port a	1	66	66
	FC-VI port b	2	66	66
	FC-VI port c	1	67	67
	FC-VI port d	2	67	67
	FC-VI-2 port a	1	84	72
	FC-VI-2 port b	2	84	72
	FC-VI-2 port c	1	85	73
	FC-VI-2 port d	2	85	73
	HBA port a	1	72	74
	HBA port b	2	72	74
	HBA port c	1	73	75
	HBA port d	2	73	75

Component	Port	Connects to FC switch	G630, G630-1 port	G730 port
controller_x_8	FC-VI port a	1	70	70
	FC-VI port b	2	70	70
	FC-VI port c	1	71	71
	FC-VI port d	2	71	71
	FC-VI-2 port a	1	86	76
	FC-VI-2 port b	2	86	76
	FC-VI-2 port c	1	87	77
	FC-VI-2 port d	2	87	77
	HBA port a	1	76	78
	HBA port b	2	76	78
	HBA port c	1	77	79
	HBA port d	2	77	79

Brocade port usage for FC-to-SAS bridges in a MetroCluster FC configuration

Learn about the port assignments required to cable Brocade FC switches to FC-to-SAS bridges. The port assignments vary depending on whether the bridges use one or two FC ports.



Brocade 7810 switches only support one DR group.

Shelf configurations using FibreBridge 7500N or 7600N bridges using both FC ports (FC1 and FC2)

MetroCluster 1 or DR group 1

The following table shows the supported shelf configurations in MetroCluster 1 or DR group 1 for FibreBridge 7500N or 7600N bridges using both FC ports (FC1 and FC2) on Brocade switches. You should be aware of the following when using this configuration table:

- On Brocade 6505, G610, G710, G620, G620-1, G630, and G630-1 switches, you can cable additional bridges to ports 12-15.
- On Brocade 6510 and DCX 8510-8 switches, you can cable additional bridges to ports 12-19.
- On Brocade 6520 switches, you can cable additional bridges to ports 12-21 and 24-45.
- On Brocade 7810 and 7840 switches, MetroCluster 1 or DR group 1 only supports two bridge stacks.
- On Brocade G720 and G730 switches, you can cable additional bridges to ports 16-21.

Component		Port	Connec ts to FC switch 	6505, G610, G710 port	6510, DCX 8510-8 port	6520 port	7810 port	7840 port	G620, G620-1, G630, G630-1 port	G720, G730 port
Stack 1 b	bridge_x _1a	FC1	1	8	8	8	8	8	8	10
		FC2	2	8	8	8	8	8	8	10
	bridge_x _1b	FC1	1	9	9	9	9	9	9	11
		FC2	2	9	9	9	9	9	9	11
Stack 2	bridge_x _2a	FC1	1	10	10	10	10	10	10	14
		FC2	2	10	10	10	10	10	10	14
	bridge_x _2b	FC1	1	11	11	11	11	11	11	15
		FC2	2	11	11	11	11	11	11	15

MetroCluster 2 or DR group 2

The following table shows the supported shelf configurations in MetroCluster 2 or DR group 2 for FibreBridge 7500N or 7600N bridges using both FC ports (FC1 and FC2) on Brocade switches. You should be aware of the following when using this configuration table:

- On Brocade 6510, DCX 8510-8, and 7840 switches, MetroCluster 2 or DR group 2 only supports two bridge stacks.
- On Brocade 6520 switches, you can cable additional bridges to ports 60-69 and 72-93.
- On Brocade G620, G620-1, G630, and G630-1 switches, you can cable additional bridges to ports 32-35.
- On Brocade G720 and G730 switches, you can cable additional bridges to ports 36-39.
- Port usage for the Brocade 6505, G610, and G710 switches in an eight-node MetroCluster configuration is not shown. Due to the limited number of ports, you assign ports on a site-by-site basis depending on the controller model and the number of ISLs and bridge pairs that you're using.

Component		Port	Connec ts to FC switch 	6505, G610, G710 port	6510, DCX 8510-8 port	6520 port	7810 port	7840 port	G620, G620-1, G630, G630-1 port	G720, G730 port
Stack 1 bridge _1a bridge _1b	bridge_x _1a	FC1	1	n/a	32	56	n/a	20	26	32
		FC2	2	n/a	32	56	n/a	20	26	32
	bridge_x _1b	FC1	1	n/a	33	57	n/a	21	27	33
		FC2	2	n/a	33	57	n/a	21	27	33
Stack 2	bridge_x _2a	FC1	1	n/a	34	58	n/a	22	30	34
		FC2	2	n/a	34	58	n/a	22	30	34
	bridge_x _2b	FC1	1	n/a	35	59	n/a	23	31	35
		FC2	2	n/a	35	59	n/a	23	31	35

MetroCluster 3 or DR group 3

The following table shows the supported shelf configurations in MetroCluster 3 or DR group 3 for FibreBridge 7500N or 7600N bridges using both FC ports (FC1 and FC2) on Brocade switches. You should be aware of the following when using this configuration table:

- On Brocade G630 and G630-1 switches, you can cable additional bridges to ports 60-63.
- On Brocade G730 switches, you can cable additional bridges to ports 64, 65, 68, and 69.

Component		Port	Connects to FC switch	G630, G630-1 port	G730 port
Stack 1	bridge_x_1a	FC1	1	56	58
		FC2	2	56	58
	bridge_x_1b	FC1	1	57	59
		FC2	2	57	59

Component		Port	Connects to FC switch	G630, G630-1 port	G730 port
Stack 2	bridge_x_2a	FC1	1	58	62
		FC2	2	58	62
	bridge_x_2b	FC1	1	59	63
		FC2	2	59	63

MetroCluster 4 or DR group 4

The following table shows the supported shelf configurations in MetroCluster 4 or DR group 4 for FibreBridge 7500N or 7600N bridges using both FC ports (FC1 and FC2) on Brocade switches. You should be aware of the following when using this configuration table:

- On Brocade G630 and G630-1 switches, you can cable additional bridges to ports 80-83.
- On Brocade G730 switches, you can cable additional bridges to ports 84-95.

Comp	oonent	Port	Connects to FC switch	G630, G630-1 port	G730 port
Stack 1	bridge_x_1a	FC1	1	74	80
		FC2	2	74	80
	bridge_x_1b	FC1	1	75	81
		FC2	2	75	81
Stack 2	bridge_x_2a	FC1	1	78	82
		FC2	2	78	82
	bridge_x_2b	FC1	1	79	83
		FC2	2	79	83

Shelf configurations using FibreBridge 7500N or 7600N using one FC port (FC1 or FC2) only

MetroCluster 1 or DR group 1

The following table shows the supported shelf configurations in MetroCluster 1 or DR group 1 using FibreBridge 7500N or 7600N and only one FC port (FC1 or FC2) on Brocade switches. You should be aware of the following when using this configuration table:

• On Brocade 6505, G610, G710, G620, G620-1, G630, and G630-1 switches, additional bridges ports 12-

15.

- On Brocade 6510 and DCX 8510-8 switches, you can cable additional bridges to ports 12-19.
- On Brocade 6520 switches, you can cable additional bridges to ports 16-21 and 24-45.
- On Brocade G720 and G730 switches, you can cable additional bridges to ports 16-21.

Compon ent	Port	Connect s to FC switch	6505, G610, G710 port	6510, DCX 8510-8 port	6520 port	7810 port	7840 port	G620, G620-1, G630, G630-1 port	G720, G730 port
Stack 1	bridge_x _1a	1	8	8	8	8	8	8	10
	bridge_x _1b	2	8	8	8	8	8	8	10
Stack 2	bridge_x _2a	1	9	9	9	9	9	9	11
	bridge_x _2b	2	9	9	9	9	9	9	11
Stack 3	bridge_x _3a	1	10	10	10	10	10	10	14
	bridge_x _3b	2	10	10	10	10	10	10	14
Stack 4	bridge_x _4a	1	11	11	11	11	11	11	15
	bridge_x _4b	2	11	11	11	11	11	11	15

MetroCluster 2 or DR group 2

The following table shows the supported shelf configurations in MetroCluster 2 or DR group 2 for FibreBridge 7500N or 7600N bridges using one FC port (FC1 or FC2) on Brocade switches. You should be aware of the following when using this configuration table:

- On Brocade 6520 switches, you can cable additional bridges to ports 60-69 and 72-93.
- On Brocade G620, G620-1, G630, G630-1 switches, you can cable additional bridges to ports 32-35.
- On Brocade G720 and G730 switches, you can cable additional bridges to ports 36-39.
- Port usage for the Brocade 6505, G610, and G710 switches in an eight-node MetroCluster configuration is not shown. Due to the limited number of ports, you assign ports on a site-by-site basis depending on the controller model and the number of ISLs and bridge pairs that you're using.

Compon ent	Port	Connect s to FC switch	6505, G610, G710 port	6510, DCX 8510-8 port	6520 port	7810 port	7840 port	G620, G620-1, G630, G630-1 port	G720, G730 port
Stack 1	bridge_x _1a	1	n/a	32	56	n/a	20	26	32
	bridge_x _1b	2	n/a	32	56	n/a	20	26	32
Stack 2	bridge_x _2a	1	n/a	33	57	n/a	21	27	33
	bridge_x _2b	2	n/a	33	57	n/a	21	27	33
Stack 3	bridge_x _3a	1	n/a	34	58	n/a	22	30	34
	bridge_x _3b	2	n/a	34	58	n/a	22	30	34
Stack 4	bridge_x _4a	1	n/a	35	59	n/a	23	31	35
	bridge_x _4b	2	n/a	35	59	n/a	23	31	35

MetroCluster 3 or DR group 3

The following table shows the supported shelf configurations in MetroCluster 3 or DR group 3 for FibreBridge 7500N or 7600N bridges using one FC port (FC1 or FC2) on Brocade switches. You should be aware of the following when using this configuration table:

- On Brocade G630 and G630-1 switches, you can cable additional bridges to ports 60-63.
- On Brocade G730 switches, you can cable additional bridges to ports 64, 65, 68, 69.

Component	Port	Connects to FC switch	G630, G630-1 port	G730 port
Stack 1	bridge_x_1a	1	56	58
	bridge_x_1b	2	56	58

Component	Port	Connects to FC switch	G630, G630-1 port	G730 port
Stack 2	bridge_x_2a	1	57	59
	bridge_x_2b	2	57	59
Stack 3	bridge_x_3a	1	58	62
	bridge_x_3b	2	58	62
Stack 4	bridge_x_4a	1	59	63
	bridge_x_4b	2	59	63

MetroCluster 4 or DR group 4

The following table shows the supported shelf configurations in MetroCluster 4 or DR group 4 for FibreBridge 7500N or 7600N bridges using one FC port (FC1 or FC2) on Brocade switches. You should be aware of the following when using this configuration table:

- On Brocade G630 and G630-1 switches, you can cable additional bridges to ports 80-83.
- On Brocade G730 switches, you can cable additional bridges to ports 84-95.

Component	Port	Connects to FC switch	G630, G630-1 port	G730 port
Stack 1	bridge_x_1a	1	74	80
	bridge_x_1b	2	74	80
Stack 2	bridge_x_2a	1	75	81
	bridge_x_2b	2	75	81
Stack 3	bridge_x_3a	1	78	82
	bridge_x_3b	2	78	82
Stack 4	bridge_x_4a	1	79	83
	bridge_x_4b	2	79	83

Brocade port usage for ISLs in a MetroCluster FC configuration

Learn about the port assignments required to cable Brocade FC switches to ISLs.

• AFF A900 and FAS9500 systems support eight ISLs. Eight ISLs are supported on the Brocade 6510, G620, G620-1, G630, G630-1, G720, and G730 switches.

 (\mathbf{i})

 Brocade 6520 switches supports eight ISLs, but do not support AFF A900 and FAS9500 systems.

ISL port	6505, G610, G710 port	6520 port	7810 port	7840 (10- Gbps) port	7840 (40- Gbps) port	6510, G620, G620-1, G630, G630- 1, G720, G730 port
ISL port 1	20	22	ge2	ge2	ge0	40
ISL port 2	21	23	ge3	ge3	ge1	41
ISL port 3	22	46	ge4	ge10	n/a	42
ISL port 4	23	47	ge5	ge11	n/a	43
ISL port 5	n/a	70	ge6	n/a	n/a	44
ISL port 6	n/a	71	ge7	n/a	n/a	45
ISL port 7	n/a	94	n/a	n/a	n/a	46
ISL port 8	n/a	95	n/a	n/a	n/a	47

Cisco port usage for controllers in a MetroCluster FC configuration

Learn about the port assignments required to cable Cisco 9124V, 9148S, 9148V, 9250i, and 9396S FC switches to your controllers.

The tables show the maximum supported configurations, with eight controller modules in two DR groups. For smaller configurations, ignore the rows for the additional controller modules.

- For Cisco 9132T, see Cisco 9132T port usage for controllers in a MetroCluster FC configuration.
- Cisco 9124V and 9250i switches are not supported in eight-node MetroCluster configurations.

MetroCluster 1 or DR group 1

The following table shows the supported controller configurations in MetroCluster 1 or DR group 1 on Cisco switches (excluding 9132T).

(i)

Component	Port	Connects to FC switch	9124V port	9148S port	9148V port	9250i port	9396S port
controller_x _1	FC-VI port a	1	1	1	1	1	1
	FC-VI port b	2	1	1	1	1	1
	FC-VI port c	1	2	2	2	2	2
	FC-VI port d	2	2	2	2	2	2
	FC-VI-2 port a	1	3	n/a	3	n/a	n/a
	FC-VI-2 port b	2	3	n/a	3	n/a	n/a
	FC-VI-2 port c	1	4	n/a	4	n/a	n/a
	FC-VI-2 port d	2	4	n/a	4	n/a	n/a
	HBA port a	1	13	3	13	3	3
	HBA port b	2	13	3	13	3	3
	HBA port c	1	14	4	14	4	4
	HBA port d	2	14	4	14	4	4

Component	Port	Connects to FC switch	9124V port	9148S port	9148V port	9250i port	9396S port
controller_x _2	FC-VI port a	1	5	5	5	5	5
	FC-VI port b	2	5	5	5	5	5
	FC-VI port c	1	6	6	6	6	6
	FC-VI port d	2	6	6	6	6	6
	FC-VI-2 port a	1	7	n/a	7	n/a	n/a
	FC-VI-2 port b	2	7	n/a	7	n/a	n/a
	FC-VI-2 port c	1	8	n/a	8	n/a	n/a
	FC-VI-2 port d	2	8	n/a	8	n/a	n/a
	HBA port a	1	15	7	15	7	7
	HBA port b	2	15	7	15	7	7
	HBA port c	1	16	8	16	8	8
	HBA port d	2	16	8	16	8	8

MetroCluster 2 or DR group 2

The following table shows the supported controller configurations in MetroCluster 2 or DR group 2 on Cisco switches (excluding 9132T).

Component	Port	Connects to FC switch	9124V port	9148S port	9148V port	9250i port	9396S port
controller_x _3	FC-VI port a	1	n/a	25	25	n/a	49
	FC-VI port b	2	n/a	25	25	n/a	49
	FC-VI port c	1	n/a	26	26	n/a	50
	FC-VI port d	2	n/a	26	26	n/a	50
	FC-VI-2 port a	1	n/a	n/a	27	n/a	n/a
	FC-VI-2 port b	2	n/a	n/a	27	n/a	n/a
	FC-VI-2 port c	1	n/a	n/a	28	n/a	n/a
	FC-VI-2 port d	2	n/a	n/a	28	n/a	n/a
	HBA port a	1	n/a	27	37	n/a	51
	HBA port b	2	n/a	27	37	n/a	51
	HBA port c	1	n/a	28	38	n/a	52
	HBA port d	2	n/a	28	38	n/a	52

Component	Port	Connects to FC switch	9124V port	9148S port	9148V port	9250i port	9396S port
controller_x _4	FC-VI port a	1	n/a	29	29	n/a	53
	FC-VI port b	2	n/a	29	29	n/a	53
	FC-VI port c	1	n/a	30	30	n/a	54
	FC-VI port d	2	n/a	30	30	n/a	54
	FC-VI-2 port a	1	n/a	n/a	31	n/a	n/a
	FC-VI-2 port b	2	n/a	n/a	31	n/a	n/a
	FC-VI-2 port c	1	n/a	n/a	32	n/a	n/a
	FC-VI-2 port d	2	n/a	n/a	32	n/a	n/a
	HBA port a	1	n/a	31	39	n/a	55
	HBA port b	2	n/a	31	39	n/a	55
	HBA port c	1	n/a	32	40	n/a	56
	HBA port d	1	n/a	32	40	n/a	56

Cisco port usage for FC-to-SAS bridges in a MetroCluster FC configuration

Learn about the port assignments required to cable Cisco 9124V, 9148S, 9148V, 9250i, and 9396S FC switches to FC-to-SAS bridges. The port assignments vary depending on whether the bridges use one or two FC ports.



For Cisco 9132T, see Cisco 9132t port usage for FC-to-SAS bridges in a MetroCluster FC configuration.

Shelf configurations using FibreBridge 7500N or 7600N using both FC ports (FC1 and FC2)

MetroCluster 1 or DR group 1

The following table shows the supported shelf configurations in MetroCluster 1 or DR group 1 using FibreBridge 7500N or 7600N bridges and both FC ports (FC1 and FC2) on Cisco switches (excluding 9132T). You should be aware of the following when using this configuration table:

- On Cisco 9250i switches, you can cable additional MetroCluster 1 or DR group 1 bridges to ports 17-40.
- On Cisco 9396S switches, you can cable additional MetroCluster 1 or DR group 1 bridges to ports 17-32.

Comp	oonent	Port	Connects to FC switch	9124V port	9148S port	9148V port	9250i port	9396S port
Stack 1	bridge_x_ 1a	FC1	1	17	9	17	9	9
		FC2	2	17	9	17	9	9
	bridge_x_ 1b	FC1	1	18	10	18	10	10
		FC2	2	18	10	18	10	10
Stack 2	Stack 2 bridge_x_ 2a	FC1	1	19	11	19	11	11
bridg 2b		FC2	2	19	11	19	11	11
	bridge_x_ 2b	FC1	1	20	12	20	12	12
		FC2	2	20	12	20	12	12
Stack 3	bridge_x_ 3a	FC1	1	21	13	21	13	13
		FC2	2	21	13	21	13	13
	bridge_x_ 3b	FC1	1	22	14	22	14	14
		FC2	2	22	14	22	14	14
Stack 4	bridge_x_ 4a	FC1	1	23	15	23	15	15
		FC2	2	23	15	23	15	15
	bridge_x_ 4b	FC1	1	24	16	24	16	16
		FC2	2	24	16	24	16	16

MetroCluster 2 or DR group 2

The following table shows the supported shelf configurations in MetroCluster 2 or DR group 2 using FibreBridge 7500N or 7600N and both FC ports (FC1 and FC2) on Cisco switches (excluding 9132T). You should be aware of the following when using the cabling tables:

• Cisco 9124V and 9250i switches are not supported for eight-node MetroCluster configurations.

• On Cisco 9396S switches, you can cable additional MetroCluster 2 (DR group 2) bridges to ports 65-80.

Comp	oonent	Port	Connects to FC switch	9124V port	9148S port	9148V port	9250i port	9396S port
Stack 1	bridge_x_ 1a	FC1	1	n/a	33	41	n/a	57
		FC2	2	n/a	33	41	n/a	57
	bridge_x_ 1b	FC1	1	n/a	34	42	n/a	58
		FC2	2	n/a	34	42	n/a	58
Stack 2 bridge_x_ 2a bridge_x_ 2b	FC1	1	n/a	35	43	n/a	59	
		FC2	2	n/a	35	43	n/a	59
	bridge_x_ 2b	FC1	1	n/a	36	44	n/a	60
		FC2	2	n/a	36	44	n/a	60
Stack 3	bridge_x_ 3a	FC1	1	n/a	37	45	n/a	61
		FC2	2	n/a	37	45	n/a	61
	bridge_x_ 3b	FC1	1	n/a	38	46	n/a	62
		FC2	2	n/a	38	46	n/a	62
Stack 4	bridge_x_ 4a	FC1	1	n/a	39	47	n/a	63
		FC2	2	n/a	39	47	n/a	63
	bridge_x_ 4b	FC1	1	n/a	40	48	n/a	64
		FC2	2	n/a	40	48	n/a	64

Shelf configurations using FibreBridge 7500N or 7600N using one FC port (FC1 or FC2) only

MetroCluster 1 or DR group 1

The following table shows the supported shelf configurations in MetroCluster 1 or DR group 1 for FibreBridge 7500N or 7600N bridges using one FC port (FC1 or FC2) on Cisco switches (excluding 9132T). The reference configuration file (RCF) doesn't support one FC port on FibreBridge bridges, so you must configure the backend fibre channel switches manually.

Configure the Cisco FC switches manually

You should be aware of the following when using the cabling tables:

- On Cisco 9250i switches, you can cable additional MetroCluster 1 or DR group 1 bridges to ports 17-40.
- On Cisco 9396S switches, you can cable additional MetroCluster 1 or DR group 1 bridges to ports 17-32.

Component	Port	Connects to FC switch	9124V port	9148S port	9148V port	9250i port	9396S port
Stack 1	bridge_x_1a	1	17	9	17	9	9
	bridge_x_1b	2	17	9	17	9	9
Stack 2	bridge_x_2a	1	18	10	18	10	10
	bridge_x_2b	2	18	10	18	10	10
Stack 3	bridge_x_3a	1	19	11	19	11	11
	bridge_x_3b	2	19	11	19	11	11
Stack 4	bridge_x_4a	1	20	12	20	12	12
	bridge_x_4b	2	20	12	20	12	12
Stack 5	bridge_x_5a	1	21	13	21	13	13
	bridge_x_5b	2	21	13	21	13	13
Stack 6	bridge_x_6a	1	22	14	22	14	14
	bridge_x_6b	2	22	14	22	14	14
Stack 7	bridge_x_7a	1	23	15	23	15	15
	bridge_x_7b	2	23	15	23	15	15
Stack 8	bridge_x_8a	1	24	16	24	16	16
	bridge_x_8b	2	24	16	24	16	16

MetroCluster 2 or DR group 2

The following table shows the supported shelf configurations in MetroCluster 2 or DR group 2 for FibreBridge 7500N or 7600N bridges using one FC port (FC1 or FC2) on Cisco switches (excluding 9132T). You should be aware of the following when using this configuration table:

- The Cisco 9124V and 9250i switches are not supported for eight-node MetroCluster configurations.
- On Cisco 9396S switches, you can cable additional MetroCluster 2 or DR group 2 bridges to ports 65-80.

Component	Port	Connects to FC switch	9124V port	9148S port	9148V port	9250i port	9396S port
Stack 1	bridge_x_1a	1	n/a	33	41	n/a	57
	bridge_x_1b	2	n/a	33	41	n/a	57
Stack 2	bridge_x_2a	1	n/a	34	42	n/a	58
	bridge_x_2b	2	n/a	34	42	n/a	58
Stack 3	bridge_x_3a	1	n/a	35	43	n/a	59
	bridge_x_3b	2	n/a	35	43	n/a	59
Stack 4	bridge_x_4a	1	n/a	36	44	n/a	60
	bridge_x_4b	2	n/a	36	44	n/a	60
Stack 5	bridge_x_5a	1	n/a	37	45	n/a	61
	bridge_x_5b	2	n/a	37	45	n/a	61
Stack 6	bridge_x_6a	1	n/a	38	46	n/a	62
	bridge_x_6b	2	n/a	38	46	n/a	62
Stack 7	bridge_x_7a	1	n/a	39	47	n/a	63
	bridge_x_7b	2	n/a	39	47	n/a	63
Stack 8	bridge_x_8a	1	n/a	40	48	n/a	64
	bridge_x_8b	2	n/a	40	48	n/a	64

Cisco port usage for ISLs in a MetroCluster FC configuration

Learn about the port assignments required to cable Cisco 9124V, 9148S, 9148V, 9250i, and 9396S FC switches to ISLs.

The following table shows ISL port usage. ISL port usage is the same on all switches in the configuration.

• For Cisco 9132T, see Cisco 9132T port usage for ISLs in a MetroCluster FC configuration.



The Cisco 9250i switch requires a 24 port license.

ISL port	9124V port	9148S port	9148V port	9250i port	9396S port
ISL port 1	9	20	9	12	44
ISL port 2	10	24	10	16	48
ISL port 3	11	44	11	20	92
ISL port 4	12	48	12	24	96
ISL port 5	n/a	n/a	33	n/a	n/a
ISL port 6	n/a	n/a	34	n/a	n/a
ISL port 7	n/a	n/a	35	n/a	n/a
ISL port 8	n/a	n/a	36	n/a	n/a

Cisco 9132T port usage for controllers in a MetroCluster FC configuration

Learn about the port assignments required to cable Cisco 9132T FC switches to your controllers.

The following table shows controller configurations using FibreBridge 7500N or 7600N using both FC ports (FC1 and FC2). The tables show the maximum supported configurations with four and eight controller modules in two DR groups.



For eight-node configurations, you must perform the zoning manually because RCFs are not provided.

MetroCluster 1 or DR group 1

The following table shows the supported controller configurations for MetroCluster 1 or DR group 1 on Cisco 9132T switches. You should be aware of the following when using this configuration table:

• AFF A900 and FAS9500 systems have eight FC-VI ports (a, b, c, and d for FC-VI-1 and FC-VI-2).

Component	Port	Connects to FC_switch	9132T 1x LEM (Four- node)	9132T 2x LEM (Four- node)	9132T 2x LEM (Eight- node)
controller_x_1	FC-VI port a	1	LEM1-1	LEM1-1	LEM1-1
	FC-VI port b	2	LEM1-1	LEM1-1	LEM1-1
	FC-VI port c	1	LEM1-2	LEM1-2	LEM1-2
	FC-VI port d	2	LEM1-2	LEM1-2	LEM1-2
	FC-VI-2 port a	1	LEM1-3	LEM1-3	n/a
	FC-VI-2 port b	2	LEM1-3	LEM1-3	n/a
	FC-VI-2 port c	1	LEM1-4	LEM1-4	n/a
	FC-VI-2 port d	2	LEM1-4	LEM1-4	n/a
	HBA port a	1	LEM1-5	LEM1-5	LEM1-3
	HBA port b	2	LEM1-5	LEM1-5	LEM1-3
	HBA port c	1	LEM1-6	LEM1-6	LEM1-4
	HBA port d	2	LEM1-6	LEM1-6	LEM1-4

Component	Port	Connects to FC_switch	9132T 1x LEM (Four- node)	9132T 2x LEM (Four- node)	9132T 2x LEM (Eight- node)
controller_x_2	FC-VI port a	1	LEM1-7	LEM1-7	LEM1-5
	FC-VI port b	2	LEM1-7	LEM1-7	LEM1-5
	FC-VI port c	1	LEM1-8	LEM1-8	LEM1-6
	FC-VI port d	2	LEM1-8	LEM1-8	LEM1-6
	FC-VI-2 port a	1	LEM1-9	LEM1-9	n/a
	FC-VI-2 port b	2	LEM1-9	LEM1-9	n/a
	FC-VI-2 port c	1	LEM1-10	LEM1-10	n/a
	FC-VI-2 port d	2	LEM1-10	LEM1-10	n/a
	HBA port a	1	LEM1-11	LEM1-11	LEM1-7
	HBA port b	2	LEM1-11	LEM1-11	LEM1-7
	HBA port c	1	LEM1-12	LEM1-12	LEM1-8
	HBA port d	2	LEM1-12	LEM1-12	LEM1-8

MetroCluster 2 or DR group 2

The following table shows the supported Cisco 9132T controller configurations for MetroCluster 2 or DR group 2 on Cisco 9132T switches. You should be aware of the following when using this configuration table:

- AFF A900 and FAS9500 systems have eight FC-VI ports (a, b, c, and d for FC-VI-1 and FC-VI-2).
- MetroCluster 2 or DR group 2 is not supported on Cisco 9132T switches for AFF A900 and FAS9500 systems.
- MetroCluster 2 or DR group 2 is only supported in eight-node MetroCluster configurations

Component	Port	Connects to FC_switch	9132T 1x LEM (Four- node)	9132T 2x LEM (Four- node)	9132T 2x LEM (Eight- node)
controller_x_3	FC-VI port a	1	n/a	n/a	LEM2-1
	FC-VI port b	2	n/a	n/a	LEM2-1
	FC-VI port c	1	n/a	n/a	LEM2-2
	FC-VI port d	2	n/a	n/a	LEM2-2
	FC-VI-2 port a	1	n/a	n/a	n/a
	FC-VI-2 port b	2	n/a	n/a	n/a
	FC-VI-2 port c	1	n/a	n/a	n/a
	FC-VI-2 port d	2	n/a	n/a	n/a
	HBA port a	1	n/a	n/a	LEM2-3
	HBA port b	2	n/a	n/a	LEM2-3
	HBA port c	1	n/a	n/a	LEM2-4
	HBA port d	2	n/a	n/a	LEM2-4

Component	Port	Connects to FC_switch	9132T 1x LEM (Four- node)	9132T 2x LEM (Four- node)	9132T 2x LEM (Eight- node)
controller_x_4	FC-VI-1 port a	1	n/a	n/a	LEM2-5
	FC-VI-1 port b	2	n/a	n/a	LEM2-5
	FC-VI-1 port c	1	n/a	n/a	LEM2-6
	FC-VI-1 port d	2	n/a	n/a	LEM2-6
	FC-VI-2 port a	1	n/a	n/a	n/a
	FC-VI-2 port b	2	n/a	n/a	n/a
	FC-VI-2 port c	1	n/a	n/a	n/a
	FC-VI-2 port d	2	n/a	n/a	n/a
	HBA port a	1	n/a	n/a	LEM2-7
	HBA port b	2	n/a	n/a	LEM2-7
	HBA port c	1	n/a	n/a	LEM2-8
	HBA port d	2	n/a	n/a	LEM2-8

Cisco 9132T port usage for FC-to-SAS bridges in a MetroCluster FC configuration

Learn about the port assignments required to cable Cisco 9132T FC switches to FC-to-SAS bridges using both FC ports.

Only one (1) bridge stack is supported using Cisco 9132T switches with 1xLEM Module.

MetroCluster 1 or DR group 1

(;`

The following table shows the supported shelf configurations in MetroCluster 1 or DR group 1 for FibreBridge 7500N or 7600N bridges using both FC ports (FC1 and FC2) on Cisco 9132T switches. You should be aware of the following when using this configuration table:

• In four-node configurations, you can cable additional bridges to ports LEM2-1 through LEM2-8 on Cisco 9132T switches with 2xLEMs.

Component		Port	Connects to FC_switch	9132T 1x LEM (Four- node)	9132T 2x LEM (Four- node)	9132T 2x LEM (Eight- node)
Stack 1	bridge_x_1a	FC1	1	LEM1-13	LEM1-13	LEM1-9
		FC2	2	LEM1-13	LEM1-13	LEM1-9
	bridge_x_1b	FC1	1	LEM1-14	LEM1-14	LEM1-10
		FC2	2	LEM1-14	LEM1-14	LEM1-10
Stack 2	bridge_x_2a	FC1	1	n/a	LEM1-15	LEM1-11
		FC2	2	n/a	LEM1-15	LEM1-11
	bridge_x_2b	FC1	1	n/a	LEM1-16	LEM1-12
		FC2	2	n/a	LEM1-16	LEM1-12

MetroCluster 2 or DR group 2

The following table shows the supported shelf configurations in MetroCluster 2 or DR group 2 for FibreBridge 7500N or 7600N bridges using both FC ports (FC1 and FC2) on Cisco 9132T switches. You should be aware of the following when using this configuration table:

• In eight-node configurations, you can cable additional bridges to ports LEM2-13 through LEM2-16 on Cisco 9132T switches with 2x LEMs.

Component		Port	Connects to FC_switch	9132T 1x LEM (Four- node)	9132T 2x LEM (Four- node)	9132T 2x LEM (Eight- node)
Stack 1	bridge_x_1a	FC1	1	n/a	n/a	LEM1-9
		FC2	2	n/a	n/a	LEM1-9
	bridge_x_1b	FC1	1	n/a	n/a	LEM1-10
		FC2	2	n/a	n/a	LEM1-10

Component		Port	Connects to FC_switch	9132T 1x LEM (Four- node)	9132T 2x LEM (Four- node)	9132T 2x LEM (Eight- node)
Stack 2	bridge_x_2a	FC1	1	n/a	n/a	LEM1-11
		FC2	2	n/a	n/a	LEM1-11
	bridge_x_2b	FC1	1	n/a	n/a	LEM1-12
		FC2	2	n/a	n/a	LEM1-12

Cisco 9132T port usage for ISLs in a MetroCluster FC configuration

Learn about the port assignments required to cable Cisco 9132T FC switches to ISLs.

The following table shows ISL port usage for a Cisco 9132T switch.

ISL port	9132T 1x LEM (Four- node)	9132T 2x LEM (Four- node)	9132T 2x LEM (Eight- node)
ISL port 1	LEM1-15	LEM2-9	LEM1-13
ISL port 2	LEM1-16	LEM2-10	LEM1-14
ISL port 3	n/a	LEM2-11	LEM1-15
ISL port 4	n/a	LEM2-12	LEM1-16
ISL port 5	n/a	LEM2-13	n/a
ISL port 6	n/a	LEM2-14	n/a
ISL port 7	n/a	LEM2-15	n/a
ISL port 8	n/a	LEM2-16	n/a

Cabling the cluster interconnect in eight- or four-node configurations

In eight-node or four-node MetroCluster configurations, you must cable the cluster interconnect between the local controller modules at each site.

About this task

This task is not required on two-node MetroCluster configurations.

This task must be performed at both MetroCluster sites.

Step

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

Related information

ONTAP Hardware Systems Documentation

Network and LIF management

Cabling the cluster peering connections

You must cable the controller module ports used for cluster peering so that they have connectivity with the cluster on the partner site.

About this task

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

At least two ports on each controller module should be used for cluster peering.

The recommended minimum bandwidth for the ports and network connectivity is 1 GbE.

Step

1. Identify and cable at least two ports for cluster peering and verify they have network connectivity with the partner cluster.

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

Related information

Cluster and SVM peering express configuration

Each MetroCluster site is configured as a peer to its partner site. You should be familiar with the prerequisites and guidelines for configuring the peering relationships and when deciding whether to use shared or dedicated ports for those relationships.

Cluster peering

Cabling the HA interconnect

If you have an eight- or a four-node MetroCluster configuration and the storage controllers within the HA pairs are in separate chassis, you must cable the HA interconnect between the controllers.

About this task

- This task does not apply to two-node MetroCluster configurations.
- This task must be performed at both MetroCluster sites.
- The HA interconnect must be cabled only if the storage controllers within 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.

ONTAP Hardware Systems Documentation

- 2. If the MetroCluster site includes two HA pairs, repeat the previous steps on the second HA pair.
- 3. Repeat this task at the MetroCluster partner site.

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 new controller at both MetroCluster sites.

You can connect the controller and cluster switch management ports to existing switches in your network or to new dedicated network switches such as NetApp CN1601 cluster management switches.

Step

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

ONTAP Hardware Systems Documentation

Configure the FC switches

FC switch configuration overview

You can configure Cisco and Brocade FC switches by using RCF files, or, if necessary, you can manually configure the switches.

If you	Use the procedure
Have an RCF that meets your requirements	 Configure Brocade FC switches with RCF files Configure Cisco FC switches with RCF files
Do not have an RCF or have an RCF that does not meet your requirements	Configure the Brocade FC switches manuallyConfigure the Cisco FC switches manually

Configure Brocade FC switches with RCF files

Resetting the Brocade FC switch to factory defaults

Before installing a new software version and RCF files, you must erase the current switch configuration and perform basic configuration.

About this task

You must repeat these steps on each of the FC switches in the MetroCluster fabric configuration.

Steps

- 1. Log in to the switch as an administrator.
- 2. Disable the Brocade Virtual Fabrics (VF) feature:

fosconfig options

```
FC_switch_A_1:admin> fosconfig --disable vf
WARNING: This is a disruptive operation that requires a reboot to take
effect.
Would you like to continue [Y/N]: y
```

- 3. Disconnect the ISL cables from the ports on the switch.
- 4. Disable the switch:

switchcfgpersistentdisable

FC_switch_A_1:admin> switchcfgpersistentdisable

5. Disable the configuration:

cfgDisable

```
FC_switch_A_1:admin> cfgDisable
You are about to disable zoning configuration. This action will disable
any previous zoning configuration enabled.
Do you want to disable zoning configuration? (yes, y, no, n): [no] y
Updating flash ...
Effective configuration is empty. "No Access" default zone mode is ON.
```

6. Clear the configuration:

cfgClear

FC_switch_A_1:admin> cfgClear
The Clear All action will clear all Aliases, Zones, FA Zones
and configurations in the Defined configuration.
Run cfgSave to commit the transaction or cfgTransAbort to
cancel the transaction.
Do you really want to clear all configurations? (yes, y, no, n): [no] y

7. Save the configuration:

cfgSave

FC_switch_A_1:admin> cfgSave
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] y
Updating flash ...

8. Set the default configuration:

configDefault

FC switch A 1:admin> configDefault WARNING: This is a disruptive operation that requires a switch reboot. Would you like to continue [Y/N]: y Executing configdefault...Please wait 2020/10/05-08:04:08, [FCR-1069], 1016, FID 128, INFO, FC switch A 1, The FC Routing service is enabled. 2020/10/05-08:04:08, [FCR-1068], 1017, FID 128, INFO, FC switch A 1, The FC Routing service is disabled. 2020/10/05-08:04:08, [FCR-1070], 1018, FID 128, INFO, FC_switch_A_1, The FC Routing configuration is set to default. Committing configuration ... done. 2020/10/05-08:04:12, [MAPS-1113], 1019, FID 128, INFO, FC switch A 1, Policy dflt conservative policy activated. 2020/10/05-08:04:12, [MAPS-1145], 1020, FID 128, INFO, FC switch A 1, FPI Profile dflt fpi profile is activated for E-Ports. 2020/10/05-08:04:12, [MAPS-1144], 1021, FID 128, INFO, FC switch A 1, FPI Profile dflt fpi profile is activated for F-Ports. The switch has to be rebooted to allow the changes to take effect. 2020/10/05-08:04:12, [CONF-1031], 1022, FID 128, INFO, FC switch A 1, configDefault completed successfully for switch.

9. Set the port configuration to default for all ports:

portcfgdefault port-number

FC_switch_A_1:admin> portcfgdefault <port number>

You must complete this step for each port.

10. If you are running a version earlier than FOS 9.0, verify that the switch is using the dynamic Port on Demand (POD) method.



In Fabric OS 9.0 and later, the license method is dynamic by default. The static license method is not supported.

For Brocade Fabric OS versions before 8.0, you run the following commands as admin, and for versions 8.0 and later, you run them as root.

a. Run the license command:

```
licenseport --show.
```

FC_switch_A_1:admin> license --show -port
24 ports are available in this switch
Full POD license is installed
Dynamic POD method is in use

b. Enable the root user if it is disabled by Brocade.

```
FC_switch_A_1:admin> userconfig --change root -e yes
FC_switch_A_1:admin> rootaccess --set consoleonly
```

c. Run the license command:

licenseport --show.

FC_switch_A_1:root> license --show -port
24 ports are available in this switch
Full POD license is installed
Dynamic POD method is in use

d. If you are running Fabric OS 8.2.x and earlier, you must change the license method to dynamic:

licenseport --method dynamic

FC_switch_A_1:admin> licenseport --method dynamic The POD method has been changed to dynamic. Please reboot the switch now for this change to take effect

11. Reboot the switch:

fastBoot

```
FC_switch_A_1:admin> fastboot
Warning: This command would cause the switch to reboot
and result in traffic disruption.
Are you sure you want to reboot the switch [y/n]?y
```

12. Confirm that the default settings have been implemented:

switchShow

13. Verify that the IP address is set correctly:

ipAddrShow

You can set the IP address with the following command, if required:

ipAddrSet

Downloading the Brocade FC switch RCF file

You must download the reference configuration (RCF) file to each switch in the MetroCluster fabric configuration.

About this task

To use these RCF files, the system must be running ONTAP 9.1 or later and you must use the port layout for ONTAP 9.1 or later.

If you are planning to use only one of the FC ports on the FibreBridge bridges, configure the back-end fibre channel switches manually using the instructions found in the section, Port assignments for FC switches.

Steps

1. Refer to the RCF file table on the Brocade RCF download page and identify the correct RCF file for each switch in your configuration.

The RCF files must be applied to the correct switches.

2. Download the RCF files for the switches from the MetroCluster RCF download page.

The files must be placed in a location where they can be transferred to the switch. There is a separate file for each of the four switches that make up the two-switch fabric.

3. Repeat these steps on each switch in the configuration.

Installing the Brocade FC switch RCF file

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

About this task

- You must repeat these steps on each of the Brocade FC switches in the MetroCluster fabric configuration.
- If you use an xWDM configuration, you might require additional settings on the ISLs. See the xWDM

vendor documentation for more information.

Steps

1. Initiate the download and configuration process:

configDownload

Respond to the prompts as shown in the following example.

```
FC_switch_A_1:admin> configDownload
Protocol (scp, ftp, sftp, local) [ftp]:
Server Name or IP Address [host]: <user input>
User Name [user]:<user input>
Path/Filename [<home dir>/config.txt]:path to configuration file
Section (all|chassis|switch [all]): all
.
.
Do you want to continue [y/n]: y
Password: <user input>
```

After entering your password, the switch downloads and executes the configuration file.

2. Confirm that the configuration file has set the switch domain:

switchShow

Each switch is assigned a different domain number depending on which configuration file the switch used.

```
FC_switch_A_1:admin> switchShow
switchName: FC_switch_A_1
switchType: 109.1
switchState: Online
switchMode: Native
switchRole: Subordinate
switchDomain: 5
```

3. Verify that your switch is assigned the correct domain value as indicated in the following table.

Fabric	Switch	Switch domain
1	A_1	5
	B_1	7

2	A_2	6
	B_2	8

4. Change the port speed:

portcfgspeed

FC_switch_A_1:admin> portcfgspeed port number port speed

By default, all the ports are configured to operate at 16 Gbps. You might change the port speed for the following reasons:

- The interconnect switch ports speed should be changed when an 8-Gbps FC-VI adapter is used and the switch port speed should set to 8 Gbps.
- The ISL ports' speed must be changed when the ISL is not capable of running at 16 Gbps.
- 5. Calculate the ISL distance.

Due to the behavior of the FC-VI, you must set the distance 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: $1.5 \times$ real distance = distance.

If the distance is 3 km, then 1.5×3 km = 4.5. This is lower than 10; therefore, you must set the ISL to the LE distance level.

The distance is 20 km, then 1.5 × 20 km = 30. You must set the ISL to the LS distance level.

6. Set the distance for each ISL port:

portcfglongdistance port level vc_link_init -distance distance_value

A vc_link_init value of 1 uses the fillword "ARB" by default. A value of 0 uses the fillword "IDLE". The required value might vary depending on the link you use. In this example, the default is set and the distance is assumed to be 20 km. Hence, the setting is "30" with a vc_link_init value of "1", and the ISL port is "21".

Example: LS

```
FC_switch_A_1:admin> portcfglongdistance 21 LS 1 -distance 30
```

Example: LE

FC_switch_A_1:admin> portcfglongdistance 21 LE 1

7. Persistently enable the switch:

switchcfgpersistentenable

The example shows how to persistently enable FC switch_A_1.

FC_switch_A_1:admin> switchcfgpersistentenable

8. Verify if the IP address is set correctly:

ipAddrshow

FC_switch_A_1:admin> ipAddrshow

You can set the IP address, if required:

ipAddrSet

9. Set the timezone from the switch prompt:

```
tstimezone --interactive
```

You should respond to the prompts as required.

FC switch A 1:admin> tstimezone --interactive

10. Reboot the switch:

reboot

The example shows how to reboot FC switch _A_1.

FC_switch_A_1:admin> reboot

11. Verify the distance setting:

portbuffershow

A distance setting of LE appears as 10 km.

FC_Switch_A_1:admin> portbuffershow							
User	Port	Lx	Max/Resv	Buffer	Needed	Link	Remaining
Port	Туре	Mode	Buffers	Usage	Buffers	Distance	Buffers
•••							
21	Ε	-	8	67	67	30 km	
22	Ε	-	8	67	67	30 km	
23	-	8	0	-	-	466	

12. Reconnect the ISL cables to the ports on the switches where they were removed.

The ISL cables were disconnected when the factory settings were reset to the default settings.

Resetting the Brocade FC switch to factory defaults

- 13. Validate the configuration.
 - a. Verify that the switches form one fabric:

switchshow

The following example shows the output for a configuration that uses ISLs on ports 20 and 21.

```
FC switch A 1:admin> switchshow
switchName: FC_switch_A_1
switchType: 109.1
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
_____
. . .
    20 010C00 id 16G Online FC LE E-Port
20
10:00:00:05:33:8c:2e:9a "FC_switch_B_1" (downstream)(trunk master)
21
    21 010D00 id 16G Online FC LE E-Port (Trunk port,
master is Port 20)
. . .
```

b. Confirm the configuration of the fabrics:

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"

c. Verify that the ISLs are working:

islshow

FC switch A 1:admin> islshow

d. Confirm that zoning is properly replicated:

cfgshow zoneshow

Both outputs should show the same configuration information and zoning information for both switches.

e. If trunking is used, confirm the trunking:

trunkShow

FC_switch_A_1:admin> trunkshow

Configure the Cisco FC switches with RCF files

Resetting the Cisco FC switch to factory defaults

Before installing a new software version and RCFs, you must erase the Cisco switch configuration and perform basic configuration.

About this task

You must repeat these steps on each of the FC switches in the MetroCluster fabric configuration.



The outputs shown are for Cisco IP switches; however, these steps are also applicable for Cisco FC switches.

Steps

- 1. Reset the switch to factory defaults:
 - a. Erase the existing configuration:

write erase

b. Reload the switch software:

reload

The system reboots and enters the configuration wizard. During the boot, if you receive the prompt Abort Auto Provisioning and continue with normal setup?(yes/no)[n], you should respond **yes** to proceed.

- c. In the configuration wizard, enter the basic switch settings:
 - Admin password
 - Switch name
 - Out-of-band management configuration
 - Default gateway
 - SSH service (Remote Support Agent).

After completing the configuration wizard, the switch reboots.

d. When prompted, enter the user name and password to log in to the switch.

The following example shows the prompts and system responses when logging in to the switch. The angle brackets (<<<) show where you enter the information.

```
---- System Admin Account Setup ----
Do you want to enforce secure password standard (yes/no) [y]:y
**<<<**
Enter the password for "admin": password **<<<**
Confirm the password for "admin": password **<<<**
---- Basic System Configuration Dialog VDC: 1 ----
```

This setup utility will guide you through the basic configuration of the system. Setup configures only enough connectivity for management of the system.

Please register Cisco Nexus3000 Family devices promptly with your supplier. Failure to register may affect response times for initial service calls. Nexus3000 devices must be registered to receive entitled support services.

Press Enter at anytime to skip a dialog. Use ctrl-c at anytime to skip the remaining dialogs.

e. Enter basic information in the next set of prompts, including the switch name, management address, and gateway, and enter **rsa** for the SSH key as shown in the example:

```
Would you like to enter the basic configuration dialog (yes/no): yes
 Create another login account (yes/no) [n]:
 Configure read-only SNMP community string (yes/no) [n]:
  Configure read-write SNMP community string (yes/no) [n]:
 Enter the switch name : switch-name **<<<**
  Continue with Out-of-band (mgmt0) management configuration?
(yes/no) [y]:
   Mgmt0 IPv4 address : management-IP-address **<<<**
   Mgmt0 IPv4 netmask : management-IP-netmask **<<<**</pre>
 Configure the default gateway? (yes/no) [y]: y **<<<**
    IPv4 address of the default gateway : gateway-IP-address **<<<**
 Configure advanced IP options? (yes/no) [n]:
 Enable the telnet service? (yes/no) [n]:
 Enable the ssh service? (yes/no) [y]: y **<<<**
    Type of ssh key you would like to generate (dsa/rsa) [rsa]: rsa
**<<<**
   Number of rsa key bits <1024-2048> [1024]:
 Configure the ntp server? (yes/no) [n]:
 Configure default interface layer (L3/L2) [L2]:
 Configure default switchport interface state (shut/noshut)
[noshut]: shut **<<<**</pre>
  Configure CoPP system profile (strict/moderate/lenient/dense)
[strict]:
```

The final set of prompt completes the configuration:

```
The following configuration will be applied:
 password strength-check
 switchname IP switch A 1
vrf context management
ip route 0.0.0/0 10.10.99.1
exit
 no feature telnet
 ssh key rsa 1024 force
 feature ssh
 system default switchport
 system default switchport shutdown
 copp profile strict
interface mgmt0
ip address 10.10.99.10 255.255.255.0
no shutdown
Would you like to edit the configuration? (yes/no) [n]:
Use this configuration and save it? (yes/no) [y]:
2017 Jun 13 21:24:43 A1 %$ VDC-1 %$ %COPP-2-COPP POLICY: Control-
Plane is protected with policy copp-system-p-policy-strict.
Copy complete.
User Access Verification
IP switch A 1 login: admin
Password:
Cisco Nexus Operating System (NX-OS) Software
IP switch A 1#
```

2. Save the configuration:

IP_switch_A_1# copy running-config startup-config

3. Reboot the switch and wait for the switch to reload:

```
IP_switch_A_1# reload
```

4. Repeat the previous steps on the other three switches in the MetroCluster fabric configuration.
Downloading and installing the Cisco FC switch NX-OS software

You must download the switch operating system file and RCF file to each switch in the MetroCluster fabric configuration.

Before you begin

This task requires file transfer software, such as FTP, TFTP, SFTP, or SCP, to copy the files to the switches.

About this task

These steps must be repeated on each of the FC switches in the MetroCluster fabric configuration.

You must use the supported switch software version.

NetApp Hardware Universe



The outputs shown are for Cisco IP switches; however, these steps are also applicable for Cisco FC switches.

Steps

1. Download the supported NX-OS software file.

Cisco download page

2. Copy the switch software to the switch:

```
copy sftp://root@server-ip-address/tftpboot/NX-OS-file-name bootflash: vrf
management
```

In this example, the nxos.7.0.3.14.6.bin file is copied from SFTP server 10.10.99.99 to the local bootflash:

```
IP_switch_A_1# copy sftp://root@10.10.99.99/tftpboot/nxos.7.0.3.I4.6.bin
bootflash: vrf management
root@10.10.99.99's password: password
sftp> progress
Progress meter enabled
sftp> get /tftpboot/nxos.7.0.3.I4.6.bin
/bootflash/nxos.7.0.3.I4.6.bin
Fetching /tftpboot/nxos.7.0.3.I4.6.bin to /bootflash/nxos.7.0.3.I4.6.bin
/tftpboot/nxos.7.0.3.I4.6.bin 100% 666MB 7.2MB/s
01:32
sftp> exit
Copy complete, now saving to disk (please wait)...
```

3. Verify on each switch that the switch NX-OS files are present in each switch's bootflash directory:

dir bootflash

The following example shows that the files are present on IP_switch_A_1:

4. Install the switch software:

install all system bootflash:nxos.version-number.bin kickstart bootflash:nxos.version-kickstart-number.bin

```
IP switch A 1# install all system bootflash:nxos.7.0.3.I4.6.bin
kickstart bootflash:nxos.7.0.3.I4.6.bin
Installer will perform compatibility check first. Please wait.
Verifying image bootflash:/nxos.7.0.3.14.6.bin for boot variable
"kickstart".
Verifying image bootflash:/nxos.7.0.3.14.6.bin for boot variable
"system".
[##################### 100% -- SUCCESS
Performing module support checks.
Verifying image type.
Extracting "system" version from image bootflash:/nxos.7.0.3.14.6.bin.
[##################### 100% -- SUCCESS
Extracting "kickstart" version from image
bootflash:/nxos.7.0.3.I4.6.bin.
[#################### 100% -- SUCCESS
. . .
```

The switch reboot automatically after the switch software has installed.

5. Wait for the switch to reload and then log in to the switch.

After the switch has rebooted the login prompt is displayed:

```
User Access Verification
IP_switch_A_1 login: admin
Password:
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (C) 2002-2017, Cisco and/or its affiliates.
All rights reserved.
.
.
MDP database restore in progress.
IP_switch_A_1#
The switch software is now installed.
```

6. Verify that the switch software has been installed:

```
show version
```

The following example shows the output:

```
IP switch A 1# show version
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (C) 2002-2017, Cisco and/or its affiliates.
All rights reserved.
.
Software
  BIOS: version 04.24
 NXOS: version 7.0(3)14(6) **<<< switch software version**
 BIOS compile time: 04/21/2016
 NXOS image file is: bootflash:///nxos.7.0.3.I4.6.bin
 NXOS compile time: 3/9/2017 22:00:00 [03/10/2017 07:05:18]
Hardware
  cisco Nexus 3132QV Chassis
  Intel(R) Core(TM) i3- CPU @ 2.50GHz with 16401416 kB of memory.
  Processor Board ID FOC20123GPS
  Device name: A1
 bootflash: 14900224 kB
  usb1:
                      0 kB (expansion flash)
Kernel uptime is 0 day(s), 0 hour(s), 1 minute(s), 49 second(s)
Last reset at 403451 usecs after Mon Jun 10 21:43:52 2017
  Reason: Reset due to upgrade
  System version: 7.0(3) I4(1)
  Service:
plugin
  Core Plugin, Ethernet Plugin
IP switch A 1#
```

7. Repeat these steps on the remaining three FC switches in the MetroCluster fabric configuration.

Downloading and installing the Cisco FC RCF files

You must download the RCF file to each switch in the MetroCluster fabric configuration.

Before you begin

This task requires file transfer software, such as FTP, Trivial File Transfer Protocol (TFTP), SFTP, or Secure

Copy Protocol (SCP), to copy the files to the switches.

About this task

These steps must be repeated on each of the Cisco FC switches in the MetroCluster fabric configuration.

You must use the supported switch software version.

NetApp Hardware Universe

There are four RCF files, one for each of the four switches in the MetroCluster fabric configuration. You must use the correct RCF files for the switch model you are using.

Switch	RCF file
FC_switch_A_1	NX3232_v1.80_Switch-A1.txt
FC_switch_A_2	NX3232_v1.80_Switch-A2.txt
FC_switch_B_1	NX3232_v1.80_Switch-B1.txt
FC_switch_B_2	NX3232_v1.80_Switch-B2.txt



The outputs shown are for Cisco IP switches; however, these steps are also applicable for Cisco FC switches.

Steps

- 1. Download the Cisco FC RCF files from the MetroCluster RCF download page.
- 2. Copy the RCF files to the switches.
 - a. Copy the RCF files to the first switch:

```
copy sftp://root@FTP-server-IP-address/tftpboot/switch-specific-RCF
bootflash: vrf management
```

In this example, the NX3232_v1.80_Switch-A1.txt RCF file is copied from the SFTP server at 10.10.99.99 to the local bootflash. You must use the IP address of your TFTP/SFTP server and the file name of the RCF file that you need to install.

```
IP_switch_A_1# copy sftp://root@10.10.99.99/tftpboot/NX3232_v1.8T-
X1_Switch-A1.txt bootflash: vrf management
root@10.10.99.99's password: password
sftp> progress
Progress meter enabled
sftp> get /tftpboot/NX3232_v1.80_Switch-A1.txt
/bootflash/NX3232_v1.80_Switch-A1.txt
Fetching /tftpboot/NX3232_v1.80_Switch-A1.txt to
/bootflash/NX3232_v1.80_Switch-A1.txt
/tftpboot/NX3232_v1.80_Switch-A1.txt 100% 5141 5.0KB/s
00:00
sftp> exit
Copy complete, now saving to disk (please wait)...
IP_switch_A_1#
```

- b. Repeat the previous substep for each of the other three switches, being sure to copy the matching RCF file to the corresponding switch.
- 3. Verify on each switch that the RCF file is present in each switch's bootflash directory:

dir bootflash:

The following example shows that the files are present on IP_switch_A_1:

4. Copy the matching RCF file from the local bootflash to the running configuration on each switch:

copy bootflash:switch-specific-RCF.txt running-config

5. Copy the RCF files from the running configuration to the startup configuration on each switch:

copy running-config startup-config

You should see output similar to the following:

IP_switch_A_1# copy bootflash:NX3232_v1.80_Switch-A1.txt running-config IP_switch_A_1# copy running-config startup-config

6. Reload the switch:

reload

IP_switch_A_1# reload

7. Repeat the previous steps on the other three switches in the MetroCluster IP configuration.

Configure the Brocade FC switches manually

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 Secure Shell (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

In the IMT, you can use the Storage Solution field to select your MetroCluster solution. You use the **Component Explorer** to select the components and ONTAP version to refine your search. You can click **Show Results** to display the list of supported configurations that match the criteria.

- The four supported Brocade switches must be connected to two fabrics of two switches each, with each fabric spanning both sites.
- Each storage controller must have four initiator ports available to connect to the switch fabrics. Two initiator ports must be connected from each storage controller to each fabric.



You can configure FAS8020, AFF8020, FAS8200, and AFF A300 systems with two initiators ports per controller (a single initiator port to each fabric) if all the following criteria are met:

- There are fewer than four FC initiator ports available to connect the disk storage and no additional ports can be configured as FC initiators.
- All slots are in use and no FC initiator card can be added.

About this task

• You should enable Inter-Switch Link (ISL) trunking when it is supported by the links.

Considerations for using TDM/WDM equipment with fabric-attached MetroCluster configurations

- If you use an xWDM configuration, you might require additional settings on the ISLs. See the xWDM vendor documentation for more information.
- All ISLs must have the same length and same speed in one fabric.

Different lengths can be used in the different fabrics. The same speed must be used in all fabrics.

 Metro-E and TDM (SONET/SDH) are not supported, and any non-FC native framing or signaling is not supported.

Metro-E means Ethernet framing or signaling occurs either natively over a Metro distance or through some time-division multiplexing (TDM), multiprotocol label switching (MPLS), or wavelength-division multiplexing (WDM).

- TDMs, FCR (native FC Routing), or FCIP extensions are not supported for the MetroCluster FC switch fabric.
- The following Brocade FC switches support WAN ISL encryption and compression for MetroCluster FC back-end fabrics:
 - Brocade G720
 - Brocade G630
 - Brocade G620
 - Brocade 6520
- The Brocade Virtual Fabric (VF) feature is not supported.
- FC zoning based on domain port is supported, but zoning based on worldwide name (WWN) is not supported.

Review Brocade license requirements

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

About this task

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.

Step

1. Verify that the licenses are installed:

For Fabric OS 8.2.x and earlier

Run the command licenseshow.

For Fabric OS 9.0 and later

Run the command license --show.

If you do not have these licenses, you should contact your sales representative before proceeding.

Set 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.

The following example shows the command on BrocadeSwitchA:

BrocadeSwitchA:admin> switchcfgpersistentdisable

The following example shows the command on BrocadeSwitchB:

BrocadeSwitchB:admin> switchcfgpersistentdisable

3. Set the switch name:

switchname switch_name

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

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:

portcfgdefault

This must be done for all ports on the switch.

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:

cfgdisable

cfgclear

cfgsave

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
```

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
```

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
```

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:

The following example shows the commands on FC_switch_A_1:

```
FC_switch_A_1:> defzone --noaccess
FC_switch_A_1:> cfgsave
FC_switch_A_1:> exit
```

The following example shows the commands on FC_switch_B_1:

```
FC_switch_A_1:> defzone --noaccess
FC_switch_A_1:> cfgsave
FC_switch_A_1:> exit
```

10. Reboot the switch:

reboot

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

Configure basic switch settings

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

About this task

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

In this procedure, you set the unique domain ID for each switch as shown in the following example. In the example, domain IDs 5 and 7 form fabric_1, and domain IDs 6 and 8 form fabric_2.

- 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

Steps

1. Enter configuration mode:

configure

- 2. Proceed through the prompts:
 - a. Set the domain ID for the switch.
 - b. Press **Enter** in response to the prompts until you get to "RDP Polling Cycle", and then set that value to 0 to disable the polling.
 - c. Press Enter until you return to the switch prompt.

```
FC_switch_A_1:admin> configure
Fabric parameters = y
Domain_id = 5
.
.
.
.
RSCN Transmission Mode [yes, y, no, no: [no] y
End-device RSCN Transmission Mode
 (0 = RSCN with single PID, 1 = RSCN with multiple PIDs, 2 = Fabric
RSCN): (0..2) [1]
Domain RSCN To End-device for switch IP address or name change
 (0 = disabled, 1 = enabled): (0..1) [0] 1
.
.
RDP Polling Cycle(hours)[0 = Disable Polling]: (0..24) [1] 0
```

3. If you are using two or more ISLs per fabric, then you can configure either in-order delivery (IOD) of frames or out-of-order (OOD) delivery of frames.

The standard IOD settings are recommended. You should configure OOD only if necessary.

Considerations for using TDM/WDM equipment with fabric-attached MetroCluster configurations

- a. The following steps must be performed on each switch fabric to configure IOD of frames:
 - i. Enable IOD:

iodset

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

aptpolicy 1

iii. Disable Dynamic Load Sharing (DLS):

dlsreset

iv. Verify the IOD settings by using the iodshow, aptpolicy, and dlsshow commands.

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
```

- v. Repeat these steps on the second switch fabric.
- b. The following steps must be performed on each switch fabric to configure OOD of frames:
 - i. Enable OOD:

iodreset

ii. Set the Advanced Performance Tuning (APT) policy to 3:

aptpolicy 3

iii. Disable Dynamic Load Sharing (DLS):

dlsreset

iv. Verify the OOD settings:

iodshow

aptpolicy

dlsshow

For example, issue the following commands on FC_switch_A_1:

```
FC_switch_A_1:admin> iodshow
IOD is not set
FC_switch_A_1:admin> aptpolicy
Current Policy: 3 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 set by default with current routing policy
```

v. Repeat these steps on the second switch fabric.



When configuring ONTAP on the controller modules, OOD must be explicitly configured on each controller module in the MetroCluster configuration.

Configure in-order delivery or out-of-order delivery of frames on ONTAP software

4. If you are running a version earlier than FOS 9.0, verify that the switch is using the dynamic Port on Demand (POD) licensing method.



In Fabric OS 9.0 and later, the license method is dynamic by default. The static license method is not supported.

a. Run the license command:

licenseport --show

```
FC_switch_A_1:admin> license --show -port
24 ports are available in this switch
Full POD license is installed
Dynamic POD method is in use
```



Brocade FabricOS versions before 8.0 run the following commands as admin and versions 8.0 and later run them as root.

b. Enable the root user.

If the root user is already disabled by Brocade, enable the root user as shown in the following example:

```
FC_switch_A_1:admin> userconfig --change root -e yes
FC_switch_A_1:admin> rootaccess --set consoleonly
```

c. Run the license command:

```
license --show -port
```

FC_switch_A_1:root> license --show -port
24 ports are available in this switch
Full POD license is installed
Dynamic POD method is in use

d. If you are running Fabric OS 8.2.x and earlier, you must change the license method to dynamic:

licenseport --method dynamic

FC_switch_A_1:admin> licenseport --method dynamic The POD method has been changed to dynamic. Please reboot the switch now for this change to take effect

- 5. Enable the trap for T11-FC-ZONE-SERVER-MIB to provide successful health monitoring of the switches in ONTAP:
 - a. Enable the T11-FC-ZONE-SERVER-MIB:

```
snmpconfig --set mibCapability -mib_name T11-FC-ZONE-SERVER-MIB -bitmask
0x3f
```

b. Enable the T11-FC-ZONE-SERVER-MIB trap:

```
snmpconfig --enable mibcapability -mib_name SW-MIB -trap_name
swZoneConfigChangeTrap
```

- c. Repeat the previous steps on the second switch fabric.
- 6. **Optional**: If you set the community string to a value other than "public", you must configure the ONTAP Health Monitors using the community string you specify:
 - a. Change the existing community string:

snmpconfig --set snmpv1

- b. Press Enter until you see "Community (ro): [public]" text.
- c. Enter the desired community string.

On FC_switch_A_1:

```
FC switch A 1:admin> snmpconfig --set snmpv1
SNMP community and trap recipient configuration:
Community (rw): [Secret COde]
Trap Recipient's IP address : [0.0.0.0]
Community (rw): [OrigEquipMfr]
Trap Recipient's IP address : [0.0.0.0]
Community (rw): [private]
Trap Recipient's IP address : [0.0.0.0]
Community (ro): [public] mcchm <<<<< change the community string
to the desired value,
Trap Recipient's IP address : [0.0.0.0] in this example it is set
to "mcchm"
Community (ro): [common]
Trap Recipient's IP address : [0.0.0.0]
Community (ro): [FibreChannel]
Trap Recipient's IP address : [0.0.0.0]
Committing configuration....done.
FC switch A 1:admin>
```

On FC_switch_B_1:

```
FC switch B 1:admin> snmpconfig --set snmpv1
SNMP community and trap recipient configuration:
Community (rw): [Secret COde]
Trap Recipient's IP address : [0.0.0.0]
Community (rw): [OrigEquipMfr]
Trap Recipient's IP address : [0.0.0.0]
Community (rw): [private]
Trap Recipient's IP address : [0.0.0.0]
Community (ro): [public] mcchm <<<<< change the community
string to the desired value,
Trap Recipient's IP address : [0.0.0.0] in this example it is set
to "mcchm"
Community (ro): [common]
Trap Recipient's IP address : [0.0.0.0]
Community (ro): [FibreChannel]
Trap Recipient's IP address : [0.0.0.0]
Committing configuration....done.
FC switch B 1:admin>
```

```
7. Reboot the switch:
```

reboot

On FC_switch_A_1:

FC_switch_A_1:admin> reboot

On FC_switch_B_1:

FC switch B 1:admin> reboot

8. Persistently enable the switch:

switchcfgpersistentenable

On FC_switch_A_1:

FC switch A 1:admin> switchcfgpersistentenable

On FC_switch_B_1:

FC_switch_B_1:admin> switchcfgpersistentenable

Configure basic switch settings on a Brocade DCX 8510-8 switch

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

About this task

You must perform the steps on each switch at both MetroCluster sites. In this procedure, you set the 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

In the previous example, domain IDs 5 and 7 form fabric_1, and domain IDs 6 and 8 form fabric_2.



You can also use this procedure to configure the switches when you are only using one DCX 8510-8 switch per site.

Using this procedure, you should create two logical switches on each Brocade DCX 8510-8 switch. The two logical switches created on both Brocade DCX8510-8 switches will form two logical fabrics as shown in the following examples:

- LOGICAL FABRIC 1: Switch1/Blade1 and Switch 2 Blade 1
- LOGICAL FABRIC 2: Switch1/Blade2 and Switch 2 Blade 2

Steps

1. Enter the command mode:

configure

- 2. Proceed through the prompts:
 - a. Set the domain ID for the switch.
 - b. Keep selecting **Enter** until you get to "RDP Polling Cycle", and then set the value to 0 to disable the polling.
 - c. Select Enter until you return to the switch prompt.

```
FC_switch_A_1:admin> configure
Fabric parameters = y
Domain_id = `5
RDP Polling Cycle(hours)[0 = Disable Polling]: (0..24) [1] 0
`
```

- 3. Repeat these steps on all switches in fabric_1 and fabric_2.
- 4. Configure the virtual fabrics.
 - a. Enable virtual fabrics on the switch:

fosconfig --enablevf

b. Configure the system to use the same base configuration on all logical switches:

configurechassis

The following example shows the output for the configure chassis command:

```
System (yes, y, no, n): [no] n
cfgload attributes (yes, y, no, n): [no] n
Custom attributes (yes, y, no, n): [no] y
Config Index (0 to ignore): (0..1000) [3]:
```

5. Create and configure the logical switch:

scfg --create fabricID

6. Add all ports from a blade to the virtual fabric:

```
lscfg --config fabricID -slot slot -port lowest-port - highest-port
```



The blades forming a logical fabric (e.g. Switch 1 Blade 1 and Switch 3 Blade 1) need to have the same fabric ID.

```
setcontext fabricid
switchdisable
configure
<configure the switch per the above settings>
switchname unique switch name
switchenable
```

Related information

Requirements for using a Brocade DCX 8510-8 switch

Configure E-ports on Brocade FC switches using FC ports

For Brocade switches on which the Inter-Switch Links (ISL) are configured using FC ports, you must configure the switch ports on each switch fabric that connect the ISL. These ISL ports are also known as E-ports.

Before you begin

- All of the ISLs in an FC switch fabric must be configured with the same speed and distance.
- The combination of the switch port and small form-factor pluggable (SFP) must support the speed.
- The supported ISL distance depends on the FC switch model.

NetApp Interoperability Matrix Tool

In the IMT, you can use the Storage Solution field to select your MetroCluster solution. You use the **Component Explorer** to select the components and ONTAP version to refine your search. You can click **Show Results** to display the list of supported configurations that match the criteria.

• The ISL link must have a dedicated lambda, and the link must be supported by Brocade for the distance, switch type, and Fabric Operating System (FOS).

About this task

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 distance level.

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.

You must perform this task for each FC switch fabric.

The following tables show the ISL ports for different switches and different number of ISLs in a configuration running ONTAP 9.1 or 9.2. The examples shown in this section are for a Brocade 6505 switch. You should modify the examples to use ports that apply to your switch type.

You must use the required number of ISLs for your configuration.

Switch model	ISL port	Switch port
--------------	----------	-------------

Brocade 6520	ISL port 1	23
	ISL port 2	47
	ISL port 3	71
	ISL port 4	95
Brocade 6505	ISL port 1	20
	ISL port 2	21
	ISL port 3	22
	ISL port 4	23
Brocade 6510 and Brocade DCX	ISL port 1	40
8510-8	ISL port 2	41
	ISL port 3	42
	ISL port 4	43
	ISL port 5	44
	ISL port 6	45
	ISL port 7	46
	ISL port 8	47
Brocade 7810	ISL port 1	ge2 (10-Gbps)
	ISL port 2	ge3(10-Gbps)
	ISL port 3	ge4 (10-Gbps)
	ISL port 4	ge5 (10-Gbps)
	ISL port 5	ge6 (10-Gbps)
	ISL port 6	ge7 (10-Gbps)
Brocade 7840	ISL port 1	ge0 (40-Gbps) or ge2 (10-Gbps)
Note: The Brocade 7840 switch supports either two 40 Gbps VE- ports or up to four 10 Gbps VE- ports per switch for the creation of FCIP ISLs.	ISL port 2	ge1 (40-Gbps) or ge3 (10-Gbps)
	ISL port 3	ge10 (10-Gbps)
	ISL port 4	ge11 (10-Gbps)

Brocade G610, G710	ISL port 1	20
	ISL port 2	21
	ISL port 3	22
	ISL port 4	23
Brocade G620, G620-1, G630, G630-1, G720	ISL port 1	40
	ISL port 2	41
	ISL port 3	42
	ISL port 4	43
	ISL port 5	44
	ISL port 6	45
	ISL port 7	46

Steps

1. Configure the port speed:

portcfgspeed port-numberspeed

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

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

FC_switch_A_1:admin> portcfgspeed 20 16
FC_switch_A_1:admin> portcfgspeed 21 16
FC_switch_B_1:admin> portcfgspeed 20 16
FC_switch_B_1:admin> portcfgspeed 21 16

2. Configure the trunking mode for each ISL:

```
portcfgtrunkport port-number
```

• If you are configuring the ISLs for trunking (IOD), set the portcfgtrunk port-numberport-number to 1 as shown in the following example:

```
FC_switch_A_1:admin> portcfgtrunkport 20 1
FC_switch_A_1:admin> portcfgtrunkport 21 1
FC_switch_B_1:admin> portcfgtrunkport 20 1
FC switch B 1:admin> portcfgtrunkport 21 1
```

 If you do not want to configure the ISL for trunking (OOD), set portcfgtrunkport-number to 0 as shown in the following example:

```
FC_switch_A_1:admin> portcfgtrunkport 20 0
FC_switch_A_1:admin> portcfgtrunkport 21 0
FC_switch_B_1:admin> portcfgtrunkport 20 0
FC switch B 1:admin> portcfgtrunkport 21 0
```

3. Enable QoS traffic for each of the ISL ports:

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

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

```
FC_switch_A_1:admin> portcfgqos --enable 20
FC_switch_A_1:admin> portcfgqos --enable 21
FC_switch_B_1:admin> portcfgqos --enable 20
FC_switch_B_1:admin> portcfgqos --enable 21
```

4. Verify the settings:

portCfgShow command

The following example shows the output for a configuration that uses two ISLs cabled to port 20 and port 21. The Trunk Port setting should be ON for IOD and OFF for OOD:

```
14 15
Ports of Slot 0
               12 13
                               16 17 18
                                         19
                                              20 21 22 23
                                                             24
25 26 27
----+---+----+----
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
                . .
                   • •
                      . .
                                   . .
                                                 . .
                          . .
                               . .
                                      . .
                                              . .
.. .. .. ..
                                                ON
Trunk Port
                                              ON
                                          . .
                                                     • •
                . .
                   . .
                      . .
                          . .
                               . .
                                   . .
                                      . .
  •• ••
          . .
```

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 NPIV PP Limit 126 126 126 126 126 126 126 126 126 126 126 126 126 126 126 126 126 QOS E Port AE ΑE AE AE AE AE Mirror Port • • • • Rate Limit •• •• • • • • • • • • • • • • •• •• • • • • Credit Recovery 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.

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

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

If the distance is 3 km, then 1.5 × 3 km = 4.5 km. This is lower than 10 km, so the ISL must be set to the LE distance level.

If the distance is 20 km, then 1.5×20 km = 30 km. The ISL must be set to 30 km and must use the LS distance level.

6. Set the distance on each ISL port:

portcfglongdistance portdistance-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 might depend on the link being used. The commands must be repeated for each ISL port.

For an ISL distance of 3 km, as given in the example in the previous step, the setting is 4.5 km with the default vc_link_init value of 1. Because a setting of 4.5 km is lower than 10 km, the port needs to be set to the LE distance level:

```
FC_switch_A_1:admin> portcfglongdistance 20 LE 1
```

FC_switch_B_1:admin> portcfglongdistance 20 LE 1

For an ISL distance of 20 km, as given in the example in the previous step, the setting is 30 km with the default vc_link_init value of 1:

```
FC_switch_A_1:admin> portcfglongdistance 20 LS 1 -distance 30 \,
```

```
FC_switch_B_1:admin> portcfglongdistance 20 LS 1 -distance 30
```

7. Verify the distance setting:

portbuffershow

A distance level of LE appears as 10 km.

The following example shows the output for a configuration that uses ISLs on port 20 and port 21:

```
FC_switch_A_1:admin> portbuffershow
                   Max/Resv
                                                    Remaining
User Port
           Lx
                             Buffer Needed Link
                             Usage Buffers Distance Buffers
Port Type
                   Buffers
           Mode
____
     ____
           ____
                   _____
                             _____
                                            _____
. . .
                                             30km
20
     E
                      8
                              67
                                     67
            _
                      8
                                             30km
21
                              67
                                     67
      Ε
            _
. . .
23
            _
                      8
                               0
                                     _
                                             _
                                                    466
```

8. Verify that both switches form one fabric:

switchshow

The following example shows the output for a configuration that uses ISLs on port 20 and port 21:

```
FC switch A 1:admin> switchshow
switchName: FC switch A 1
switchType: 109.1
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
_____
. . .
20 20 010C00 id 16G Online FC LE E-Port
10:00:00:05:33:8c:2e:9a "FC switch B 1" (downstream) (trunk master)
21 21 010D00 id 16G Online FC LE E-Port (Trunk port, master
is Port 20)
. . .
FC switch B 1:admin> switchshow
switchName: FC switch B 1
switchType: 109.1
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
_____
. . .
20 20 030C00 id 16G Online FC LE E-Port
10:00:00:05:33:86:89:cb "FC switch A 1" (downstream) (Trunk master)
21 21 030D00 id 16G Online FC LE E-Port (Trunk port, master
is Port 20)
. . .
```

9. Confirm the configuration of the fabrics:

fabricshow

```
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. Confirm the trunking of the ISLs:

trunkshow

• If you are configuring the ISLs for trunking (IOD), you should see output similar to the following:

```
FC_switch_A_1:admin> trunkshow
1: 20-> 20 10:00:00:05:33:ac:2b:13 3 deskew 15 MASTER
    21-> 21 10:00:00:05:33:8c:2e:9a 3 deskew 16
FC_switch_B_1:admin> trunkshow
1: 20-> 20 10:00:00:05:33:86:89:cb 3 deskew 15 MASTER
    21-> 21 10:00:00:05:33:86:89:cb 3 deskew 16
```

• If you are not configuring the ISLs for trunking (OOD), you should see output similar to the following:

FC_switch_A_1:admin> trunkshow
1: 20-> 20 10:00:00:05:33:ac:2b:13 3 deskew 15 MASTER
2: 21-> 21 10:00:00:05:33:8c:2e:9a 3 deskew 16 MASTER
FC_switch_B_1:admin> trunkshow
1: 20-> 20 10:00:00:05:33:86:89:cb 3 deskew 15 MASTER
2: 21-> 21 10:00:00:05:33:86:89:cb 3 deskew 16 MASTER

11. Repeat Step 1 through Step 10 for the second FC switch fabric.

Related information

Port assignments for FC switches

Configuring 10 Gbps VE ports on Brocade FC 7840 switches

When using the 10 Gbps VE ports (which use FCIP) for ISLs, you must create IP interfaces on each port, and configure FCIP tunnels and circuits in each tunnel.

About this task

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

The examples in this procedure assume that the two Brocade 7840 switches have the following IP addresses:

- FC_switch_A_1 is local.
- FC_switch_B_1 is remote.

Steps

1. Create IP interface (ipif) addresses for the 10 Gbps ports on both switches in the fabric:

```
portcfg ipif FC_switch1_namefirst_port_name create FC_switch1_IP_address
netmask netmask number vlan 2 mtu auto
```

The following command creates ipif addresses on ports ge2.dp0 and ge3.dp0 of FC_switch_A_1:

```
portcfg ipif ge2.dp0 create 10.10.20.71 netmask 255.255.0.0 vlan 2 mtu
auto
portcfg ipif ge3.dp0 create 10.10.21.71 netmask 255.255.0.0 vlan 2 mtu
auto
```

The following command creates ipif addresses on ports ge2.dp0 and ge3.dp0 of FC_switch_B_1:

```
portcfg ipif ge2.dp0 create 10.10.20.72 netmask 255.255.0.0 vlan 2 mtu
auto
portcfg ipif ge3.dp0 create 10.10.21.72 netmask 255.255.0.0 vlan 2 mtu
auto
```

2. Verify that the ipif addresses were created successfully on both switches:

portshow ipif all

The following command shows the ipif addresses on switch FC_switch_A_1:

```
FC_switch_A_1:root> portshow ipif all
Port IP Address / Pfx MTU VLAN Flags
ge2.dp0 10.10.20.71 / 24 AUTO 2 U R M I
ge3.dp0 10.10.21.71 / 20 AUTO 2 U R M I
Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running
I=InUse
N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport
```

The following command shows the ipif addresses on switch FC_switch_B_1:

```
FC_switch_B_1:root> portshow ipif all
Port IP Address / Pfx MTU VLAN Flags
ge2.dp0 10.10.20.72 / 24 AUTO 2 U R M I
ge3.dp0 10.10.21.72 / 20 AUTO 2 U R M I
------
Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running
I=InUse
N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport
```

3. Create the first of the two FCIP tunnels using the ports on dp0:

portcfg fciptunnel

This command creates a tunnel with a single circuit.

The following command creates the tunnel on switch FC_switch_A_1:

```
portcfg fciptunnel 24 create -S 10.10.20.71 -D 10.10.20.72 -b 10000000
-B 10000000
```

The following command creates the tunnel on switch FC_switch_B_1:

```
portcfg fciptunnel 24 create -S 10.10.20.72 -D 10.10.20.71 -b 10000000
-B 10000000
```

4. Verify that the FCIP tunnels were successfully created:

portshow fciptunnel all

The following example shows that the tunnels were created and the circuits are up:

5. Create an additional circuit for dp0.

The following command creates a circuit on switch FC_switch_A_1 for dp0:

```
portcfg fcipcircuit 24 create 1 -S 10.10.21.71 -D 10.10.21.72 --min
-comm-rate 5000000 --max-comm-rate 5000000
```

The following command creates a circuit on switch FC switch B 1 for dp0:

```
portcfg fcipcircuit 24 create 1 -S 10.10.21.72 -D 10.10.21.71 --min
-comm-rate 5000000 --max-comm-rate 5000000
```

6. Verify that all circuits were successfully created:

```
portshow fcipcircuit all
```

The following command shows the circuits and their status:

```
FC switch A 1:root> portshow fcipcircuit all
Tunnel Circuit OpStatus Flags Uptime TxMBps RxMBps ConnCnt
CommRt Met/G
_____
_____
24
     0 qe2
            Up
                  ---va---4 2d12m 0.02
                                          0.03
                                                 3
10000/10000 0/-
24 1 ge3
                   ---va---4 2d12m 0.02
                                           0.04
                                                 3
             Up
10000/10000 0/-
_____
                                       _____
Flags (circuit): h=HA-Configured v=VLAN-Tagged p=PMTU i=IPSec 4=IPv4
6=IPv6
             ARL a=Auto r=Reset s=StepDown t=TimedStepDown S=SLA
```

Configure 40 Gbps VE-ports on Brocade 7810 and 7840 FC switches

When using the two 40 GbE VE-ports (which use FCIP) for ISLs, you must create IP interfaces on each port, and configure FCIP tunnels and circuits in each tunnel.

About this task

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

The examples in this procedure use two switches:

- FC switch A 1 is local.
- FC_switch_B_1 is remote.

Steps

1. Create IP interface (ipif) addresses for the 40 Gbps ports on both switches in the fabric:

portcfg ipif FC_switch_namefirst_port_name create FC_switch_IP_address netmask
netmask_number vlan 2 mtu auto

The following command creates ipif addresses on ports ge0.dp0 and ge1.dp0 of FC_switch_A_1:

```
portcfg ipif ge0.dp0 create 10.10.82.10 netmask 255.255.0.0 vlan 2 mtu
auto
portcfg ipif ge1.dp0 create 10.10.82.11 netmask 255.255.0.0 vlan 2 mtu
auto
```

The following command creates ipif addresses on ports ge0.dp0 and ge1.dp0 of FC_switch_B_1:

```
portcfg ipif ge0.dp0 create 10.10.83.10 netmask 255.255.0.0 vlan 2 mtu
auto
portcfg ipif ge1.dp0 create 10.10.83.11 netmask 255.255.0.0 vlan 2 mtu
auto
```

2. Verify that the ipif addresses were successfully created on both switches:

portshow ipif all

The following example shows the IP interfaces on FC_switch_A_1:

```
/ Pfx MTU VLAN Flags
      IP Address
Port
              ------
_____
___
____
                              / 16 AUTO 2 URM
ge0.dp0
       10.10.82.10
         10.10.82.11
                              / 16 AUTO 2
                                           URM
gel.dp0
_____
Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running
I=InUse
    N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport
```

The following example shows the IP interfaces on FC_switch_B_1:

```
IP Address
                                / Pfx MTU VLAN Flags
Port
                _____
    _____
_____
ge0.dp0
         10.10.83.10
                                 / 16 AUTO 2
                                              URM
gel.dp0
         10.10.83.11
                                / 16 AUTO 2
                                              URM
                  _____
Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running
I=InUse
     N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport
```

3. Create the FCIP tunnel on both switches:

portcfig fciptunnel

The following command creates the tunnel on FC_switch_A_1:

```
portcfg fciptunnel 24 create -S 10.10.82.10 -D 10.10.83.10 -b 10000000
-B 10000000
```

The following command creates the tunnel on FC_switch_B_1:

```
portcfg fciptunnel 24 create -S 10.10.83.10 -D 10.10.82.10 -b 10000000
-B 10000000
```

4. Verify that the FCIP tunnel has been successfully created:

portshow fciptunnel all

The following example shows that the tunnel was created and the circuits are up:

5. Create an additional circuit on each switch:

```
portcfg fcipcircuit 24 create 1 -S source-IP-address -D destination-IP-address
--min-comm-rate 10000000 --max-comm-rate 10000000
```

The following command creates a circuit on switch FC_switch_A_1 for dp0:

```
portcfg fcipcircuit 24 create 1 -S 10.10.82.11 -D 10.10.83.11 --min
-comm-rate 10000000 --max-comm-rate 10000000
```

The following command creates a circuit on switch FC switch B 1 for dp1:

```
portcfg fcipcircuit 24 create 1 -S 10.10.83.11 -D 10.10.82.11 --min
-comm-rate 10000000 --max-comm-rate 10000000
```

6. Verify that all circuits were successfully created:

```
portshow fcipcircuit all
```

The following example lists the circuits and shows that their OpStatus is up:

```
FC switch A 1:root> portshow fcipcircuit all
Tunnel Circuit OpStatus Flags Uptime TxMBps RxMBps ConnCnt
CommRt Met/G
_____
_____
                 ---va---4 2d12m 0.02 0.03
24 0 qe0 Up
                                               3
10000/10000 0/-
24
    1 ge1
          Up
                  ---va---4 2d12m 0.02
                                         0.04
                                               3
10000/10000 0/-
_____
_____
Flags (circuit): h=HA-Configured v=VLAN-Tagged p=PMTU i=IPSec 4=IPv4
6=IPv6
            ARL a=Auto r=Reset s=StepDown t=TimedStepDown S=SLA
```

Configure the non-E-ports on the Brocade switch

You must configure the non-E-ports on the FC switch. In a 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 portspeed

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, and so on).

For example, the components might have the following supported speeds:
- The SFP is capable of 4, 8, or 16 GB.
- The switch port is capable of 4, 8, or 16 GB.
- The connected HBA maximum speed is 16 GB. The highest common speed in this case is 16 GB, so the port should be configured for a speed of 16 GB.

```
FC_switch_A_1:admin> portcfgspeed 6 16
FC_switch_B_1:admin> portcfgspeed 6 16
```

2. Verify the settings:

portcfgshow

```
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 16G:

Ports of Slot 0	0	1	2	3	4	5	6	7	8
Speed	+- 16G	+ 16G	+ 16G	+ 16G	-+ 16G	-+· 16G	-+ 16G	-+· 16G	-+ 16G
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
NPIV PP Limit	126	126	126	126	126	126	126	126	126
QOS Port	AE	AE	AE	AE	AE	AE	AE	AE	ON
EX Port	•••	••	••	••	•••	••	••	••	••
Mirror Port	•••	••	••	••	•••	••	••	••	••
Rate Limit	•••	••	••	••	•••	••	••	••	••
Credit Recovery	ON	ON	ON	ON	ON	ON	ON	ON	ON
Fport Buffers	••	••	••	••	••	••	••	••	••
Eport Credits	••	••	••	••	••	••	••	••	••
Port Auto Disable	••	••	••	••	••	••	••	••	••
CSCTL mode	••	••	••	••	••	••	••	••	••
D-Port mode	••	••	••	••	••	••	••	••	••
D-Port over DWDM	••	••	••	••	••	••	••	••	••
FEC	ON	ON	ON	ON	ON	ON	ON	ON	ON
Fault Delay	0	0	0	0	0	0	0	0	0
Non-DFE		•••	••	•••	• •	•••	•••	• •	•••

Configure compression on ISL ports on a Brocade G620 switch

If you are using Brocade G620 switches and enabling compression on the ISLs, you must configure it on each E-port on the switches.

About this task

This task must be performed on the ISL ports on both switches using the ISL.

Steps

1. Disable the port on which you want to configure compression:

portdisable port-id

2. Enable compression on the port:

portCfgCompress --enable port-id

3. Enable the port to activate the configuration with compression:

portenable port-id

4. Confirm that the setting has been changed:

portcfgshow port-id

The following example enables compression on port 0.

```
FC_switch_A_1:admin> portdisable 0
FC_switch_A_1:admin> portcfgcompress --enable 0
FC_switch_A_1:admin> portenable 0
FC_switch_A_1:admin> portcfgshow 0
Area Number: 0
Octet Speed Combo: 3(16G,10G)
(output truncated)
D-Port mode: OFF
D-Port over DWDM ..
Compression: ON
Encryption: ON
```

You can use the islShow command to check that the E_port has come online with encryption or compression configured and active.

```
FC_switch_A_1:admin> islshow
    1: 0-> 0 10:00:c4:f5:7c:8b:29:86 5 FC_switch_B_1
sp: 16.000G bw: 16.000G TRUNK QOS CR_RECOV ENCRYPTION COMPRESSION
```

You can use the portEncCompShow command to see which ports are active. In this example you can see that encryption and compression are configured and active on port 0.

FC_swi	tch_A_1:admin>	portencc	ompshow		
User	Encryption		Com	pression	Config
Port	Configured	Active	Configured	Active Speed	
0	Yes	Yes	Yes	Yes	16G

Configure zoning on Brocade FC switches

You must assign the switch ports to separate zones to separate controller and storage traffic.

Zone the FC-VI ports

For each DR group in the MetroCluster, you must configure two zones for the FC-VI connections that allow controller-to-controller traffic. These zones contain the FC switch ports connecting to the controller module FC-VI ports. These zones are Quality of Service (QoS) zones.

A QoS zone name starts with the prefix QOSHid_, followed by a user-defined string to differentiate it from a regular zone. These QoS zones are the same regardless of the model of FibreBridge bridge that is being used.

Each zone contains all the FC-VI ports, one for each FC-VI cable from each controller. These zones are configured for high priority.

The following tables show the FC-VI zones for two DR groups.

	DR	group 1	2	QOSH1	FC-VI	zone for	FC-VI	port a	a /	С
--	----	---------	---	-------	-------	----------	--------------	--------	-----	---

FC switch	Site	Switch domain	6505 / 6510 port	6520 port	G620 port	Connects to
FC_switch_A_ 1	A	5	0	0	0	controller_A_1 port FC-VI a
FC_switch_A_ 1	A	5	1	1	1	controller_A_1 port FC-VI c
FC_switch_A_ 1	A	5	4	4	4	controller_A_2 port FC-VI a
FC_switch_A_ 1	A	5	5	5	5	controller_A_2 port FC-VI c
FC_switch_B_ 1	В	7	0	0	0	controller_B_1 port FC-VI a
FC_switch_B_ 1	В	7	1	1	1	controller_B_1 port FC-VI c
FC_switch_B_ 1	В	7	4	4	4	controller_B_2 port FC-VI a
FC_switch_B_ 1	В	7	5	5	5	controller_B_2 port FC-VI c

Zone in Fabric_1	Member ports
QOSH1_MC1_FAB_1_FCVI	5,0;5,1;5,4;5,5;7,0;7,1;7,4;7,5

DR group 1 : QOSH1 FC-VI zone for FC-VI port b / d

FC switch	Site	Switch domain	6505 / 6510 port	6520 port	G620 port	Connects to
FC_switch_A_ 2	A	6	0	0	0	controller_A_1 port FC-VI b
			1	1	1	controller_A_1 port FC-VI d

FC switch	Site	Switch domain	6505 / 6510 port	6520 port	G620 port	Connects to
			4	4	4	controller_A_2 port FC-VI b
			5	5	5	controller_A_2 port FC-VI d
FC_switch_B_ 2	В	8	0	0	0	controller_B_1 port FC-VI b
			1	1	1	controller_B_1 port FC-VI d
			4	4	4	controller_B_2 port FC-VI b
			5	5	5	controller_B_2 port FC-VI d

Zone in Fabric_1	Member ports
QOSH1_MC1_FAB_2_FCVI	6,0;6,1;6,4;6,5;8,0;8,1;8,4;8,5

DR group 2 : QOSH2 FC-VI zone for FC-VI port a / c

FC switch	Site	Switch domain	Switch port			Connects to
			6510	6520	G620	
FC_switch_A_ 1	A	5	24	48	18	controller_A_3 port FC-VI a
			25	49	19	controller_A_3 port FC-VI c
			28	52	22	controller_A_4 port FC-VI a
			29	53	23	controller_A_4 port FC-VI c
FC_switch_B_ 1	В	7	24	48	18	controller_B_3 port FC-VI a
			25	49	19	controller_B_3 port FC-VI c
			28	52	22	controller_B_4 port FC-VI a
			29	53	23	controller_B_4 port FC-VI c

Zone in Fabric_1	Member ports
QOSH2_MC2_FAB_1_FCVI (6510)	5,24;5,25;5,28;5,29;7,24;7,25;7,28;7,29

DR group 2 : QOSH2 FC-VI zone for FC-VI port b / d

FC switch	Site	Switch domain	6510 port	6520 port	G620 port	Connects to
FC_switch_A_ 2	A	6	24	48	18	controller_A_3 port FC-VI b
FC_switch_A_ 2	A	6	25	49	19	controller_A_3 port FC-VI d
FC_switch_A_ 2	A	6	28	52	22	controller_A_4 port FC-VI b
FC_switch_A_ 2	A	6	29	53	23	controller_A_4 port FC-VI d
FC_switch_B_ 2	В	8	24	48	18	controller_B_3 port FC-VI b
FC_switch_B_ 2	В	8	25	49	19	controller_B_3 port FC-VI d
FC_switch_B_ 2	В	8	28	52	22	controller_B_4 port FC-VI b
FC_switch_B_ 2	В	8	29	53	23	controller_B_4 port FC-VI d

Zone in Fabric_2	Member ports
QOSH2_MC2_FAB_2_FCVI (6510)	6,24;6,25;6,28;6,29;8,24;8,25;8,28;8,29
QOSH2_MC2_FAB_2_FCVI (6520)	6,48;6,49;6,52;6,53;8,48;8,49;8,52;8,53

The following table provides a summary of the FC-VI zones:

Fabric	Zone name	Member ports
FC_switch_A_1 and FC_switch_B_1	QOSH1_MC1_FAB_1_FCVI	5,0;5,1;5,4;5,5;7,0;7,1;7,4;7,5
	QOSH2_MC1_FAB_1_FCVI (6510)	5,24;5,25;5,28;5,29;7,24;7,25;7,28; 7,29
	QOSH2_MC1_FAB_1_FCVI (6520)	5,48;5,49;5,52;5,53;7,48;7,49;7,52; 7,53

FC_switch_A_2 and FC_switch_B_2	QOSH1_MC1_FAB_2_FCVI	6,0;6,1;6,4;6,5;8,0;8,1;8,4;8,5
	QOSH2_MC1_FAB_2_FCVI (6510)	6,24;6,25;6,28;6,29;8,24;8,25;8,28; 8,29
	QOSH2_MC1_FAB_2_FCVI (6520)	6,48;6,49;6,52;6,53;8,48;8,49;8,52; 8,53

Zone FibreBridge 7500N or 7600N bridges using one FC port

If you are using FibreBridge 7500N or 7600N bridges using only one of the two FC ports, you need to create storage zones for the bridge ports. You should understand the zones and associated ports before you configure the zones.

The examples show zoning for DR group 1 only. If your configuration includes a second DR group, configure the zoning for the second DR group in the same manner, using the corresponding ports of the controllers and bridges.

Required zones

You must configure one zone for each of the FC-to-SAS bridge FC ports that allows traffic between initiators on each controller module and that FC-to-SAS bridge.

Each storage zone contains nine ports:

- Eight HBA initiator ports (two connections for each controller)
- One port connecting to an FC-to-SAS bridge FC port

The storage zones use standard zoning.

The examples show two pairs of bridges connecting two stack groups at each site. Because each bridge uses one FC port, there are a total of four storage zones per fabric (eight in total).

Bridge naming

The bridges use the following example naming: bridge_site_stack grouplocation in pair

This portion of the name	Identifies the	Possible values
site	Site on which the bridge pair physically resides.	A or B
stack group	Number of the stack group to which the bridge pair connects. FibreBridge 7600N or 7500N bridges support up to four stacks in the stack group. The stack group can contain no more than 10 storage shelves.	1, 2, etc.

location in pair	Bridge within the bridge pair.A pair of bridges connect to a specific stack group.	a or b

Example bridge names for one stack group on each site:

- bridge_A_1a
- bridge_A_1b
- bridge_B_1a
- bridge_B_1b

DR Group 1 - Stack 1 at Site_A

DrGroup 1 : MC1_INIT_GRP_1_SITE_A_STK_GRP_1_TOP_FC1:

FC switch	Site	Switch domain	Brocade 6505, 6510, 6520, G620, G610, or G710 switch port	Connects to
FC_switch_A_1	A	5	2	controller_A_1 port 0a
FC_switch_A_1	A	5	3	controller_A_1 port 0c
FC_switch_A_1	A	5	6	controller_A_2 port 0a
FC_switch_A_1	А	5	7	controller_A_2 port 0c
FC_switch_A_1	А	5	8	bridge_A_1a FC1
FC_switch_B_1	В	7	2	controller_B_1 port 0a
FC_switch_B_1	В	7	3	controller_B_1 port 0c
FC_switch_B_1	В	7	6	controller_B_2 port 0a
FC_switch_B_1	В	7	7	controller_B_2 port 0c

Zone in Fabric_1	Member ports
MC1_INIT_GRP_1_SITE_A_STK_GRP_1_TOP_FC1	5,2;5,3;5,6;5,7;7,2;7,3;7,6;7,7;5,8

DrGroup 1 : MC1_INIT_GRP_1_SITE_A_STK_GRP_1_BOT_FC1:

FC switch	Site	Switch domain	Brocade 6505, 6510, 6520, G620, G610, or G710 switch port	Connects to
FC_switch_A_1	A	6	2	controller_A_1 port 0b
FC_switch_A_1	A	6	3	controller_A_1 port 0d
FC_switch_A_1	A	6	6	controller_A_2 port 0b
FC_switch_A_1	A	6	7	controller_A_2 port 0d
FC_switch_A_1	А	6	8	bridge_A_1b FC1
FC_switch_B_1	В	8	2	controller_B_1 port 0b
FC_switch_B_1	В	8	3	controller_B_1 port 0d
FC_switch_B_1	В	8	6	controller_B_2 port 0b
FC_switch_B_1	В	8	7	controller_B_2 port 0d

Zone in Fabric_2	Member ports
MC1_INIT_GRP_1_SITE_A_STK_GRP_1_BOT_FC1	6,2;6,3;6,6;6,7;8,2;8,3;8,6;8,7;6,8

DR Group 1 - Stack 2 at Site_A

DrGroup 1 : MC1_INIT_GRP_1_SITE_A_STK_GRP_2_TOP_FC1:

FC switch	Site	Switch domain	Brocade 6505, 6510, 6520, G620, G610, or G710 switch port	Connects to
FC_switch_A_1	A	5	2	controller_A_1 port 0a
FC_switch_A_1	A	5	3	controller_A_1 port 0c
FC_switch_A_1	A	5	6	controller_A_2 port 0a
FC_switch_A_1	A	5	7	controller_A_2 port 0c
FC_switch_A_1	A	5	9	bridge_A_2a FC1

FC switch	Site	Switch domain	Brocade 6505, 6510, 6520, G620, G610, or G710 switch port	Connects to
FC_switch_B_1	В	7	2	controller_B_1 port 0a
FC_switch_B_1	В	7	3	controller_B_1 port 0c
FC_switch_B_1	В	7	6	controller_B_2 port 0a
FC_switch_B_1	В	7	7	controller_B_2 port 0c

Zone in Fabric_1	Member ports
MC1_INIT_GRP_1_SITE_A_STK_GRP_2_TOP_FC1	5,2;5,3;5,6;5,7;7,2;7,3;7,6;7,7;5,9

DrGroup 1 : MC1_INIT_GRP_1_SITE_A_STK_GRP_2_BOT_FC1:

FC switch	Site	Switch domain	Brocade 6505, 6510, 6520, G620, G610, or G710 switch port	Connects to
FC_switch_A_1	A	6	2	controller_A_1 port 0b
FC_switch_A_1	A	6	3	controller_A_1 port 0d
FC_switch_A_1	A	6	6	controller_A_2 port 0b
FC_switch_A_1	A	6	7	controller_A_2 port 0d
FC_switch_A_1	А	6	9	bridge_A_2b FC1
FC_switch_B_1	В	8	2	controller_B_1 port 0b
FC_switch_B_1	В	8	3	controller_B_1 port 0d
FC_switch_B_1	В	8	6	controller_B_2 port 0b
FC_switch_B_1	В	8	7	controller_B_2 port 0d

Zone in Fabric_2	Member ports
MC1_INIT_GRP_1_SITE_A_STK_GRP_2_BOT_FC1	6,2;6,3;6,6;6,7;8,2;8,3;8,6;8,7;6,9

DR Group 1 - Stack 1 at Site_B

MC1_INIT_GRP_1_SITE_B_STK_GRP_1_TOP_FC1:

FC switch	Site	Switch domain	Brocade 6505, 6510, 6520, G620, G610, or G710 switch	Connects to
FC_switch_A_1	A	5	2	controller_A_1 port 0a
FC_switch_A_1	A	5	3	controller_A_1 port 0c
FC_switch_A_1	A	5	6	controller_A_2 port 0a
FC_switch_A_1	A	5	7	controller_A_2 port 0c
FC_switch_B_1	В	7	2	controller_B_1 port 0a
FC_switch_B_1	В	7	3	controller_B_1 port 0c
FC_switch_B_1	В	7	6	controller_B_2 port 0a
FC_switch_B_1	В	7	7	controller_B_2 port 0c
FC_switch_B_1	В	7	8	bridge_B_1a FC1

Zone in Fabric_1	Member ports
MC1_INIT_GRP_1_SITE_B_STK_GRP_1_TOP_FC1	5,2;5,3;5,6;5,7;7,2;7,3;7,6;7,7;7,8

DrGroup 1 : MC1_INIT_GRP_1_SITE_B_STK_GRP_1_BOT_FC1:

FC switch	Site	Switch domain	Brocade 6505, 6510, 6520, G620, G610, or G710 switch	Connects to
FC_switch_A_1	A	6	2	controller_A_1 port 0b
FC_switch_A_1	A	6	3	controller_A_1 port 0d
FC_switch_A_1	A	6	6	controller_A_2 port 0b
FC_switch_A_1	A	6	7	controller_A_2 port 0d

FC switch	Site	Switch domain	Brocade 6505, 6510, 6520, G620, G610, or G710 switch	Connects to
FC_switch_B_1	В	8	2	controller_B_1 port 0b
FC_switch_B_1	В	8	3	controller_B_1 port 0d
FC_switch_B_1	В	8	6	controller_B_2 port 0b
FC_switch_B_1	В	8	7	controller_B_2 port 0d
FC_switch_B_1	В	8	8	bridge_B_1b FC1

Zone in Fabric_2	Member ports
MC1_INIT_GRP_1_SITE_B_STK_GRP_1_BOT_FC1	5,2;5,3;5,6;5,7;7,2;7,3;7,6;7,7;8,8

DR Group 1 - Stack 2 at Site_B

DrGroup 1 : MC1_INIT_GRP_1_SITE_B_STK_GRP_2_TOP_FC1:

FC switch	Site	Switch domain	Brocade 6505, 6510, 6520, G620, G610, or G710 switch port	Connects to
FC_switch_A_1	A	5	2	controller_A_1 port 0a
FC_switch_A_1	A	5	3	controller_A_1 port 0c
FC_switch_A_1	A	5	6	controller_A_2 port 0a
FC_switch_A_1	A	5	7	controller_A_2 port 0c
FC_switch_B_1	В	7	2	controller_B_1 port 0a
FC_switch_B_1	В	7	3	controller_B_1 port 0c
FC_switch_B_1	В	7	6	controller_B_2 port 0a
FC_switch_B_1	В	7	7	controller_B_2 port 0c
FC_switch_B_1	В	7	9	bridge_b_2a FC1

Zone in Fabric_1

Member ports

MC1_INIT_GRP_1_SITE_b_STK_GRP_2_TOP_FC1 5,2;5,3;5,6;5,7;7,2;7,3;7,6;7,7;7,9

DrGroup 1 : MC1_INIT_GRP_1_SITE_B_STK_GRP_2_BOT_FC1:

FC switch	Site	Switch domain	Brocade 6505, 6510, 6520, G620, G610, or G710 switch port	Connects to
FC_switch_A_1	A	6	2	controller_A_1 port 0b
FC_switch_A_1	A	6	3	controller_A_1 port 0d
FC_switch_A_1	A	6	6	controller_A_2 port 0b
FC_switch_A_1	A	6	7	controller_A_2 port 0d
FC_switch_B_1	В	8	2	controller_B_1 port 0b
FC_switch_B_1	В	8	3	controller_B_1 port 0d
FC_switch_B_1	В	8	6	controller_B_2 port 0b
FC_switch_B_1	В	8	7	controller_B_2 port 0d
FC_switch_B_1	В	8	9	bridge_B_1b FC1

Zone in Fabric_2	Member ports
MC1_INIT_GRP_1_SITE_B_STK_GRP_2_BOT_FC1	6,2;6,3;6,6;6,7;8,2;8,3;8,6;8,7;8,9

Summary of storage zones

Fabric	Zone name	Member ports
FC_switch_A_1 and FC_switch_B_1	MC1_INIT_GRP_1_SITE_A_STK_ GRP_1_TOP_FC1	5,2;5,3;5,6;5,7;7,2;7,3;7,6;7,7;5,8
	MC1_INIT_GRP_1_SITE_A_STK_ GRP_2_TOP_FC1	5,2;5,3;5,6;5,7;7,2;7,3;7,6;7,7;5,9
	MC1_INIT_GRP_1_SITE_B_STK_ GRP_1_TOP_FC1	5,2;5,3;5,6;5,7;7,2;7,3;7,6;7,7;7,8
	MC1_INIT_GRP_1_SITE_B_STK_ GRP_2_TOP_FC1	5,2;5,3;5,6;5,7;7,2;7,3;7,6;7,7;7,9

FC_switch_A_2 and FC_switch_B_2	MC1_INIT_GRP_1_SITE_A_STK_ GRP_1_BOT_FC1	6,2;6,3;6,6;6,7;8,2;8,3;8,6;8,7;6,8
	MC1_INIT_GRP_1_SITE_A_STK_ GRP_2_BOT_FC1	6,2;6,3;6,6;6,7;8,2;8,3;8,6;8,7;6,9
	MC1_INIT_GRP_1_SITE_B_STK_ GRP_1_BOT_FC1	6,2;6,3;6,6;6,7;8,2;8,3;8,6;8,7;8,8
	MC1_INIT_GRP_1_SITE_B_STK_ GRP_2_BOT_FC1	6,2;6,3;6,6;6,7;8,2;8,3;8,6;8,7;8,9

Zone FibreBridge 7500N bridges using both FC ports

If you are using FibreBridge 7500N bridges with both FC ports, you need to create storage zones for the bridge ports. You should understand the zones and associated ports before you configure the zones.

Required zones

You must configure one zone for each of the FC-to-SAS bridge FC ports that allows traffic between initiators on each controller module and that FC-to-SAS bridge.

Each storage zone contains five ports:

- Four HBA initiator ports (one connection for each controller)
- One port connecting to an FC-to-SAS bridge FC port

The storage zones use standard zoning.

The examples show two pairs of bridges connecting two stack groups at each site. Because each bridge uses one FC port, there are a total of eight storage zones per fabric (sixteen in total).

Bridge naming

The bridges use the following example naming: bridge_site_stack grouplocation in pair

This portion of the name	Identifies the	Possible values
site	Site on which the bridge pair physically resides.	A or B
stack group	Number of the stack group to which the bridge pair connects. FibreBridge 7600N or 7500N bridges support up to four stacks in the stack group. The stack group can contain no more than 10 storage shelves.	1, 2, etc.

location in pair	Bridge within the bridge pair. A pair of bridges connect to a specific stack group.	a or b

Example bridge names for one stack group on each site:

- bridge_A_1a
- bridge_A_1b
- bridge_B_1a
- bridge_B_1b

DR Group 1 - Stack 1 at Site_A

DrGroup 1 : MC1_INIT_GRP_1_SITE_A_STK_GRP_1_TOP_FC1:

FC switch	Site	Switch domain	6505 / 6510 / G610 /G710 / G620 port	6520 port	Connects to
FC_switch_A_1	A	5	2	2	controller_A_1 port 0a
FC_switch_A_1	A	5	6	6	controller_A_2 port 0a
FC_switch_A_1	A	5	8	8	bridge_A_1a FC1
FC_switch_B_1	В	7	2	2	controller_B_1 port 0a
FC_switch_B_1	В	7	6	6	controller_B_2 port 0a

Zone in Fabric_1	Member ports
MC1_INIT_GRP_1_SITE_A_STK_GRP_1_TOP_FC1	5,2;5,6;7,2;7,6;5,8

DrGroup 1 : MC1_INIT_GRP_2_SITE_A_STK_GRP_1_TOP_FC1:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710 port	6520 port	G620 port	Connects to
FC_switch_A _1	A	5	3	3	3	controller_A_ 1 port 0c

FC_switch_A _1	A	5	7	7	7	controller_A_ 2 port 0c
FC_switch_A _1	A	5	9	9	9	bridge_A_1b FC1
FC_switch_B _1	В	7	3	3	3	controller_B_ 1 port 0c
FC_switch_B _1	В	7	7	7	7	controller_B_ 2 port 0c

Zone in Fabric_2	Member ports
MC1_INIT_GRP_2_SITE_A_STK_GRP_1_BOT_FC1	5,3;5,7;7,3;7,7;5,9

DrGroup 1 : MC1_INIT_GRP_1_SITE_A_STK_GRP_1_BOT_FC1:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710	6520	G620	Connects to
FC_switch_A _2	A	6	2	2	2	controller_A_ 1 port 0b
FC_switch_A _2	A	6	6	6	6	controller_A_ 2 port 0b
FC_switch_A _2	A	6	8	8	8	bridge_A_1a FC2
FC_switch_B _2	В	8	2	2	2	controller_B_ 1 port 0b
FC_switch_B _2	В	8	6	6	6	controller_B_ 2 port 0b

Zone in Fabric_1	Member ports
MC1_INIT_GRP_1_SITE_A_STK_GRP_1_TOP_FC2	6,2;6,6;8,2;8,6;6,8

DrGroup 1 : MC1_INIT_GRP_2_SITE_A_STK_GRP_1_BOT_FC2:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710	6520	G620	Connects to
FC_switch_A _2	A	6	3	3	3	controller_A_ 1 port 0d

FC_switch_A _2	A	6	7	7	7	controller_A_ 2 port 0d
FC_switch_A _2	A	6	9	9	9	bridge_A_1b FC2
FC_switch_B _2	В	8	3	3	3	controller_B_ 1 port 0d
FC_switch_B _2	В	8	7	7	7	controller_B_ 2 port 0d

Zone in Fabric_2	Member ports
MC1_INIT_GRP_2_SITE_A_STK_GRP_1_BOT_FC2	6,3;6,7;8,3;8,7;6,9

DR Group 1 - Stack 2 at Site_A

DrGroup 1 : MC1_INIT_GRP_1_SITE_A_STK_GRP_2_TOP_FC1:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710 port	6520 port	G620 port	Connects to
FC_switch_A _1	A	5	2	2	2	controller_A_ 1 port 0a
FC_switch_A _1	A	5	6	6	6	controller_A_ 2 port 0a
FC_switch_A _1	A	5	10	10	10	bridge_A_2a FC1
FC_switch_B _1	В	7	2	2	2	controller_B_ 1 port 0a
FC_switch_B _1	В	7	6	6	6	controller_B_ 2 port 0a

Zone in Fabric_1 hh	Member ports
MC1_INIT_GRP_1_SITE_A_STK_GRP_2_TOP_FC1	5,2;5,6;7,2;7,6;5,10

DrGroup 1 : MC1_INIT_GRP_2_SITE_A_STK_GRP_2_TOP_FC1:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710 port	6520 port	G620 port	Connects to
FC_switch_A _1	A	5	3	3	3	controller_A_ 1 port 0c
FC_switch_A_ 1	A	5	7	7	7	controller_A_ 2 port 0c
FC_switch_A_ 1	A	5	11	11	11	bridge_A_2b FC1
FC_switch_B _1	В	7	3	3	3	controller_B_ 1 port 0c
FC_switch_B _1	В	7	7	7	7	controller_B_ 2 port 0c

Zone in Fabric_2	Member ports		
MC1_INIT_GRP_2_SITE_A_STK_GRP_2_BOT_FC1	5,3;5,7;7,3;7,7;5,11		

DrGroup 1 : MC1_INIT_GRP_1_SITE_A_STK_GRP_2_BOT_FC2:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710 port	6520 port	G620 port	Connects to
FC_switch_A _2	A	6	2	0	0	controller_A_ 1 port 0b
FC_switch_A _2	A	6	6	4	4	controller_A_ 2 port 0b
FC_switch_A _2	A	6	10	10	10	bridge_A_2a FC2
FC_switch_B _2	В	8	2	2	2	controller_B_ 1 port 0b
FC_switch_B _2	В	8	6	6	6	controller_B_ 2 port 0b

Zone in Fabric_1	Member ports
MC1_INIT_GRP_1_SITE_A_STK_GRP_2_TOP_FC2	6,2;6,6;8,2;8,6;6,10

DrGroup 1 : MC1_INIT_GRP_2_SITE_A_STK_GRP_2_BOT_FC2:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710 port	6520 port	G620 port	Connects to
FC_switch_A _2	A	6	3	3	3	controller_A_ 1 port 0d
FC_switch_A _2	A	6	7	7	7	controller_A_ 2 port 0d
FC_switch_A _2	A	6	11	11	11	bridge_A_2b FC2
FC_switch_B _2	В	8	3	3	3	controller_B_ 1 port 0d
FC_switch_B _2	В	8	7	7	7	controller_B_ 2 port 0d

Zone in Fabric_2	Member ports
MC1_INIT_GRP_2_SITE_A_STK_GRP_2_BOT_FC2	6,3;6,7;8,3;8,7;6,11

DR Group 1 - Stack 1 at Site_B

DrGroup 1 : MC1_INIT_GRP_1_SITE_B_STK_GRP_1_TOP_FC1:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710 port	6520 port	G620 port	Connects to…
FC_switch_A _1	A	5	2	2	2	controller_A_ 1 port 0a
FC_switch_A _1	A	5	6	6	6	controller_A_ 2 port 0a
FC_switch_B _1	В	7	2	2	8	controller_B_ 1 port 0a
FC_switch_B _1	В	7	6	6	2	controller_B_ 2 port 0a
FC_switch_B _1	В	7	8	8	6	bridge_B_1a FC1

Zone in Fabric_1	Member ports
MC1_INIT_GRP_1_SITE_B_STK_GRP_1_TOP_FC1	5,2;5,6;7,2;7,6;7,8

DrGroup 1 : MC1_INIT_GRP_2_SITE_B_STK_GRP_1_TOP_FC1:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710 port	6520 port	G620 port	Connects to
FC_switch_A _1	A	5	3	3	3	controller_A_ 1 port 0c
FC_switch_A _1	A	5	7	7	7	controller_A_ 2 port 0c
FC_switch_B _1	В	7	3	3	9	controller_B_ 1 port 0c
FC_switch_B _1	В	7	7	7	3	controller_B_ 2 port 0c
FC_switch_B _1	В	7	9	9	7	bridge_B_1b FC1

Zone in Fabric_2	Member ports		
MC1_INIT_GRP_2_SITE_B_STK_GRP_1_BOT_FC1	5,3;5,7;7,3;7,7;7,9		

DrGroup 1 : MC1_INIT_GRP_1_SITE_B_STK_GRP_1_BOT_FC2:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710 port	6520 port	G620 port	Connects to
FC_switch_A _2	A	6	2	2	2	controller_A_ 1 port 0b
FC_switch_A _2	A	6	6	6	6	controller_A_ 2 port 0b
FC_switch_B _2	В	8	2	2	2	controller_B_ 1 port 0b
FC_switch_B _2	В	8	6	6	6	controller_B_ 2 port 0b

Zone in Fabric_1	Member ports
MC1_INIT_GRP_1_SITE_B_STK_GRP_1_TOP_FC2	6,2;6,6;8,2;8,6;8,8

DrGroup 1 : MC1_INIT_GRP_2_SITE_B_STK_GRP_1_BOT_FC2:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710 port	6520 port	G620 port	Connects to
FC_switch_A _2	A	6	3	3	3	controller_A_ 1 port 0d
FC_switch_A _2	A	6	7	7	7	controller_A_ 2 port 0d
FC_switch_B _2	В	8	3	3	3	controller_B_ 1 port 0d
FC_switch_B _2	В	8	7	7	7	controller_B_ 2 port 0d
FC_switch_B _2	В	8	9	9	9	bridge_A_1b FC2

Zone in Fabric_2	Member ports
MC1_INIT_GRP_2_SITE_B_STK_GRP_1_BOT_FC2	6,3;6,7;8,3;8,7;8,9

DR Group 1 - Stack 2 at Site_B

DrGroup 1 : MC1_INIT_GRP_1_SITE_B_STK_GRP_2_TOP_FC1:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710 port	6520 port	G620 port	Connects to…
FC_switch_A _1	A	5	2	2	2	controller_A_ 1 port 0a
FC_switch_A _1	A	5	6	6	6	controller_A_ 2 port 0a

FC_switch_B _1	В	7	2	2	2	controller_B_ 1 port 0a
FC_switch_B _1	В	7	6	6	6	controller_B_ 2 port 0a
FC_switch_B _1	В	7	10	10	10	bridge_B_2a FC1

Zone in Fabric_1	Member ports		
MC1_INIT_GRP_1_SITE_B_STK_GRP_2_TOP_FC1	5,2;5,6;7,2;7,6;7,10		

DrGroup 1 : MC1_INIT_GRP_2_SITE_B_STK_GRP_2_TOP_FC1:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710 port	6520 port	G620 port	Connects to…
FC_switch_A _1	A	5	3	3	3	controller_A_ 1 port 0c
FC_switch_A _1	A	5	7	7	7	controller_A_ 2 port 0c
FC_switch_B _1	В	7	3	3	3	controller_B_ 1 port 0c
FC_switch_B _1	В	7	7	7	7	controller_B_ 2 port 0c
FC_switch_B _1	В	7	11	11	11	bridge_B_2b FC1

Zone in Fabric_2 hh	Member ports
MC1_INIT_GRP_2_SITE_B_STK_GRP_2_BOT_FC1	5,3;5,7;7,3;7,7;7,11

DrGroup 1 : MC1_INIT_GRP_1_SITE_B_STK_GRP_2_BOT_FC2:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710 port	6520 port	G620 port	Connects to
FC_switch_A _2	A	6	2	2	2	controller_A_ 1 port 0b

FC_switch_A _2	A	6	6	6	6	controller_A_ 2 port 0b
FC_switch_B _2	В	8	2	2	2	controller_B_ 1 port 0b
FC_switch_B _2	В	8	6	6	6	controller_B_ 2 port 0b
FC_switch_B _2	В	8	10	10	10	bridge_B_2a FC2

Zone in Fabric_1	Member ports
MC1_INIT_GRP_1_SITE_B_STK_GRP_2_TOP_FC2	6,2;6,6;8,2;8,6;8,10

DrGroup 1 : MC1_INIT_GRP_2_SITE_B_STK_GRP_2_BOT_FC2:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710 port	6520 port	G620 port	Connects to…
FC_switch_A _2	A	6	3	3	3	controller_A_ 1 port 0d
FC_switch_A _2	A	6	7	7	7	controller_A_ 2 port 0d
FC_switch_B _2	В	8	3	3	3	controller_B_ 1 port 0d
FC_switch_B _2	В	8	7	7	7	controller_B_ 2 port 0d
FC_switch_B _2	В	8	11	11	11	bridge_B_2b FC2

Zone in Fabric_2	Member ports
MC1_INIT_GRP_2_SITE_B_STK_GRP_2_BOT_FC2	6,3;6,7;8,3;8,7;8,11

Summary of storage zones

Fabric	Zone name	Member ports
--------	-----------	--------------

FC_switch_A_1 and FC_switch_B_1	MC1_INIT_GRP_1_SITE_A_STK_ GRP_1_TOP_FC1	5,2;5,6;7,2;7,6;5,8
FC_switch_A_1 and FC_switch_B_1	MC1_INIT_GRP_2_SITE_A_STK_ GRP_1_BOT_FC1	5,3;5,7;7,3;7,7;5,9
FC_switch_A_1 and FC_switch_B_1	MC1_INIT_GRP_1_SITE_A_STK_ GRP_2_TOP_FC1	5,2;5,6;7,2;7,6;5,10
FC_switch_A_1 and FC_switch_B_1	MC1_INIT_GRP_2_SITE_A_STK_ GRP_2_BOT_FC1	5,3;5,7;7,3;7,7;5,11
FC_switch_A_1 and FC_switch_B_1	MC1_INIT_GRP_1_SITE_B_STK_ GRP_1_TOP_FC1	5,2;5,6;7,2;7,6;7,8
FC_switch_A_1 and FC_switch_B_1	MC1_INIT_GRP_2_SITE_B_STK_ GRP_1_BOT_FC1	5,3;5,7;7,3;7,7;7,9
FC_switch_A_1 and FC_switch_B_1	MC1_INIT_GRP_1_SITE_B_STK_ GRP_2_TOP_FC1	5,2;5,6;7,2;7,6;7,10
FC_switch_A_1 and FC_switch_B_1	MC1_INIT_GRP_2_SITE_B_STK_ GRP_2_BOT_FC1	5,3;5,7;7,3;7,7;7,11
FC_switch_A_2 and FC_switch_B_2	MC1_INIT_GRP_1_SITE_A_STK_ GRP_1_TOP_FC2	6,2;6,6;8,2;8,6;6,8
FC_switch_A_2 and FC_switch_B_2	MC1_INIT_GRP_2_SITE_A_STK_ GRP_1_BOT_FC2	6,3;6,7;8,3;8,7;6,9
FC_switch_A_2 and FC_switch_B_2	MC1_INIT_GRP_1_SITE_A_STK_ GRP_2_TOP_FC2	6,2;6,6;8,2;8,6;6,10
FC_switch_A_2 and FC_switch_B_2	MC1_INIT_GRP_2_SITE_A_STK_ GRP_2_BOT_FC2	6,3;6,7;8,3;8,7;6,11
FC_switch_A_2 and FC_switch_B_2	MC1_INIT_GRP_1_SITE_B_STK_ GRP_1_TOP_FC2	6,2;6,6;8,2;8,6;8,8
FC_switch_A_2 and FC_switch_B_2	MC1_INIT_GRP_2_SITE_B_STK_ GRP_1_BOT_FC2	6,3;6,7;8,3;8,7;8,9
FC_switch_A_2 and FC_switch_B_2	MC1_INIT_GRP_1_SITE_B_STK_ GRP_2_TOP_FC2	6,2;6,6;8,2;8,6;8,10
FC_switch_A_2 and FC_switch_B_2	MC1_INIT_GRP_2_SITE_B_STK_ GRP_2_BOT_FC2	6,3;6,7;8,3;8,7;8,11

Zone the Brocade FC switches

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

About this task

The following steps use the standard zoning for the MetroCluster configuration.

Zoning for FC-VI ports

Zoning for FibreBridge 7500N or 7600N bridges using one FC port

Zoning for FibreBridge 7500N bridges using both FC ports

Steps

1. Create the FC-VI zones on each switch:

```
zonecreate "QOSH1 FCVI 1", member;member ...
```

In this example a QOS FCVI zone is created containing ports 5,0;5,1;5,4;5,5;7,0;7,1;7,4;7,5:

```
Switch_A_1:admin> zonecreate "QOSH1_FCVI_1",
"5,0;5,1;5,4;5,5;7,0;7,1;7,4;7,5"
```

2. Configure the storage zones on each switch.

You can configure zoning for the fabric from one switch in the fabric. In the example that follows, zoning is configured on Switch_A_1.

a. Create the storage zone for each switch domain in the switch fabric:

zonecreate name, member;member ...

In this example a storage zone for a FibreBridge 7500N using both FC ports is being created. The zones contains ports 5,2;5,6;7,2;7,6;5,16:

Switch_A_1:admin> zonecreate
"MC1_INIT_GRP_1_SITE_A_STK_GRP_1_TOP_FC1", "5,2;5,6;7,2;7,6;5,16"

b. Create the configuration in the first switch fabric:

```
cfgcreate config name, zone; zone ...
```

In this example a configuration with the name CFG_1 and the two zones QOSH1_MC1_FAB_1_FCVI and MC1_INIT_GRP_1_SITE_A_STK_GRP_1_TOP_FC1 is created

Switch_A_1:admin> cfgcreate "CFG_1", "QOSH1_MC1_FAB_1_FCVI; MC1 INIT GRP 1 SITE A STK GRP 1 TOP FC1" c. Add zones to the configuration, if desired:

cfgadd config_namezone;zone...

d. Enable the configuration:

cfgenable config name

Switch_A_1:admin> cfgenable "CFG_1"

e. Save the configuration:

cfgsave

Switch A 1:admin> cfgsave

f. Validate the zoning configuration:

zone --validate

```
Switch A 1:admin> zone --validate
Defined configuration:
cfg: CFG 1 QOSH1 MC1 FAB 1 FCVI ;
MC1_INIT_GRP_1_SITE A STK GRP 1 TOP FC1
zone: QOSH1 MC1 FAB 1 FCVI
5,0;5,1;5,4;5,5;7,0;7,1;7,4;7,5
zone: MC1 INIT GRP 1 SITE A STK GRP 1 TOP FC1
5,2;5,6;7,2;7,6;5,16
Effective configuration:
cfg: CFG 1
zone: QOSH1 MC1 FAB 1 FCVI
5,0
5,1
5,4
5,5
7,0
7,1
7,4
7,5
zone: MC1 INIT GRP 1 SITE A STK GRP 1 TOP FC1
5,2
5,6
7,2
7,6
5,16
_____
~ - Invalid configuration
* - Member does not exist
# - Invalid usage of broadcast zone
```

Set ISL encryption on Brocade 6510 or G620 switches

On Brocade 6510 or G620 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 or G620 switches.



Support for ISL encryption on Brocade G620 switches is only supported on ONTAP 9.4 and later.

- You must have selected two switches from the same fabric.
- You must have reviewed the Brocade documentation for your switch and Fabric Operating System version to confirm the bandwidth and port limits.

About this task

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

Disable 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.

Steps

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

```
fosconfig --disable vf
```

2. Reboot the switch.

Set 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: 2048

Set the authentication policy

You must set the authentication policy and associated parameters.

About this task

The commands must be executed at the switch console.

Steps

- 1. Set the authentication secret:
 - a. Begin the setup process:

secAuthSecret --set

This command initiates a series of prompts that you respond to in the following steps:

- b. Provide the worldwide name (WWN) of the other switch in the fabric for the "Enter peer WWN, Domain, or switch name" parameter.
- c. Provide the peer secret for the "Enter peer secret" parameter.
- d. Provide the local secret for the "Enter local secret" parameter.
- e. Enter y for the "Are you done" parameter.

The following is an example of setting the authentication secret:

```
brcd> 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.
```

2. Set the authentication group to 4:

authUtil --set -g 4

3. Set the authentication type to "dhchap":

authUtil --set -a dhchap

The system displays the following output:

Authentication is set to dhchap.

4. Set the authentication policy on the switch to on:

authUtil --policy -sw on

The system 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

Enable ISL encryption on Brocade switches

After setting the authentication policy and the authentication secret, you must enable ISL encryption on the ports for it to take effect.

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. Enable encryption on all of the ISL ports:

portCfgEncrypt --enable port_number

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

portCfgEncrypt --enable 8

portCfgEncrypt --enable 12

2. Enable the switch:

switchenable

3. Verify that the ISL is up and working:

islshow

4. Verify that encryption is enabled:

portenccompshow

The following example shows that 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	

What to do next

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

Configuring the Cisco FC switches manually

Each Cisco switch in the MetroCluster configuration must be configured appropriately for the ISL and storage connections.

Before you begin

The following requirements apply to the Cisco FC switches:

- You must use four supported Cisco switches of the same model with the same NX-OS version and licensing.
- 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.

- The switch must support connectivity to the ATTO FibreBridge model.
- You cannot use encryption or compression in the Cisco FC storage fabric. It is not supported in the MetroCluster configuration.

In the NetApp Interoperability Matrix Tool (IMT), you can use the Storage Solution field to select your MetroCluster solution. You use the **Component Explorer** to select the components and ONTAP version to refine your search. You can click **Show Results** to display the list of supported configurations that match the criteria.

About this task

The following requirement applies to the Inter-Switch Link (ISL) connections:

• All ISLs must have the same length and same speed in one fabric.

Different lengths of ISLs can be used in the different fabrics. The same speed must be used in all fabrics.

The following requirement applies to the storage connections:

• Each storage controller must have four initiator ports available to connect to the switch fabrics.

Two initiator ports must be connected from each storage controller to each fabric.

You can configure FAS8020, AFF8020, FAS8200, and AFF A300 systems with two initiators ports per controller (a single initiator port to each fabric) if all of the following criteria are met:



• There are fewer than four FC initiator ports available to connect the disk storage and no additional ports can be configured as FC initiators.

• All slots are in use and no FC initiator card can be added.

Related information

NetApp Interoperability Matrix Tool

Cisco switch license requirements

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

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

• ENTERPRISE_PKG

This license enables you to use the QoS feature on 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 license 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 and threshold monitoring. 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 with the installation.



The Cisco MDS 9250i switches have two fixed 1/10 GbE IP storage services ports. No additional licenses are required for these ports. The Cisco SAN Extension over IP application package is a standard license on these switches that enables features such as FCIP and compression.

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. 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. 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.

Example

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, configure the basic switch settings:

setup

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 ONTAP Health Monitors.

You can set the community string to a value other than "public", but you must configure the ONTAP Health Monitors using the community string you specify.

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.

Before you begin

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. Display the 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. Enter configuration mode:

config t

- 3. Remove the license from the selected port:
 - a. Select the port to be unlicensed:

interface interface-name

b. Remove the license from the port:

no port-license acquire

c. Exit the port configuration interface:

exit

- 4. Acquire the license for the selected port:
 - a. Select the port to be unlicensed:

interface interface-name

b. Make the port eligible to acquire a license:

port-license

c. Acquire the license on the port:

port-license acquire

d. Exit the port configuration interface:

exit

- 5. Repeat for any additional ports.
- 6. 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 or 9148S switch

In Cisco MDS 9148 or 9148S 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 or 9148S 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:

```
show port-resources module blade number
```
2. License and acquire the required port in a port group:

```
config t
interface port_number
shut
port-license acquire
no shut
```

For example, the following command sequence 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.

About this task

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.

Refer to the following sections to identify the F-ports (switch-to-node) for your configuration:

• Port assignments for FC switches

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

Steps

1. 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:

switchport mode F

5. Set the ports to fixed speed:

switchport speed speed-value

speed-value is either 8000 or 16000

6. Set the rate mode of the switch port to dedicated:

switchport rate-mode dedicated

7. Restart the port:

no shutdown

8. Exit configuration mode:

end

Example

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.

About this task

This task is not required if the F-ports are not in the same port group as the ISL port.

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. Enter configuration mode:

config t

2. 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:

switchport mode F

5. Set the buffer-to-buffer credit of the non-E ports to 1:

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:

show port-resources module 1

10. Exit configuration mode:

exit

- 11. Repeat these steps on the other switch in the fabric.
- 12. Verify the settings:

Example

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 1
FC 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 B2B Credit Bandwidth Rate Mode
                       Buffers (Gbps)
_____
                           32
                                8.0 dedicated
fc1/37
fc1/38
                            1
                                 8.0 dedicated
                            1 8.0 dedicated
fc1/39
. . .
FC switch B 1# port-resource module
. . .
Port-Group 11
Available dedicated buffers are 93
_____
Interfaces in the Port-Group B2B Credit Bandwidth Rate Mode
                      Buffers (Gbps)
_____
                                 8.0 dedicated
fc1/37
                           32
fc1/38
                            1
                                 8.0 dedicated
fc1/39
                                 8.0 dedicated
                            1
. . .
```

Creating and configuring 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.

About this task

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.

The examples of this task 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. Enter configuration mode if you have not done so already:

config t

b. Edit the VSAN database:

vsan database

c. Set the VSAN ID:

vsan *vsan-ID*

d. Set the VSAN name:

vsan vsan-ID name vsan_name

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

vsan vsan-ID interface interface_name

For the FC-VI VSAN, the ports connecting the local FC-VI ports will be added.

b. Exit configuration mode:

end

c. Copy the running-config to the startup-config:

copy running-config startup-config

In the following example, the ports are fc1/1 and fc1/13:

```
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_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(config) # end
FC_switch_B_1(config) # end
FC_switch_B_1(config) # end
```

3. Verify port membership of the VSAN:

show vsan member

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 or out-of-order delivery of frames:

The standard IOD settings are recommended. You should configure OOD only if necessary.

Considerations for using TDM/WDM equipment with fabric-attached MetroCluster configurations

- The following steps must be performed to configure in-order delivery of frames:
 - a. Enter configuration mode:

conf t

b. Enable the in-order guarantee of exchanges for the VSAN:

in-order-guarantee vsan vsan-ID



For FC-VI VSANs (FCVI_1_10 and FCVI_2_30), you must enable in-order guarantee of frames and exchanges only on VSAN 10.

c. Enable load balancing for the VSAN:

vsan vsan-ID loadbalancing src-dst-id

d. Exit configuration mode:

end

e. Copy the running-config to the startup-config:

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

The commands to configure in-order delivery of frames 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 commands to configure in-order delivery of frames 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
```

- The following steps must be performed to configure out-of-order delivery of frames:
 - a. Enter configuration mode:

conf t

b. Disable the in-order guarantee of exchanges for the VSAN:

```
no in-order-guarantee vsan vsan-ID
```

c. Enable load balancing for the VSAN:

vsan vsan-ID loadbalancing src-dst-id

d. Exit configuration mode:

end

e. Copy the running-config to the startup-config:

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

The commands to configure out-of-order delivery of frames on FC_switch_A_1:

```
FC_switch_A_1# config t
FC_switch_A_1(config) # no 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 commands to configure out-of-order delivery of frames on FC_switch_B_1:

```
FC_switch_B_1# config t
FC_switch_B_1(config)# no 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
```



When configuring ONTAP on the controller modules, OOD must be explicitly configured on each controller module in the MetroCluster configuration.

Configuring in-order delivery or out-of-order delivery of frames on ONTAP software

- 5. Set QoS policies for the FC-VI VSAN:
 - a. Enter configuration mode:

conf t

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

```
qos enable
```

qos class-map class_name match-any

c. Add the class map created in a previous step to the policy map:

class *class_name*

d. Set the priority:

priority high

e. Add the VSAN to the policy map created previously in this procedure:

qos service policy policy_name vsan vsan-id

f. Copy the updated configuration to the startup configuration:

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

The commands to set the QoS policies 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 commands to set the QoS policies 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:

vsan *vsan-ID*

b. Set the VSAN name:

vsan vsan-ID name vsan name

The commands to configure the storage VSAN 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 commands to configure the storage VSAN 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.

The commands to add ports to the storage VSAN 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 fc1/17
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 fc1/29
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) # vsan 20 interface fc1/37
FC_switch_A_1(config) # end
FC_switch_A_1(config) # end
```

The commands to add ports to the storage VSAN 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 fc1/17
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 fc1/25
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) # vsan 20 interface fc1/37
FC_switch_B_1(config) # end
FC_switch_B_1(config) # end
```

Configuring E-ports

You must configure the switch ports that connect the ISL (these are the E-Ports).

About this task

The procedure you use depends on which switch you are using:

- Configuring the E-ports on the Cisco FC switch
- Configuring FCIP ports for a single ISL on Cisco 9250i FC switches
- Configuring FCIP ports for a dual ISL on Cisco 9250i FC switches

Configuring the E-ports on the Cisco FC switch

You must configure the FC switch ports that connect the inter-switch link (ISL).

About this task

These are the E-ports, and configuration must be done for each port. To do so, you must calculate the correct number of buffer-to-buffer credits (BBCs).

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 following table) by the distance in kilometers between the switches. An 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

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
- $^{\circ}$ Distance in kilometers between switches is 30 km $\,$
- 3 x 30 = 90
 - 1. Enter configuration mode:

config t

2. Specify the port you are configuring:

interface port-name

3. Shut down the port:

shutdown

4. Set the rate mode of the port to "dedicated":

switchport rate-mode dedicated

5. Set the speed for the port:

switchport speed speed-value

6. Set the buffer-to-buffer credits for the port:

switchport fcrxbbcredit number_of_buffers

7. Set the port to E mode:

switchport mode E

8. Enable the trunk mode for the port:

switchport trunk mode on

9. Add the ISL virtual storage area networks (VSANs) to the trunk:

switchport trunk allowed vsan 10

switchport trunk allowed vsan add 20

10. Add the port to port channel 1:

channel-group 1

11. Repeat the previous steps for the matching ISL port on the partner switch in the fabric.

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
```

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

no shutdown

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

interface port-channel number

switchport trunk allowed vsan add native san id

15. 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).

```
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] : fc1/36
    Member[2] : fc1/40
    Interface last changed at Thu Oct 16 11:48:00 2014
```

1. Exit interface configuration on both switches:

```
end
```

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

```
copy running-config startup-config

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
```

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

Related information

You need to verify that you are using the specified port assignments when you cable the FC switches. Refer to Port assignments for FC switches

Configuring FCIP ports for a single ISL on Cisco 9250i FC switches

You must configure the FCIP switch ports that connect the ISL (E-ports) by creating FCIP profiles and interfaces, and then assigning them to the IPStorage1/1 GbE interface.

About this task

This task is only for configurations using a single ISL per switch fabric, using the IPStorage1/1 interface on each switch.

This task must be performed on each FC switch.

Two FCIP profiles are created on each switch:

- Fabric 1
 - FC_switch_A_1 is configured with FCIP profiles 11 and 111.
 - FC_switch_B_1 is configured with FCIP profiles 12 and 121.
- Fabric 2
 - $\circ\,$ FC_switch_A_2 is configured with FCIP profiles 13 and 131.
 - FC_switch_B_2 is configured with FCIP profiles 14 and 141.

Steps

1. Enter configuration mode:

config t

2. Enable FCIP:

feature fcip

3. Configure the IPStorage1/1 GbE interface:

a. Enter configuration mode:

conf t

b. Specify the IPStorage1/1 interface:

interface IPStorage1/1

c. Specify the IP address and subnet mask:

interface ip-address subnet-mask

d. Specify the MTU size of 2500:

switchport mtu 2500

e. Enable the port:

no shutdown

f. Exit configuration mode:

exit

The following example shows the configuration of an IPStorage1/1 port:

```
conf t
interface IPStorage1/1
  ip address 192.168.1.201 255.255.255.0
  switchport mtu 2500
  no shutdown
exit
```

- 4. Configure the FCIP profile for FC-VI traffic:
 - a. Configure an FCIP profile and enter FCIP profile configuration mode:

fcip profile FCIP-profile-name

The profile name depends on which switch is being configured.

b. Assign the IP address of the IPStorage1/1 interface to the FCIP profile:

ip address *ip-address*

c. Assign the FCIP profile to TCP port 3227:

port 3227

d. Set the TCP settings:

```
tcp keepalive-timeout 1
tcp max-retransmissions 3
max-bandwidth-mbps 5000 min-available-bandwidth-mbps 4500 round-trip-time-ms
3
tcp min-retransmit-time 200
tcp keepalive-timeout 1
tcp pmtu-enable reset-timeout 3600
tcp sack-enable``no tcp cwm
```

The following example shows the configuration of the FCIP profile:

```
conf t
fcip profile 11
  ip address 192.168.1.333
  port 3227
  tcp keepalive-timeout 1
tcp max-retransmissions 3
max-bandwidth-mbps 5000 min-available-bandwidth-mbps 4500 round-trip-
time-ms 3
  tcp min-retransmit-time 200
  tcp keepalive-timeout 1
  tcp pmtu-enable reset-timeout 3600
  tcp sack-enable
  no tcp cwm
```

- 5. Configure the FCIP profile for storage traffic:
 - a. Configure an FCIP profile with the name 111 and enter FCIP profile configuration mode:

fcip profile 111

b. Assign the IP address of the IPStorage1/1 interface to the FCIP profile:

ip address *ip-address*

c. Assign the FCIP profile to TCP port 3229:

port 3229

d. Set the TCP settings:

```
tcp keepalive-timeout 1
```

```
tcp max-retransmissions 3
max-bandwidth-mbps 5000 min-available-bandwidth-mbps 4500 round-trip-time-ms
3
tcp min-retransmit-time 200
tcp keepalive-timeout 1
tcp pmtu-enable reset-timeout 3600
tcp sack-enable``no tcp cwm
```

The following example shows the configuration of the FCIP profile:

```
conf t
fcip profile 111
  ip address 192.168.1.334
  port 3229
  tcp keepalive-timeout 1
tcp max-retransmissions 3
max-bandwidth-mbps 5000 min-available-bandwidth-mbps 4500 round-trip-
time-ms 3
  tcp min-retransmit-time 200
  tcp keepalive-timeout 1
  tcp pmtu-enable reset-timeout 3600
  tcp sack-enable
  no tcp cwm
```

6. Create the first of two FCIP interfaces:

```
interface fcip 1
```

This interface is used for FC-IV traffic.

a. Select the profile 11 created previously:

```
use-profile 11
```

b. Set the IP address and port of the IPStorage1/1 port on the partner switch:

peer-info ipaddr partner-switch-port-ip port 3227

c. Select TCP connection 2:

tcp-connection 2

d. Disable compression:

```
no ip-compression
```

e. Enable the interface:

no shutdown

f. Configure the control TCP connection to 48 and the data connection to 26 to mark all packets on that differentiated services code point (DSCP) value:

```
qos control 48 data 26
```

g. Exit the interface configuration mode:

exit

The following example shows the configuration of the FCIP interface:

```
interface fcip 1
  use-profile 11
# the port # listed in this command is the port that the remote switch
is listening on
  peer-info ipaddr 192.168.32.334   port 3227
  tcp-connection 2
  no ip-compression
  no shutdown
  gos control 48 data 26
exit
```

7. Create the second of two FCIP interfaces:

interface fcip 2

This interface is used for storage traffic.

a. Select the profile 111 created previously:

use-profile 111

b. Set the IP address and port of the IPStorage1/1 port on the partner switch:

peer-info ipaddr partner-switch-port-ip port 3229

c. Select TCP connection 2:

tcp-connection 5

d. Disable compression:

no ip-compression

e. Enable the interface:

no shutdown

f. Configure the control TCP connection to 48 and data connection to 26 to mark all packets on that differentiated services code point (DSCP) value:

qos control 48 data 26

g. Exit the interface configuration mode:

exit

The following example shows the configuration of the FCIP interface:

```
interface fcip 2
  use-profile 11
# the port # listed in this command is the port that the remote switch
is listening on
  peer-info ipaddr 192.168.32.33e port 3229
  tcp-connection 5
  no ip-compression
  no shutdown
   gos control 48 data 26
exit
```

- 8. Configure the switchport settings on the fcip 1 interface:
 - a. Enter configuration mode:

config t

b. Specify the port you are configuring:

interface fcip 1

c. Shut down the port:

shutdown

d. Set the port to E mode:

switchport mode E

e. Enable the trunk mode for the port:

switchport trunk mode on

f. Set the trunk allowed vsan to 10:

switchport trunk allowed vsan 10

g. Set the speed for the port:

switchport speed speed-value

- 9. Configure the switchport settings on the fcip 2 interface:
 - a. Enter configuration mode:

config t

b. Specify the port you are configuring:

interface fcip 2

c. Shut down the port:

shutdown

d. Set the port to E mode:

switchport mode E

e. Enable the trunk mode for the port:

switchport trunk mode on

f. Set the trunk allowed vsan to 20:

switchport trunk allowed vsan 20

g. Set the speed for the port:

switchport speed speed-value

10. Repeat the previous steps on the second switch.

The only differences are the appropriate IP addresses and unique FCIP profile names.

- When configuring the first switch fabric, FC_switch_B_1 is configured with FCIP profiles 12 and 121.
- When configuring the first switch fabric, FC_switch_A_2 is configured with FCIP profiles 13 and 131 and FC_switch_B_2 is configured with FCIP profiles 14 and 141.
- 11. Restart the ports on both switches:

no shutdown

12. Exit the interface configuration on both switches:

end

13. Copy the updated configuration to the startup configuration on both switches:

copy running-config startup-config

```
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
```

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

Configuring FCIP ports for a dual ISL on Cisco 9250i FC switches

You must configure the FCIP switch ports that connect the ISL (E-ports) by creating FCIP profiles and interfaces, and then assigning them to the IPStorage1/1 and IPStorage1/2 GbE interfaces.

About this task

This task is only for configurations that use a dual ISL per switch fabric, using the IPStorage1/1 and IPStorage1/2 GbE interfaces on each switch.

This task must be performed on each FC switch.



The task and examples use the following profile configuration tables:

- Fabric 1 profile configuration table
- Fabric 2 profile configuration table

Fabric 1 profile configuration table

Switch	IPStorage	IP	Port type	FCIP	FCIP	Port	Peer	VSAN ID
fabric	interface	Address		interface	profile		IP/port	

FC_switch _A_1	IPStorage 1/1	a.a.a.a	FC-VI	fcip 1	15	3220	c.c.c/323 0	10
			Storage	fcip 2	20	3221	c.c.c.c/323 1	20
	IPStorage 1/2	b.b.b.b	FC-VI	fcip 3	25	3222	d.d.d.d/32 32	10
			Storage	fcip 4	30	3223	d.d.d.d/32 33	20
FC_switch _B_1	IPStorage 1/1	C.C.C.C	FC-VI	fcip 1	15	3230	a.a.a.a/32 20	10
			Storage	fcip 2	20	3231	a.a.a.a/32 21	20
	IPStorage 1/2	ge d.d.d.d	FC-VI	fcip 3	25	3232	b.b.b.b/32 22	10
			Storage	fcip 4	30	3233	b.b.b.b/32 23	20

Fabric 2 profile configuration table

Switch fabric	IPStorage interface	IP Address	Port type	FCIP interface	FCIP profile	Port	Peer IP/port	VSAN ID
FC_switch _A_2	IPStorage 1/1	e.e.e	FC-VI	fcip 1	15	3220	g.g.g.g/32 30	10
			Storage	fcip 2	20	3221	g.g.g.g/32 31	20
	IPStorage f. 1/2	f.f.f	FC-VI	fcip 3	25	3222	h.h.h.h/32 32	10
			Storage	fcip 4	30	3223	h.h.h.h/32 33	20

FC_switch _B_2	IPStorage 1/1	g.g.g.g	FC-VI	fcip 1	15	3230	e.e.e.e/32 20	10
			Storage	fcip 2	20	3231	e.e.e.e/32 21	20
	IPStorage 1/2	h.h.h.h	FC-VI	fcip 3	25	3232	f.f.f.f/3222	10
			Storage	fcip 4	30	3233	f.f.f.f/3223	20

Steps

1. Enter configuration mode:

config t

2. Enable FCIP:

feature fcip

- 3. On each switch, configure the two IPStorage interfaces ("IPStorage1/1" and "IPStorage1/2"):
 - a. Enter configuration mode:

conf t

b. Specify the IPStorage interface to create:

interface *ipstorage*

The *ipstorage* parameter value is "IPStorage1/1" or "IPStorage1/2".

c. Specify the IP address and subnet mask of the IPStorage interface previously specified:

interface *ip-address* subnet-mask



On each switch, the IPStorage interfaces "IPStorage1/1" and "IPStorage1/2" must have different IP addresses.

d. Specify the MTU size as 2500:

switchport mtu 2500

e. Enable the port:

no shutdown

f. Exit configuration mode:

exit

g. Repeat Substep "a" through Substep "f" to configure the IPStorage1/2 GbE interface with a different IP

address.

- 4. Configure the FCIP profiles for FC-VI and storage traffic with the profile names given in the profile configuration table:
 - a. Enter configuration mode:

conf t

b. Configure the FCIP profiles with the following profile names:

fcip profile FCIP-profile-name

The following list provides the values for the *FCIP*-profile-name parameter:

- 15 for FC-VI on IPStorage1/1
- 20 for storage on IPStorage1/1
- 25 for FC-VI on IPStorage1/2
- 30 for storage on IPStorage1/2
- c. Assign the FCIP profile ports according to the profile configuration table:

port port_number

d. Set the TCP settings:

```
tcp keepalive-timeout 1
```

tcp max-retransmissions 3

```
max-bandwidth-mbps 5000 min-available-bandwidth-mbps 4500 round-trip-time-ms
3
```

tcp min-retransmit-time 200

tcp keepalive-timeout 1

tcp pmtu-enable reset-timeout 3600

tcp sack-enable

no tcp cwm

5. Create FCIP interfaces:

interface fcip FCIP_interface

The FCIP interface parameter value is "1", "2", "3", or "4" as shown in the profile configuration table.

a. Map interfaces to the previously created profiles:

use-profile profile

b. Set the peer IP address and peer profile port number:

peer-info peer IPstorage ipaddr port peer_profile_port_number

c. Select the TCP connections:

tcp-connection connection-#

The *connection*-# parameter value is "2" for FC-VI profiles and "5" for storage profiles.

d. Disable compression:

```
no ip-compression
```

e. Enable the interface:

no shutdown

f. Configure the control TCP connection to "48" and the data connection to "26" to mark all packets that have differentiated services code point (DSCP) value:

qos control 48 data 26

g. Exit configuration mode:

exit

- 6. Configure the switchport settings on each FCIP interface:
 - a. Enter configuration mode:

config t

b. Specify the port that you are configuring:

interface fcip 1

c. Shut down the port:

shutdown

d. Set the port to E mode:

switchport mode E

e. Enable the trunk mode for the port:

switchport trunk mode on

f. Specify the trunk that is allowed on a specific VSAN:

switchport trunk allowed vsan vsan_id

The vsan_id parameter value is "VSAN 10" for FC-VI profiles and "VSAN 20" for storage profiles.

g. Set the speed for the port:

switchport speed speed-value

```
h. Exit configuration mode:
```

exit

7. Copy the updated configuration to the startup configuration on both switches:

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

The following examples show the configuration of FCIP ports for a dual ISL in fabric 1 switches FC_switch_A_1 and FC_switch_B_1.

For FC_switch_A_1:

```
FC switch A 1# config t
FC switch A 1(config) # no in-order-guarantee vsan 10
FC switch A 1(config-vsan-db) # end
FC switch A 1# copy running-config startup-config
# fcip settings
feature fcip
conf t
interface IPStorage1/1
# IP address: a.a.a.a
# Mask: y.y.y.y
 ip address <a.a.a y.y.y.y>
 switchport mtu 2500
 no shutdown
exit
conf t
fcip profile 15
 ip address <a.a.a.a>
 port 3220
 tcp keepalive-timeout 1
tcp max-retransmissions 3
max-bandwidth-mbps 5000 min-available-bandwidth-mbps 4500 round-trip-time-
ms 3
 tcp min-retransmit-time 200
 tcp keepalive-timeout 1
 tcp pmtu-enable reset-timeout 3600
 tcp sack-enable
 no tcp cwm
conf t
fcip profile 20
```

```
ip address <a.a.a.a>
  port 3221
  tcp keepalive-timeout 1
tcp max-retransmissions 3
max-bandwidth-mbps 5000 min-available-bandwidth-mbps 4500 round-trip-time-
ms 3
 tcp min-retransmit-time 200
 tcp keepalive-timeout 1
 tcp pmtu-enable reset-timeout 3600
 tcp sack-enable
  no tcp cwm
conf t
interface IPStorage1/2
# IP address: b.b.b.b
# Mask: y.y.y.y
 ip address <b.b.b.b y.y.y.y>
  switchport mtu 2500
 no shutdown
exit
conf t
fcip profile 25
 ip address <b.b.b>
 port 3222
tcp keepalive-timeout 1
tcp max-retransmissions 3
max-bandwidth-mbps 5000 min-available-bandwidth-mbps 4500 round-trip-time-
ms 3
 tcp min-retransmit-time 200
 tcp keepalive-timeout 1
 tcp pmtu-enable reset-timeout 3600
 tcp sack-enable
 no tcp cwm
conf t
fcip profile 30
 ip address <b.b.b>
 port 3223
tcp keepalive-timeout 1
tcp max-retransmissions 3
max-bandwidth-mbps 5000 min-available-bandwidth-mbps 4500 round-trip-time-
ms 3
 tcp min-retransmit-time 200
 tcp keepalive-timeout 1
  tcp pmtu-enable reset-timeout 3600
```

```
tcp sack-enable
 no tcp cwm
interface fcip 1
 use-profile 15
# the port # listed in this command is the port that the remote switch is
listening on
peer-info ipaddr <c.c.c.> port 3230
 tcp-connection 2
 no ip-compression
 no shutdown
 qos control 48 data 26
exit
interface fcip 2
 use-profile 20
# the port # listed in this command is the port that the remote switch is
listening on
 peer-info ipaddr <c.c.c.> port 3231
 tcp-connection 5
 no ip-compression
 no shutdown
 qos control 48 data 26
exit
interface fcip 3
 use-profile 25
# the port # listed in this command is the port that the remote switch is
listening on
peer-info ipaddr < d.d.d.d > port 3232
 tcp-connection 2
 no ip-compression
 no shutdown
 qos control 48 data 26
exit
interface fcip 4
 use-profile 30
# the port # listed in this command is the port that the remote switch is
listening on
peer-info ipaddr < d.d.d.d > port 3233
 tcp-connection 5
 no ip-compression
 no shutdown
 gos control 48 data 26
exit
```

```
conf t
interface fcip 1
shutdown
switchport mode E
switchport trunk mode on
switchport trunk allowed vsan 10
no shutdown
exit
conf t
interface fcip 2
shutdown
switchport mode E
switchport trunk mode on
switchport trunk allowed vsan 20
no shutdown
exit
conf t
interface fcip 3
shutdown
switchport mode E
switchport trunk mode on
switchport trunk allowed vsan 10
no shutdown
exit
conf t
interface fcip 4
shutdown
switchport mode E
switchport trunk mode on
switchport trunk allowed vsan 20
no shutdown
exit
```

For FC_switch_B_1:

```
FC_switch_A_1# config t
FC_switch_A_1(config)# in-order-guarantee vsan 10
FC_switch_A_1(config-vsan-db)# end
FC_switch_A_1# copy running-config startup-config
# fcip settings
```

```
feature fcip
conf t
interface IPStorage1/1
# IP address: c.c.c.c
# Mask: y.y.y.y
  ip address <c.c.c y.y.y.y>
  switchport mtu 2500
  no shutdown
exit
conf t
fcip profile 15
 ip address <c.c.c.<>
 port 3230
 tcp keepalive-timeout 1
tcp max-retransmissions 3
max-bandwidth-mbps 5000 min-available-bandwidth-mbps 4500 round-trip-time-
ms 3
 tcp min-retransmit-time 200
 tcp keepalive-timeout 1
 tcp pmtu-enable reset-timeout 3600
 tcp sack-enable
 no tcp cwm
conf t
fcip profile 20
 ip address <c.c.c.<>
 port 3231
 tcp keepalive-timeout 1
tcp max-retransmissions 3
max-bandwidth-mbps 5000 min-available-bandwidth-mbps 4500 round-trip-time-
ms 3
 tcp min-retransmit-time 200
 tcp keepalive-timeout 1
 tcp pmtu-enable reset-timeout 3600
 tcp sack-enable
 no tcp cwm
conf t
interface IPStorage1/2
# IP address: d.d.d.d
# Mask: y.y.y.y
  ip address <b.b.b.b y.y.y.y>
  switchport mtu 2500
  no shutdown
```

```
exit
conf t
fcip profile 25
  ip address <d.d.d.d>
  port 3232
tcp keepalive-timeout 1
tcp max-retransmissions 3
max-bandwidth-mbps 5000 min-available-bandwidth-mbps 4500 round-trip-time-
ms 3
  tcp min-retransmit-time 200
 tcp keepalive-timeout 1
 tcp pmtu-enable reset-timeout 3600
  tcp sack-enable
  no tcp cwm
conf t
fcip profile 30
  ip address <d.d.d.d>
  port 3233
tcp keepalive-timeout 1
tcp max-retransmissions 3
max-bandwidth-mbps 5000 min-available-bandwidth-mbps 4500 round-trip-time-
ms 3
  tcp min-retransmit-time 200
 tcp keepalive-timeout 1
  tcp pmtu-enable reset-timeout 3600
 tcp sack-enable
  no tcp cwm
interface fcip 1
  use-profile 15
# the port # listed in this command is the port that the remote switch is
listening on
peer-info ipaddr <a.a.a.a> port 3220
 tcp-connection 2
 no ip-compression
 no shutdown
  gos control 48 data 26
exit
interface fcip 2
  use-profile 20
# the port # listed in this command is the port that the remote switch is
listening on
 peer-info ipaddr <a.a.a.a> port 3221
```

```
tcp-connection 5
  no ip-compression
 no shutdown
  qos control 48 data 26
exit
interface fcip 3
  use-profile 25
# the port # listed in this command is the port that the remote switch is
listening on
 peer-info ipaddr < b.b.b.b > port 3222
 tcp-connection 2
 no ip-compression
 no shutdown
 qos control 48 data 26
exit
interface fcip 4
 use-profile 30
# the port # listed in this command is the port that the remote switch is
listening on
peer-info ipaddr < b.b.b.b > port 3223
 tcp-connection 5
 no ip-compression
 no shutdown
 qos control 48 data 26
exit
conf t
interface fcip 1
shutdown
switchport mode E
switchport trunk mode on
switchport trunk allowed vsan 10
no shutdown
exit
conf t
interface fcip 2
shutdown
switchport mode E
switchport trunk mode on
switchport trunk allowed vsan 20
no shutdown
exit
```

conf t interface fcip 3 shutdown switchport mode E switchport trunk mode on switchport trunk allowed vsan 10 no shutdown exit conf t interface fcip 4 shutdown switchport mode E switchport trunk mode on switchport trunk allowed vsan 20 no shutdown exit

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.

The following steps use the zoning described in the section Zoning for a FibreBridge 7500N in a four-node MetroCluster configuration. Refer to Zoning for FC-VI ports.

Steps

- 1. Clear the existing zones and zone set, if present.
 - a. Determine which zones and zone sets are active:

show zoneset active

FC switch A 1# show zoneset active

FC switch B 1# show zoneset active

b. Disable the active zone sets identified in the previous step:

no zoneset activate name zoneset name vsan vsan id

The following example shows two zone sets 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 zone sets are deactivated, clear the zone database:

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

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 "permit":

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
```
```
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 them.



Perform these steps on only one switch in each fabric.

The zoning depends on the model FC-to-SAS bridge you are using. For details, see the section for your model bridge. The examples show Brocade switch ports, so adjust your ports accordingly.

- Zoning for FibreBridge 7500N or 7600N bridges using one FC port
- Zoning for FibreBridge 7500N bridges using both FC ports

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

a. Create the storage zones:

zone name STOR-zone-name vsan vsanid

b. Add storage ports to the zone:

member portswitch WWN

c. Activate the zone set:

```
zoneset activate name STOR-zone-name-setname vsan vsan-id
```

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 a storage zone set and add the storage zones to the new set.



Perform these steps on only one switch in the fabric.

a. Create the storage zone set:

zoneset name STOR-zone-set-name vsan vsan-id

b. Add storage zones to the zone set:

member STOR-zone-name

c. Activate the zone set:

zoneset activate name STOR-zone-set-name vsan vsanid

```
FC_switch_A_1# conf t
FC_switch_A_1(config) # zoneset name STORI_Zoneset_1_20 vsan 20
FC_switch_A_1(config-zoneset) # member STOR_Zone_1_20_25
...
FC_switch_A_1(config-zoneset) # exit
FC_switch_A_1(config) # zoneset activate name STOR_ZoneSet_1_20 vsan
20
FC_switch_A_1(config) # exit
FC_switch_A_1(config) # exit
FC_switch_A_1# copy running-config startup-config
```

6. Create FCVI zones and add the FCVI ports to them.

Each FCVI zone contains the FCVI ports from all the controllers of one DR Group.



Perform these steps on only one switch in the fabric.

The zoning depends on the model FC-to-SAS bridge you are using. For details, see the section for your model bridge. The examples show Brocade switch ports, so adjust your ports accordingly.

- Zoning for FibreBridge 7500N or 7600N bridges using one FC port
- Zoning for FibreBridge 7500N bridges using both FC ports

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

a. Create the FCVI zones:

zone name FCVI-zone-name vsan vsanid

b. Add FCVI ports to the zone:

member FCVI-zone-name

c. Activate the zone set:

zoneset activate name FCVI-zone-name-set-name vsan vsanid

```
FC_switch_A_1# conf t
FC_switch_A_1(config)# zone name FCVI_Zone_1_10_25 vsan 10
FC_switch_A_1(config-zone)# member interface fc1/1
swwn20:00:00:05:9b:24:cb:78
FC_switch_A_1(config-zone)# member interface fc1/2
swwn20:00:00:05:9b:24:cb:78
FC_switch_A_1(config-zone)# member interface fc1/1
swwn20:00:00:05:9b:24:12:99
FC_switch_A_1(config-zone)# member interface fc1/2
swwn20:00:00:05:9b:24:12:99
FC_switch_A_1(config-zone)# end
FC_switch_A_1(config-zone)# end
FC_switch_A_1# copy running-config startup-config
```

7. Create an FCVI zone set and add the FCVI zones to it:



Perform these steps on only one switch in the fabric.

a. Create the FCVI zone set:

```
zoneset name FCVI zone set name vsan vsan-id
```

b. Add FCVI zones to the zone set:

member FCVI zonename

c. Activate the zone set:

zoneset activate name FCVI_zone_set_name vsan vsan-id

```
FC_switch_A_1# conf t
FC_switch_A_1(config)# zoneset name FCVI_Zoneset_1_10 vsan 10
FC_switch_A_1(config-zoneset)# member FCVI_Zone_1_10_25
FC_switch_A_1(config-zoneset)# member FCVI_Zone_1_10_29
...
FC_switch_A_1(config)# exit
FC_switch_A_1(config)# zoneset activate name FCVI_ZoneSet_1_10 vsan 10
FC_switch_A_1(config)# exit
FC_switch_A_1(config)# exit
```

8. Verify the zoning:

show zone

9. 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

Issue the following command on both FC switch fabrics:

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

FC_switch_A_1# copy running-config startup-config

FC_switch_B_1# copy running-config startup-config

Install FC-to-SAS bridges and SAS disk shelves

You install and cable ATTO FibreBridge bridges and SAS disk shelves when adding new storage to the configuration.

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:



In-band management of the FC-to-SAS bridges

Beginning with ONTAP 9.5 with FibreBridge 7500N or 7600N bridges, *in-band management* of the bridges is supported as an alternative to IP management of the bridges. Beginning with ONTAP 9.8, out-of-band management is deprecated.



Beginning with ONTAP 9.8, the storage bridge command is replaced with system bridge. The following steps show the storage bridge command, but if you are running ONTAP 9.8 or later, the system bridge command is preferred.

When using in-band management, the bridges can be managed and monitored from the ONTAP CLI via the FC connection to the bridge. Physical access to the bridge via the bridge Ethernet ports is not required, reducing the security vulnerability of the bridge.

The availability of in-band management of the bridges depends on the version of ONTAP:

- Beginning with ONTAP 9.8, bridges are managed via in-band connections by default and out-of-band management of the bridges via SNMP is deprecated.
- ONTAP 9.5 through 9.7: Either in-band management or out-of-band SNMP management is supported.
- Prior to ONTAP 9.5, only out-of-band SNMP management is supported.

Bridge CLI commands can be issued from the ONTAP interface storage bridge run-cli -name bridge name -command bridge command name command at the ONTAP interface.



Using in-band management with IP access disabled is recommended to improve security by limiting physical connectivity to the bridge.

FibreBridge 7600N and 7500N bridge limits and attachment rules

Review the limits and considerations when attaching FibreBridge 7600N and 7500N bridges.

FibreBridge 7600N and 7500N bridge limits

- The maximum number of HDD and SSD drives combined is 240.
- The maximum number of SSD drives is 96.
- The maximum number of SSDs per SAS port is 48.
- The maximum number of shelves per SAS port is 10.

FibreBridge 7600N and 7500N bridge attachment rules

- Do not mix SSD and HDD drives on the same SAS port.
- Distribute the shelves evenly across the SAS ports.
- You shouldn't have DS460 shelves on the same SAS port as other shelf types (for example, DS212 or DS224 shelves).

Example configuration

The following shows an example configuration for connecting four DS224 shelves with SSD drives and six DS224 shelves with HDD drives:

SAS port	Shelves and drives
SAS port-A	2x DS224 shelves with SSD drives
SAS port-B	2x DS224 shelves with SSD drives
SAS port-C	3x DS224 shelves with HDD drives
SAS port-D	3x DS224 shelves with HDD drives

Prepare for the installation

When you are preparing to install the 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 configuration must be using supported hardware models and software versions.

In the NetApp Interoperability Matrix Tool (IMT), you can use the Storage Solution field to select your MetroCluster solution. You use the **Component Explorer** to select the components and ONTAP version to refine your search. You can click **Show Results** to display the list of supported configurations that match the criteria.

- Each FC switch must have one FC port available for one bridge to connect to it.
- You must have familiarized yourself with how to handle SAS cables and the considerations and best

practices for installing and cabling disk shelves.

The *Installation and Service Guide* for your disk shelf model describes the considerations and best practices.

 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 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 Installation and Service Guide for your disk shelf model:
 - a. Access the ATTO web site using the link provided for your FibreBridge model and download the manual and the QuickNAV utility.



The *ATTO FibreBridge Installation and Operation Manual* for your model bridge has more information about management interfaces.

You can access this and other content on the ATTO web site by using the link provided on the ATTO FibreBridge Description page.

- 2. Gather the hardware and information needed to use the recommended bridge management interfaces, the ATTO ExpressNAV GUI, and the ATTO QuickNAV utility:
 - a. Determine a non-default user name and password (for accessing the bridges).

You should change the default user name and password.

- b. If configuring for IP management of the bridges, you need the shielded Ethernet cable provided with the bridges (which connects from the bridge Ethernet management 1 port to your network).
- c. If configuring for IP management of the bridges, you need 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.

Install the FC-to-SAS bridge and SAS shelves

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

About this task

• The disk and shelf configuration at both sites should be identical.

If a non-mirrored aggregate is used, the disk and shelf configuration at each site might be different.



All disks in the disaster recovery group must use the same type of connection and be visible to all of the nodes within the disaster recovery group, regardless of the disks being used for mirrored or non-mirrored aggregate.

• 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 FibreBridge bridges.

NetApp Hardware Universe

 A mix of IOM12 modules and IOM3 modules is not supported within the same storage stack. A mix of IOM12 modules and IOM6 modules is supported within the same storage stack if your system is running a supported version of ONTAP.

In-band ACP is supported without additional cabling in the following shelves and FibreBridge 7500N or 7600N bridge:

- IOM12 (DS460C) behind a 7500N or 7600N bridge with ONTAP 9.2 and later
- IOM12 (DS212C and DS224C) behind a 7500N or 7600N bridge with ONTAP 9.1 and later

SAS shelves in MetroCluster configurations do not support ACP cabling.

Enable IP port access on the FibreBridge 7600N bridge if necessary

If you are using an ONTAP version prior to 9.5, or otherwise plan to use out-of-band access to the FibreBridge 7600N bridge using telnet or other IP port protocols and services (FTP, ExpressNAV, ICMP, or QuickNAV), you can enable the access services via the console port.

About this task

1

i.

Unlike the ATTO FibreBridge 7500N bridges, the FibreBridge 7600N bridge is shipped with all IP port protocols and services disabled.

Beginning with ONTAP 9.5, *in-band management* of the bridges is supported. This means the bridges can be configured and monitored from the ONTAP CLI via the FC connection to the bridge. Physical access to the bridge via the bridge Ethernet ports is not required and the bridge user interfaces are not required.

Beginning with ONTAP 9.8, *in-band management* of the bridges is supported by default and out-of-band SNMP management is deprecated.

This task is required if you are **not** using in-band management to manage the bridges. In this case, you need to configure the bridge via the Ethernet management port.

Steps

- 1. Access the bridge console interface by connecting a serial cable to the serial port on the FibreBridge 7600N bridge.
- 2. Using the console, enable the access services, and then save the configuration:

set closeport none

saveconfiguration

The set closeport none command enables all access services on the bridge.

3. Disable a service, if desired, by issuing the set closeport command and repeating the command as necessary until all desired services are disabled:

set closeport *service*

The set closeport command disables a single service at a time.

The parameter *service* can be specified as one of the following:

- expressnav
- ∘ ftp
- icmp
- quicknav
- ∘ snmp
- telnet

You can check whether a specific protocol is enabled or disabled by using the get closeport command.

4. If you are enabling SNMP, you must also issue following command:

set SNMP enabled

SNMP is the only protocol that requires a separate enable command.

5. Save the configuration:

saveconfiguration

Configure the FC-to-SAS bridges

Before cabling your model of the FC-to-SAS bridges, you must configure the settings in the FibreBridge software.

Before you begin

You should decide whether you will be using in-band management of the bridges.



Beginning with ONTAP 9.8, the storage bridge command is replaced with system bridge. The following steps show the storage bridge command, but if you are running ONTAP 9.8 or later, the system bridge command is preferred.

About this task

If you will be using in-band management of the bridge rather than IP management, the steps for configuring the Ethernet port and IP settings can be skipped, as noted in the relevant steps.

Steps

1. Configure the serial console port on the ATTO FibreBridge by setting the port speed to 115000 bauds:

```
get serialportbaudrate
SerialPortBaudRate = 115200
Ready.
set serialportbaudrate 115200
Ready. *
saveconfiguration
Restart is necessary....
Do you wish to restart (y/n) ? y
```

2. If configuring for in-band management, connect a cable from FibreBridge RS-232 serial port to the serial (COM) port on a personal computer.

The serial connection will be used for initial configuration, and then in-band management via ONTAP and the FC ports can be used to monitor and manage the bridge.

3. If configuring for IP management, connect the Ethernet management 1 port on each bridge to your network by using an Ethernet cable.

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

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.

4. If configuring for IP management, configure the Ethernet management 1 port for each bridge by following the procedure in section 2.0 of the *ATTO FibreBridge Installation and Operation Manual* for your bridge model.

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

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.

5. Configure the bridge.

You should make note of the user name and password that you designate.



Do not configure time synchronization on ATTO FibreBridge 7600N or 7500N. The time synchronization for ATTO FibreBridge 7600N or 7500N is set to the cluster time after the bridge is discovered by ONTAP. It is also synchronized periodically once a day. The time zone used is GMT and is not changeable.

a. If configuring for IP management, configure the IP settings of the bridge.

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

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

If using the CLI, you must run the following commands:

set ipaddress mpl ip-address
set ipsubnetmask mpl subnet-mask
set ipgateway mpl x.x.x.x
set ipdhcp mpl disabled
set ethernetspeed mpl 1000

b. Configure the bridge name.

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

Example bridge names for one stack group on each site:

- bridge_A_1a
- bridge_A_1b
- bridge_B_1a
- bridge_B_1b

If using the CLI, you must run the following command:

set bridgename bridge name

c. If running ONTAP 9.4 or earlier, enable SNMP on the bridge:

set SNMP enabled

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

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

The supported FC data rate depends on your model bridge.

- The FibreBridge 7600N bridge supports up to 32, 16, or 8 Gbps.
- The FibreBridge 7500N bridge supports up to 16, 8, or 4 Gbps.



The FCDataRate speed you select is limited to the maximum speed supported by both the bridge and the FC port of the controller module to which the bridge port connects. Cabling distances must not exceed the limitations of the SFPs and other hardware.

If using the CLI, you must run the following command:

```
set FCDataRate <port-number> <port-speed>
```

b. If you are configuring a FibreBridge 7500N bridge, configure the connection mode that the port uses to "ptp".



The FCConnMode setting is not required when configuring a FibreBridge 7600N bridge.

If using the CLI, you must run the following command:

```
set FCConnMode <port-number> ptp
```

- c. If you are configuring a FibreBridge 7600N or 7500N bridge, you must configure or disable the FC2 port.
 - If you are using the second port, you must repeat the previous substeps for the FC2 port.
 - If you are not using the second port, then you must disable the port:

```
FCPortDisable <port-number>
```

The following example shows the disabling of FC port 2:

```
FCPortDisable 2
Fibre Channel Port 2 has been disabled.
```

d. If you are configuring a FibreBridge 7600N or 7500N bridge, disable the unused SAS ports:

```
SASPortDisable sas-port
```



SAS ports A through D are enabled by default. You must disable the SAS ports that are not being used.

If only SAS port A is used, then SAS ports B, C, and D must be disabled. The following example shows the disabling of SAS port B. You must similarly disable SAS ports C and D:

```
SASPortDisable b
SAS Port B has been disabled.
```

7. Secure access to the bridge and save the bridge's configuration. Choose an option from below depending on the version of ONTAP your system is running.

ONTAP version	Steps
ONTAP 9.5 or later	 a. View the status of the bridges: storage bridge show The output shows which bridge is not secured. b. Secure the bridge: securebridge
ONTAP 9.4 or earlier	 a. View the status of the bridges: storage bridge show The output shows which bridge is not secured. b. Check the status of the unsecured bridge's ports: info The output shows the status of Ethernet ports MP1 and MP2. c. If Ethernet port MP1 is enabled, run: set EthernetPort mp1 disabled If Ethernet port MP2 is also enabled, repeat the previous substep for port MP2. d. Save the bridge's configuration. You must run the following commands: SaveConfiguration FirmwareRestart You are prompted to restart the bridge.

8. After completing MetroCluster configuration, use the flashimages command to check your version of FibreBridge firmware and, if the bridges are not using the latest supported version, update the firmware on all bridges in the configuration.

Maintain MetroCluster Components

Cable disk shelves to the bridges

You must use the correct FC-to-SAS bridges for cabling your disk shelves.

Choices

- Cable a FibreBridge 7600N or 7500N bridge with disk shelves using IOM12 modules
- Cable a FibreBridge 7600N or 7500N bridge with disk shelves using IOM6 or IOM3 modules

Cable a FibreBridge 7600N or 7500N bridge with disk shelves using IOM12 modules

After configuring the bridge, you can start cabling your new system.

About this task

For disk shelves, you insert a SAS cable connector with the pull tab oriented down (on the underside of the connector).

Steps

- 1. Daisy-chain the disk shelves in each stack:
 - a. Beginning with the logical first shelf in the stack, connect IOM A port 3 to IOM A port 1 on the next shelf until each IOM A in the stack is connected.
 - b. Repeat the previous substep for IOM B.
 - c. Repeat the previous substeps for each stack.

The *Installation and Service Guide* for your disk shelf model provides detailed information about daisychaining disk shelves.

- 2. Power on the disk shelves, and then set the shelf IDs.
 - · You must power-cycle each disk shelf.
 - Shelf IDs must be unique for each SAS disk shelf within each MetroCluster DR group (including both sites).
- 3. Cable disk shelves to the FibreBridge bridges.
 - a. For the first stack of disk shelves, cable IOM A of the first shelf to SAS port A on FibreBridge A, and cable IOM B of the last shelf to SAS port A on FibreBridge B.
 - b. For additional shelf stacks, repeat the previous step using the next available SAS port on the FibreBridge bridges, using port B for the second stack, port C for the third stack, and port D for the fourth stack.
 - c. During cabling, attach the stacks based on IOM12 and IOM3/IOM6 modules to the same bridge as long as they are connected to separate SAS ports.



Each stack can use different models of IOM, but all disk shelves within a stack must use the same model.

The following illustration shows disk shelves connected to a pair of FibreBridge 7600N or 7500N bridges:



Cable a FibreBridge 7600N or 7500N bridge with shelves using IOM6 or IOM3 modules

After configuring the bridge, you can start cabling your new system. The FibreBridge 7600N or 7500N bridge uses mini-SAS connectors and supports shelves that use IOM6 or IOM3 modules.

About this task

IOM3 modules are not supported with FibreBridge 7600N bridges.

For disk shelves, you insert a SAS cable connector with the pull tab oriented down (on the underside of the connector).

Steps

- 1. Daisy-chain the shelves in each stack.
 - a. For the first stack of shelves, cable IOM A square port of the first shelf to SAS port A on FibreBridge A.
 - b. For the first stack of shelves, cable IOM B circle port of the last shelf to SAS port A on FibreBridge B.

The *Installation and Service Guide* for your shelf model provides detailed information about daisy-chaining shelves.

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

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



2. For additional shelf stacks, repeat the previous steps using the next available SAS port on the FibreBridge bridges, using port B for a second stack, port C for a third stack, and port D for a fourth stack.





Verify bridge connectivity and cabling the bridge FC ports

You should verify that each bridge can detect all of the disk drives, and then cable each bridge to the local FC switches.

Steps

1. Verify that each bridge can detect all of the disk drives and disk shelves it is connected to:

If you are using the…	Then
-----------------------	------

ATTO ExpressNAV GUI	a. In a supported web browser, enter the IP address of a bridge in the browser box.		
	You are brought to the ATTO FibreBridge homepage of the bridge for which you entered the IP address, which has a link.		
	b. Click the link, and then enter your user name and the password that you designated when you configured the bridge.		
	The ATTO FibreBridge status page of the bridge appears with a menu to the left.		
	c. Click Advanced.		
	d. View the connected devices by using the sastargets command, and then click Submit .		
Serial port connection	View the connected devices:		
	sastargets		

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

Tgt	VendorID	ProductID	Туре	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	3QP1G7A900009939FNTT
8	NETAPP	X410_S15K6288A15	DISK	3QP1FY0T00009940G8PA
9	NETAPP	X410_S15K6288A15	DISK	3QP1FXW600009940VERQ



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 of the output.

2. Verify that the command output shows that 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 1 for each remaining bridge.

Not correct	a. Check for loose SAS cables or correct the SAS cabling by repeating the cabling.
	Cable disk shelves to the bridges
	b. Repeat Step 1.

3. Cable each bridge to the local FC switches, using the cabling in the table for your configuration and switch model and FC-to-SAS bridge model:



The second FC port connection on the FibreBridge 7500N bridge should not be cabled until zoning has been completed.

See the port assignments for your version of ONTAP.

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

Related information

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

Port assignments for FC switches

Secure or unsecure the FibreBridge bridge

To easily disable potentially unsecure Ethernet protocols on a bridge, beginning with ONTAP 9.5 you can secure the bridge. This disables the bridge's Ethernet ports. You can also reenable Ethernet access.

About this task

- Securing the bridge disables telnet and other IP port protocols and services (FTP, ExpressNAV, ICMP, or QuickNAV) on the bridge.
- This procedure uses out-of-band management using the ONTAP prompt, which is available beginning with ONTAP 9.5.

You can issue the commands from the bridge CLI if you are not using out-of-band management.

- The unsecurebridge command can be used to re-enable the Ethernet ports.
- In ONTAP 9.7 and earlier, running the securebridge command on the ATTO FibreBridge might not update the bridge status correctly on the partner cluster. If this occurs, run the securebridge command from the partner cluster.



Beginning with ONTAP 9.8, the storage bridge command is replaced with system bridge. The following steps show the storage bridge command, but if you are running ONTAP 9.8 or later, the system bridge command is preferred.

Steps

- 1. From the ONTAP prompt of the cluster containing the bridge, secure or unsecure the bridge.
 - The following command secures bridge_A_1:

cluster A> storage bridge run-cli -bridge bridge A 1 -command securebridge

• The following command unsecures bridge_A_1:

cluster_A> storage bridge run-cli -bridge bridge_A_1 -command unsecurebridge

2. From the ONTAP prompt of the cluster containing the bridge, save the bridge configuration:

```
storage bridge run-cli -bridge bridge-name -command saveconfiguration
```

The following command secures bridge_A_1:

```
cluster_A> storage bridge run-cli -bridge bridge_A_1 -command
saveconfiguration
```

3. From the ONTAP prompt of the cluster containing the bridge, restart the bridge's firmware:

```
storage bridge run-cli -bridge bridge-name -command firmwarerestart
```

The following command secures bridge_A_1:

```
cluster_A> storage bridge run-cli -bridge bridge_A_1 -command firmwarerestart
```

Configuring the MetroCluster software in ONTAP

You must set up each node in the MetroCluster configuration in ONTAP, including the node-level configurations and the configuration of the nodes into two sites. You must also implement the MetroCluster relationship between the two sites.



Gathering required information

You need to gather the required IP addresses for the controller modules before you begin the configuration process.

IP network information worksheet 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 (switched clusters)

When you cable the system, you need a host name and management IP address for each cluster switch. This information is not needed if you are using a two-node switchless cluster or have a two-node MetroCluster configuration (one node at each site).

Cluster switch	Host name	IP address	Network mask	Default gateway
Interconnect 1				
Interconnect 2				
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 2		
Not required if using two-node MetroCluster configuration (one node at each site).		
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 addresses of two intercluster LIFs, including 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 IC LIF 1				
Node 1 IC LIF 2				
Node 2 IC LIF 1 Not required for two- node MetroCluster configurations (one node at each site).				
Node 2 IC LIF 2 Not required for two- node MetroCluster configurations (one node at each site).				

Site A time server information

You must synchronize the time, which requires 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 requires 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 SP information

You must enable access to the Service Processor (SP) 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 Not required for two-node MetroCluster configurations (one node at each site).			

IP network information worksheet 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 switch information (switched clusters)

When you cable the system, you need a host name and management IP address for each cluster switch. This information is not needed if you are using a two-node switchless cluster or you have a two-node MetroCluster configuration (one node at each site).

Interconnect 1		
Interconnect 2		
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	
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				
Not required for two- node MetroCluster configurations (one node at each site).				
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 addresses of two intercluster LIFs, including 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 IC LIF 1				
Node 1 IC LIF 2				
Node 2 IC LIF 1 Not required for two- node MetroCluster				
node at each site).				
Node 2 IC LIF 2 Not required for two- node MetroCluster configurations (one node at each site).				

Site B time server information

You must synchronize the time, which requires 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 requires 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 SP information

You must enable access to the Service Processor (SP) 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)			
Not required for two-node MetroCluster configurations (one node at each site).			

Similarities and differences between standard 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 of the MetroCluster components must be cabled and configured. However, 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 the root aggregate.	Same in both types of clusters		
Configure nodes in the cluster as HA pairs	Same in both types of clusters		
Set up the cluster on one node in the cluster.	Same in both types of clusters		
Join the other node to the cluster.	Same in both types of clusters		
Create a mirrored root aggregate.	Optional	Required	
Peer the clusters.	Optional	Required	

Verifying and configuring the HA state of components in Maintenance mode

When configuring a storage system in a MetroCluster FC configuration, you must make sure that the highavailability (HA) state of the controller module and chassis components is mcc or mcc-2n so that these components boot properly. Although this value should be preconfigured on systems received from the factory, you should still verify the setting before proceeding.

If the HA state of the controller module and chassis is incorrect, you cannot configure the MetroCluster without re-initializing the node. You must correct the setting using this procedure, and then initialize the system by using one of the following procedures:



- In a MetroCluster IP configuration, follow the steps in Restore system defaults on a controller module.
- In a MetroCluster FC configuration, follow the steps in Restore system defaults and configuring the HBA type on a controller module.

Before you begin

Verify that the system is in Maintenance mode.

Steps

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

ha-config show

The correct HA state depends on your MetroCluster configuration.

MetroCluster configuration type	HA state for all components
Eight or four node MetroCluster FC configuration	mcc
Two-node MetroCluster FC configuration	mcc-2n
Eight or four node MetroCluster IP configuration	mccip

2. If the displayed system state of the controller is not correct, set the correct HA state for your configuration on the controller module:

MetroCluster configuration type	Command
Eight or four node MetroCluster FC configuration	ha-config modify controller mcc
Two-node MetroCluster FC configuration	ha-config modify controller mcc-2n
Eight or four node MetroCluster IP configuration	ha-config modify controller mccip

3. If the displayed system state of the chassis is not correct, set the correct HA state for your configuration on the chassis:

MetroCluster configuration type	Command
Eight or four node MetroCluster FC configuration	ha-config modify chassis mcc
Two-node MetroCluster FC configuration	ha-config modify chassis mcc-2n
Eight or four node MetroCluster IP configuration	ha-config modify chassis mccip

4. Boot the node to ONTAP:

boot_ontap

5. Repeat this entire procedure to verify the HA state on each node in the MetroCluster configuration.

Restoring system defaults and configuring the HBA type on a controller module

About this task

To ensure a successful MetroCluster installation, reset and restore defaults on the controller modules.

Important

This task is only required for stretch configurations using FC-to-SAS bridges.

Steps

1. At the LOADER prompt, return the environmental variables to their default setting:

set-defaults

- 2. Boot the node into Maintenance mode, then configure the settings for any HBAs in the system:
 - a. Boot into Maintenance mode:

boot_ontap maint

b. Check the current settings of the ports:

ucadmin show

c. Update the port settings as needed.

If you have this type of HBA and desired mode	Use this command
CNA FC	ucadmin modify -m fc -t initiator adapter_name
CNA Ethernet	ucadmin modify -mode cna <pre>adapter_name</pre>
FC target	fcadmin config -t target <pre>adapter_name</pre>

3. Exit Maintenance mode:

halt

After you run the command, wait until the node stops at the LOADER prompt.

4. Boot the node back into Maintenance mode to enable the configuration changes to take effect:

boot_ontap maint

5. Verify the changes you made:

If you have this type of HBA	Use this command
CNA	ucadmin show
FC	fcadmin show

6. Exit Maintenance mode:

halt

After you run the command, wait until the node stops at the LOADER prompt.

7. Boot the node to the boot menu:

boot_ontap menu

After you run the command, wait until the boot menu is shown.

8. Clear the node configuration by typing "wipeconfig" at the boot menu prompt, and then press Enter.

The following screen shows the boot menu prompt:

```
Please choose one of the following:
     (1) Normal Boot.
     (2) Boot without /etc/rc.
     (3) Change password.
     (4) Clean configuration and initialize all disks.
     (5) Maintenance mode boot.
     (6) Update flash from backup config.
     (7) Install new software first.
     (8) Reboot node.
     (9) Configure Advanced Drive Partitioning.
     Selection (1-9)? wipeconfig
 This option deletes critical system configuration, including cluster
membership.
 Warning: do not run this option on a HA node that has been taken over.
 Are you sure you want to continue?: yes
 Rebooting to finish wipeconfig request.
```

Configuring FC-VI ports on a X1132A-R6 quad-port card on FAS8020 systems

If you are using the X1132A-R6 quad-port card on a FAS8020 system, you can enter Maintenance mode to configure the 1a and 1b ports for FC-VI and initiator usage. This is not required on MetroCluster systems received from the factory, in which the ports are set appropriately for your configuration.

About this task

This task must be performed in Maintenance mode.



Converting an FC port to an FC-VI port with the ucadmin command is only supported on the FAS8020 and AFF 8020 systems. Converting FC ports to FCVI ports is not supported on any other platform.

Steps

1. Disable the ports:

```
storage disable adapter 1a storage disable adapter 1b
```

*> storage disable adapter 1a
Jun 03 02:17:57 [controller_B_1:fci.adapter.offlining:info]: Offlining
Fibre Channel adapter 1a.
Host adapter 1a disable succeeded
Jun 03 02:17:57 [controller_B_1:fci.adapter.offline:info]: Fibre Channel
adapter 1a is now offline.
*> storage disable adapter 1b
Jun 03 02:18:43 [controller_B_1:fci.adapter.offlining:info]: Offlining
Fibre Channel adapter 1b.
Host adapter 1b disable succeeded
Jun 03 02:18:43 [controller_B_1:fci.adapter.offline:info]: Fibre Channel
adapter 1b disable succeeded

2. Verify that the ports are disabled:

ucadmin show

*> ucadmin show					
	Current	Current	Pending	Pending	Admin
Adapter	Mode	Туре	Mode	Туре	Status
•••					
1a	fc	initiator	-	-	offline
1b	fc	initiator	-	-	offline
1c	fc	initiator	-	-	online
1d	fc	initiator	-	-	online

3. Set the a and b ports to FC-VI mode:

ucadmin modify -adapter 1a -type fcvi

The command sets the mode on both ports in the port pair, 1a and 1b (even though only 1a is specified in the command).

```
*> ucadmin modify -t fcvi 1a
Jun 03 02:19:13 [controller_B_1:ucm.type.changed:info]: FC-4 type has
changed to fcvi on adapter 1a. Reboot the controller for the changes to
take effect.
Jun 03 02:19:13 [controller_B_1:ucm.type.changed:info]: FC-4 type has
changed to fcvi on adapter 1b. Reboot the controller for the changes to
take effect.
```

4. Confirm that the change is pending:

*> ucadm	in show				
	Current	Current	Pending	Pending	Admin
Adapter	Mode	Туре	Mode	Туре	Status
• • •					
1a	fc	initiator	-	fcvi	offline
1b	fc	initiator	-	fcvi	offline
1c	fc	initiator	-	-	online
1d	fc	initiator	-	-	online

- 5. Shut down the controller, and then reboot into Maintenance mode.
- 6. Confirm the configuration change:

ucadmin show local

Node	Adapter	Mode	Туре	Mode	Туре	Status
controller_B_1						
	1a	fc	fcvi	-	-	online
controller B 1						
	1b	fc	fcvi	_	-	online
controller B 1						
	1c	fc	initiator	_	_	online
controller B 1						
	1d	fc	initiator	_	_	online
6 entries were	displaye	ed.				

Verifying disk assignment in Maintenance mode in an eight-node or a four-node configuration

Before fully booting the system to 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 pool has 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 shelves at Site A

Disk shelf (sample_shelf_name)	Belongs to…	And is assigned to that node's
Disk shelf 1 (shelf_A_1_1)	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 1 (shelf_A_3_1)	Node A 3	Pool 0
Disk shelf 2 (shelf_A_3_3)		
Disk shelf 3 (shelf_B_3_1)	Node B 3	Pool 1
Disk shelf 4 (shelf_B_3_3)		
Disk shelf 5 (shelf_A_4_1)	Node A 4	Pool 0
Disk shelf 6 (shelf_A_4_3)		
Disk shelf 7 (shelf_B_4_1)	Node B 4	Pool 1
Disk shelf 8 (shelf_B_4_3)		

Disk shelves at Site B

Disk shelf (sample_shelf_name)	Belongs to…	And is assigned to that node's
Disk shelf 9 (shelf_B_1_2)	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)		

Disk shelf 1 (shelf_B_3_2)	Node A 3	Pool 0
Disk shelf 2 (shelf_B_3_4)		
Disk shelf 3 (shelf_A_3_2)	Node B 3	Pool 1
Disk shelf 4 (shelf_A_3_4)		
Disk shelf 5 (shelf_B_4_2)	Node A 4	Pool 0
Disk shelf 6 (shelf_B_4_4)		
Disk shelf 7 (shelf_A_4_2)	Node B 4	Pool 1
Disk shelf 8 (shelf_A_4_4)		

Steps

1. Confirm the shelf assignments:

disk show -v

2. If necessary, explicitly assign disks on the attached disk shelves to the appropriate pool:

disk assign

Using wildcards in the command enables you to assign all of the disks on a disk shelf with one command. You can identify the disk shelf IDs and bays for each disk with the storage show disk -x command.

Assigning disk ownership in non-AFF systems

If the MetroCluster nodes do not have the disks correctly assigned, or if you are using DS460C disk shelves in your configuration, you must assign disks to each of the nodes in the MetroCluster configuration on a shelf-by-shelf basis. You will create a configuration in which each node has the same number of disks in its local and remote disk pools.

Before you begin

The storage controllers must be in Maintenance mode.

About this task

If your configuration does not include DS460C disk shelves, this task is not required if disks were correctly assigned when received from the factory.



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.

If your configuration includes DS460C disk shelves, you should manually assign the disks using the following guidelines for each 12-disk drawer:

Assign these disks in the drawer	To this node and pool
0 - 2	Local node's pool 0
3 - 5	HA partner node's pool 0
6 - 8	DR partner of the local node's pool 1
9 - 11	DR partner of the HA partner's pool 1

This disk assignment pattern ensures that an aggregate is minimally affected in case a drawer goes offline.

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

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

If storage controller Controller_A_2 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-nameshelf-name -p pool

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

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.

Assigning disk ownership in AFF systems

If you are using AFF systems in a configuration with mirrored aggregates and the nodes do not have the disks (SSDs) correctly assigned, you should assign half the disks on each shelf to one local node and the other half of the disks to its HA partner node. You should create a configuration in which each node has the same number of disks in its local and remote disk pools.

Before you begin

The storage controllers must be in Maintenance mode.

About this task

This does not apply to configurations which have unmirrored aggregates, an active/passive configuration, or that have an unequal number of disks in local and remote pools.

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



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 disks to the nodes located at the first site (site A):

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

a. On the first node, systematically assign half the disks on each shelf to pool 0 and the other half to the HA partner's pool 0:

disk assign -shelf <shelf-name> -p <pool> -n <number-of-disks>

If storage controller Controller_A_1 has four shelves, each with 8 SSDs, you issue the following commands:

```
*> disk assign -shelf FC_switch_A_1:1-4.shelf1 -p 0 -n 4
*> disk assign -shelf FC_switch_A_1:1-4.shelf2 -p 0 -n 4
*> disk assign -shelf FC_switch_B_1:1-4.shelf1 -p 1 -n 4
*> disk assign -shelf FC switch B 1:1-4.shelf2 -p 1 -n 4
```

b. Repeat the process for the second node at the local site, systematically assigning half the disks on each shelf to pool 1 and the other half to the HA partner's pool 1:

disk assign -disk disk-name -p pool

If storage controller Controller_A_1 has four shelves, each with 8 SSDs, you issue the following commands:

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

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

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

a. On the first node at the remote site, systematically assign half the disks on each shelf to pool 0 and the other half to the HA partner's pool 0:

disk assign -disk disk-name -p pool

If storage controller Controller_B_1 has four shelves, each with 8 SSDs, you issue the following commands:

*> disk assign -shelf FC_switch_B_1:1-5.shelf1 -p 0 -n 4
*> disk assign -shelf FC_switch_B_1:1-5.shelf2 -p 0 -n 4
*> disk assign -shelf FC_switch_A_1:1-5.shelf1 -p 1 -n 4
*> disk assign -shelf FC_switch_A_1:1-5.shelf2 -p 1 -n 4

b. Repeat the process for the second node at the remote site, systematically assigning half the disks on each shelf to pool 1 and the other half to the HA partner's pool 1:

disk assign -disk disk-name -p pool

If storage controller Controller_B_2 has four shelves, each with 8 SSDs, you issue the following commands:

```
*> disk assign -shelf FC_switch_B_1:1-5.shelf3 -p 0 -n 4
*> disk assign -shelf FC_switch_B_1:1-5.shelf4 -p 0 -n 4
*> disk assign -shelf FC_switch_A_1:1-5.shelf3 -p 1 -n 4
*> disk assign -shelf FC_switch_A_1:1-5.shelf4 -p 1 -n 4
```

4. Confirm the disk assignments:

storage show disk

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 disk assignment in Maintenance mode in a two-node configuration

Before fully booting the system to ONTAP, you can optionally boot the system to Maintenance mode and verify the disk assignment on the nodes. The disks should be assigned to create a fully symmetric configuration with both sites owning their own disk shelves and serving data, where each node and each pool have an equal number of mirrored disks assigned to them.

Before you begin

The system must be in Maintenance mode.

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 (example name)	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 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)			

If your configuration includes DS460C disk shelves, you should manually assign the disks using the following guidelines for each 12-disk drawer:

Assign these disks in the drawer	To this node and pool
1 - 6	Local node's pool 0
7 - 12	DR partner's pool 1

This disk assignment pattern minimizes the effect on an aggregate if a drawer goes offline.

Steps

1. If your system was received from the factory, confirm the shelf assignments:

disk show -v

If necessary, you can explicitly assign disks on the attached disk shelves to the appropriate pool by using the disk assign command.

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. If you have not done so, boot each system into Maintenance mode.
- b. On the node on site A, systematically assign the local disk shelves to pool 0 and the remote disk shelves to pool 1:

```
disk assign -shelf disk_shelf_name -p pool
```

If storage controller node_A_1 has four shelves, you issue the following commands:

```
*> disk assign -shelf shelf_A_1_1 -p 0
*> disk assign -shelf shelf_A_1_3 -p 0
*> disk assign -shelf shelf_A_1_2 -p 1
*> disk assign -shelf shelf_A_1_4 -p 1
```

c. On the node at the remote site (site B), systematically assign its local disk shelves to pool 0 and its remote disk shelves to pool 1:

disk assign -shelf disk shelf name -p pool

If storage controller node_B_1 has four shelves, you issue the following commands:

```
*> disk assign -shelf shelf_B_1_2 -p 0
*> disk assign -shelf shelf_B_1_4 -p 0
*> disk assign -shelf shelf_B_1_1 -p 1
*> disk assign -shelf shelf_B_1_3 -p 1
```

d. Show the disk shelf IDs and bays for each disk:

disk show -v

Setting up ONTAP

You must set up ONTAP on each controller module.

If you need to netboot the new controllers, see Netbooting the new controller modules in the *MetroCluster Upgrade, Transition, and Expansion Guide*.

Choices

Setting up ONTAP in a two-node MetroCluster configuration

• Setting up ONTAP in an eight-mode or four-node MetroCluster configuration

Setting up ONTAP in a two-node MetroCluster configuration

In a two-node MetroCluster configuration, on each cluster you must boot up the node, exit the Cluster Setup wizard, and use the cluster setup command to configure the node into a single-node cluster.

Before you begin

You must not have configured the Service Processor.

About this task

This task is for two-node MetroCluster configurations using native NetApp storage.

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

For more general information about setting up ONTAP, see Set up ONTAP.

Steps

1. Power on the first node.



You must repeat this step on the node at the disaster recovery (DR) site.

The node boots, and then the Cluster Setup wizard starts on the console, informing you that AutoSupport will be enabled automatically.

```
::> Welcome to the cluster setup wizard.
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 cluster setup wizard.
     Any changes you made before quitting will be saved.
You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.
This system will send event messages and periodic reports to NetApp
Technical
Support. To disable this feature, enter
autosupport modify -support disable
within 24 hours.
Enabling AutoSupport can significantly speed problem determination and
resolution, should a problem occur on your system.
For further information on AutoSupport, see:
http://support.netapp.com/autosupport/
Type yes to confirm and continue {yes}: yes
Enter the node management interface port [eOM]:
Enter the node management interface IP address [10.101.01.01]:
Enter the node management interface netmask [101.010.101.0]:
Enter the node management interface default gateway [10.101.01.0]:
Do you want to create a new cluster or join an existing cluster?
{create, join}:
```

2. Create a new cluster:

create

3. Choose whether the node is to be used as a single node cluster.

Do you intend for this node to be used as a single node cluster? {yes, no} [yes]:

4. Accept the system default yes by pressing Enter, or enter your own values by typing no, and then pressing

Enter.

5. Follow the prompts to complete the **Cluster Setup** wizard, pressing Enter to accept the default values or typing your own values and then pressing Enter.

The default values are determined automatically based on your platform and network configuration.

6. After you complete the **Cluster Setup** wizard and it exits, verify that the cluster is active and the first node is healthy: `

cluster show

The following example shows a cluster in which the first node (cluster1-01) is healthy and eligible to participate:

```
cluster1::> cluster show
Node Health Eligibility
cluster1-01 true true
```

If it becomes necessary to change any of the settings you entered for the admin SVM or node SVM, you can access the Cluster Setup wizard by using the cluster setup command.

Setting up ONTAP in an eight-node or four-node MetroCluster configuration

After you boot 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 ONTAP CLI to create aggregates and create the MetroCluster configuration.

Before you begin

You must have cabled the MetroCluster configuration.

About this task

This task is for eight-node or four-node MetroCluster configurations using native NetApp storage.

New MetroCluster systems are preconfigured; you do not need to perform these steps. However, you should configure the AutoSupport tool.

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 instead.

Steps

1. If you have not already done so, power up each node and let them boot completely.

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

The output should be similar to the following:

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

The prompts are similar to the following:

```
Enter the node management interface port: [eOM]:
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 -fields mode

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

storage failover modify -mode ha -node localhost

This command configures high availability mode but does not enable storage failover. Storage failover is automatically enabled when the MetroCluster configuration is performed later in the configuration process.

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

network port show

The following example shows output for cluster_A:

cluster_A::> network port show						
						Speed
(Mbps)						
Node	Port	IPspace	Broadcast Domain	Link M	FU .	Admin/Oper
node_A					1 = 0	
	**e0a	Cluster	Cluster	up	150	0
auto/10	000					
	e0b	Cluster	Cluster	up	1500	
auto/10	**000					
	eOc	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000
	e0e	Default	Default	up	1500	auto/1000
	eOf	Default	Default	up	1500	auto/1000
	eOg	Default	Default	up	1500	auto/1000
node_A	_2					
	**e0a	Cluster	Cluster	up	150	0
auto/10	000					
	e0b	Cluster	Cluster	up	1500	
auto/10	**000					
	eOc	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000
	e0e	Default	Default	up	1500	auto/1000
	eOf	Default	Default	up	1500	auto/1000
	eOg	Default	Default	up	1500	auto/1000
14 ent:	ries were d	displayed.				

- 6. If you are creating a two-node switchless cluster (a cluster without cluster interconnect switches), enable the switchless-cluster networking mode:
 - a. Change to the advanced privilege level:

set -privilege advanced

You can respond $_{\rm Y}$ 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 boot.

8. Use the System Setup tool to configure each node and create the cluster, but do not create aggregates.



You create mirrored aggregates in later tasks.

After you finish

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

Configuring the clusters into a MetroCluster configuration

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

About this task

Before you run metrocluster configure, HA mode and DR mirroring are not enabled and you might see an error message related to this expected behavior. You enable HA mode and DR mirroring later when you run the command metrocluster configure to implement the configuration.

Peering the clusters

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

Configuring intercluster LIFs

You must create intercluster LIFs on ports used for communication between the MetroCluster partner clusters. You can use dedicated ports or ports that also have data traffic.

Choices

- · Configuring intercluster LIFs on dedicated ports
- · Configuring intercluster LIFs on shared data ports

Configuring intercluster LIFs on dedicated ports

You can configure intercluster LIFs on dedicated ports. Doing so typically increases the available bandwidth for replication traffic.

Steps

1. List the ports in the cluster:

network port show

For complete command syntax, see the man page.

The following example shows the network ports in "cluster01":

cluster01::> network port show						
						Speed
(Mbps)						
Node	Port	IPspace	Broadcast Domain	Link	MTU	Admin/Oper
cluste	r01-01					
	e0a	Cluster	Cluster	up	1500	auto/1000
	e0b	Cluster	Cluster	up	1500	auto/1000
	eOc	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000
	e0e	Default	Default	up	1500	auto/1000
	eOf	Default	Default	up	1500	auto/1000
cluste	r01-02					
	e0a	Cluster	Cluster	up	1500	auto/1000
	e0b	Cluster	Cluster	up	1500	auto/1000
	eOc	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000
	e0e	Default	Default	up	1500	auto/1000
	eOf	Default	Default	up	1500	auto/1000

2. Determine which ports are available to dedicate to intercluster communication:

network interface show -fields home-port, curr-port

For complete command syntax, see the man page.

The following example shows that ports "e0e" and "e0f" have not been assigned LIFs:

```
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 eOa
                                e0a
Cluster cluster01-02 clus2 e0b
                                 e0b
cluster01
      cluster mgmt
                       e0c
                                 e0c
cluster01
      cluster01-01_mgmt1 e0c
                                 e0c
cluster01
                        e0c
      cluster01-02 mgmt1
                                 e0c
```

3. Create a failover group for the dedicated ports:

network interface failover-groups create -vserver system_SVM -failover-group
failover_group -targets physical_or_logical_ports

The following example assigns ports "e0e" and "e0f" to the failover group intercluster01 on the system "SVMcluster01":

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

4. Verify that the failover group was created:

network interface failover-groups show

For complete command syntax, see the man page.

```
cluster01::> network interface failover-groups show
                              Failover
Vserver
              Group
                              Targets
-----
 _____
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 intercluster LIFs on the system SVM and assign them to the failover group.

ONTAP 9.6 and later

```
network interface create -vserver system_SVM -lif LIF_name -service-policy
default-intercluster -home-node node -home-port port -address port_IP
-netmask netmask -failover-group failover_group
```

ONTAP 9.5 and earlier

```
network interface create -vserver system_SVM -lif LIF_name -role
intercluster -home-node node -home-port port -address port_IP -netmask
netmask -failover-group failover_group
```

For complete command syntax, see the man page.

The following example creates intercluster LIFs "cluster01_icl01" and "cluster01_icl02" in the failover group "intercluster01":

```
cluster01::> network interface create -vserver cluster01 -lif
cluster01_icl01 -service-
policy default-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 -service-
policy default-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:

ONTAP 9.6 and later

Run the command: network interface show -service-policy default-intercluster

ONTAP 9.5 and earlier

Run the command: network interface show -role intercluster

For complete command syntax, see the man page.

cluster01::> network interface show -service-policy default-intercluster Logical Status Network Current Current Is Interface Admin/Oper Address/Mask Node Port Vserver Home _____ _ _____ _____ ___ cluster01 cluster01 icl01 up/up 192.168.1.201/24 cluster01-01 e0e true cluster01 icl02 192.168.1.202/24 cluster01-02 eOf up/up true

7. Verify that the intercluster LIFs are redundant:

ONTAP 9.6 and later

Run the command: network interface show -service-policy default-intercluster -failover

ONTAP 9.5 and earlier

```
Run the command: network interface show -role intercluster -failover
```

For complete command syntax, see the man page.

The following example shows that the intercluster LIFs "cluster01_icl01" and "cluster01_icl02" on the SVM "e0e" port will fail over to the "e0f" port.

<pre>cluster01::> network interface show -service-policy default-intercluster -failover</pre>						
	Logical	Home	Failover	Failover		
Vserver	Interface	Node:Port	Policy	Group		
cluster0	1					
	cluster01_icl01	cluster01-01:e0e lo	cal-only			
interclu	ster01					
		Failover Targets:	cluster01-01:e0	e,		
			cluster01-01:e0	f		
cluster01_icl02 cluster01-02:e0e local-only						
interclu	ster01					
		Failover Targets:	cluster01-02:e0	e,		
			cluster01-02:e0	f		

Related information

Considerations when using dedicated ports

When determining whether using a dedicated port for intercluster replication is the correct intercluster network solution, you should consider configurations and requirements such as LAN type, available WAN bandwith, replication interval, change rate, and number of ports.

Configuring intercluster LIFs on shared data ports

You can configure intercluster LIFs on ports shared with the data network. Doing so reduces the number of ports you need for intercluster networking.

Steps

1. List the ports in the cluster:

network port show

For complete command syntax, see the man page.

The following example shows the network ports in cluster01:

cluster01::> network port show						
						Speed
(Mbps)						
Node	Port	IPspace	Broadcast Domain	Link	MTU	Admin/Oper
cluste	r01-01					
	e0a	Cluster	Cluster	up	1500	auto/1000
	e0b	Cluster	Cluster	up	1500	auto/1000
	eOc	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000
cluste	r01-02					
	e0a	Cluster	Cluster	up	1500	auto/1000
	e0b	Cluster	Cluster	up	1500	auto/1000
	eOc	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000

2. Create intercluster LIFs on the system SVM:

ONTAP 9.6 and later

Run the command: network interface create -vserver system_SVM -lif LIF_name -service-policy default-intercluster -home-node node -home-port port -address port IP -netmask netmask

ONTAP 9.5 and earlier

Run the command: network interface create -vserver system_SVM -lif LIF_name -role intercluster -home-node node -home-port port -address port_IP -netmask netmask

For complete command syntax, see the man page. The following example creates intercluster LIFs cluster01_icl01 and cluster01_icl02:

```
cluster01::> network interface create -vserver cluster01 -lif
cluster01_icl01 -service-
policy default-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 -service-
policy default-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:

ONTAP 9.6 and later

Run the command: network interface show -service-policy default-intercluster

ONTAP 9.5 and earlier

Run the command: network interface show -role intercluster

For complete command syntax, see the man page.

cluster01::> network interface show -service-policy default-intercluster Logical Status Network Current Current Is Interface Admin/Oper Address/Mask Node Vserver Port 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 redundant:

ONTAP 9.6 and later

Run the command: network interface show -service-policy default-intercluster -failover

ONTAP 9.5 and earlier

```
Run the command: network interface show -role intercluster -failover
```

For complete command syntax, see the man page.

The following example shows that the intercluster LIFs "cluster01_icl01" and "cluster01_icl02" on the "e0c" port will fail over to the "e0d" port.

<pre>cluster01::> network interface show -service-policy default-intercluster -failover</pre>					
	Logical	Home	Failover	Failover	
Vserver	Interface	Node:Port	Policy	Group	
cluster0	1				
	cluster01_icl01	cluster01-01:e0c l	ocal-only		
192.168.	1.201/24				
		Failover Targets:	cluster01-01:e0c	,	
			cluster01-01:e0d		
	cluster01_icl02	cluster01-02:e0c l	ocal-only		
192.168.	1.201/24				
		Failover Targets:	cluster01-02:e0c	,	
			cluster01-02:e0d		

Related information

Considerations when sharing data ports

Creating a cluster peer relationship

You must create the cluster peer relationship between the MetroCluster clusters.

About this task

You can use the cluster peer create command to create a peer relationship between a local and remote cluster. After the peer relationship has been created, you can run cluster peer create on the remote cluster to authenticate it to the local cluster.

Before you begin

- · You must have created intercluster LIFs on every node in the clusters that are being peered.
- The clusters must be running ONTAP 9.3 or later.

Steps

1. On the destination cluster, create a peer relationship with the source cluster:

```
cluster peer create -generate-passphrase -offer-expiration MM/DD/YYYY
HH:MM:SS|1...7days|1...168hours -peer-addrs peer LIF IPs -ipspace ipspace
```

If you specify both -generate-passphrase and -peer-addrs, only the cluster whose intercluster LIFs are specified in -peer-addrs can use the generated password.

You can ignore the -ipspace option if you are not using a custom IPspace. For complete command syntax, see the man page.

The following example creates a cluster peer relationship on an unspecified remote cluster:

2. On the source cluster, authenticate the source cluster to the destination cluster:

cluster peer create -peer-addrs peer LIF IPs -ipspace ipspace

For complete command syntax, see the man page.

The following example authenticates the local cluster to the remote cluster at intercluster LIF IP addresses "192.140.112.101" and "192.140.112.102":

```
cluster01::> cluster peer create -peer-addrs
192.140.112.101,192.140.112.102
Notice: Use a generated passphrase or choose a passphrase of 8 or more
characters.
        To ensure the authenticity of the peering relationship, use a
phrase or sequence of characters that would be hard to guess.
Enter the passphrase:
Confirm the passphrase:
Clusters cluster02 and cluster01 are peered.
```

Enter the passphrase for the peer relationship when prompted.

3. Verify that the cluster peer relationship was created:

```
cluster peer show -instance
```

4. Check the connectivity and status of the nodes in the peer relationship:

cluster peer health show

cluster01::> cluster peer health show Node cluster-Name Node-Name Ping-Status 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

Creating a cluster peer relationship (ONTAP 9.2 and earlier)

You can use the cluster peer create command to initiate a request for a peering relationship between a local and remote cluster. After the peer relationship has been requested by the local cluster, you can run cluster peer create on the remote cluster to accept the relationship.

Before you begin

- You must have created intercluster LIFs on every node in the clusters being peered.
- The cluster administrators must have agreed on the passphrase that each cluster will use to authenticate itself to the other.

Steps

1. On the data protection destination cluster, create a peer relationship with the data protection source cluster:

cluster peer create -peer-addrs peer LIF IPs -ipspace ipspace

You can ignore the *-ipspace* option if you are not using a custom IPspace. For complete command syntax, see the man page.

The following example creates a cluster peer relationship with the remote cluster at intercluster LIF IP addresses "192.168.2.201" and "192.168.2.202":

```
cluster02::> cluster peer create -peer-addrs 192.168.2.201,192.168.2.202
Enter the passphrase:
Please enter the passphrase again:
```

Enter the passphrase for the peer relationship when prompted.

2. On the data protection source cluster, authenticate the source cluster to the destination cluster:

cluster peer create -peer-addrs peer LIF IPs -ipspace ipspace

For complete command syntax, see the man page.

The following example authenticates the local cluster to the remote cluster at intercluster LIF IP addresses "192.140.112.203" and "192.140.112.204":

```
cluster01::> cluster peer create -peer-addrs 192.168.2.203,192.168.2.204
Please confirm the passphrase:
Please confirm the passphrase again:
```

Enter the passphrase for the peer relationship when prompted.

3. Verify that the cluster peer relationship was created:

```
cluster peer show -instance
```

For complete command syntax, see the man page.

```
cluster01::> cluster peer show -instance
Peer Cluster Name: cluster01
Remote Intercluster Addresses: 192.168.2.201,192.168.2.202
Availability: Available
Remote Cluster Name: cluster02
Active IP Addresses: 192.168.2.201,192.168.2.202
Cluster Serial Number: 1-80-000013
```

4. Check the connectivity and status of the nodes in the peer relationship:

cluster peer health show`

For complete command syntax, see the man page.

cluster01::> cluster peer health show Node cluster-Name Node-Name Ping-Status 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 provide data protection.

About this task

By default, the root aggregate is created as RAID-DP type aggregate. You can change the root aggregate from RAID-DP to RAID4 type aggregate. The following command modifies the root aggregate for RAID4 type aggregate:

```
storage aggregate modify -aggregate aggr_name -raidtype raid4
```



On non-ADP systems, the RAID type of the aggregate can be modified from the default RAID-DP to RAID4 before or after the aggregate is mirrored.

Steps

1. Mirror the root aggregate:

storage aggregate mirror aggr_name

The following command mirrors the root aggregate for controller_A_1:

controller_A_1::> storage aggregate mirror aggr0_controller_A_1

This mirrors the aggregate, so it consists of a local plex and a remote plex located at the remote

MetroCluster site.

2. Repeat the previous step for each node in the MetroCluster configuration.

Related information

Logical storage management with the CLI

Creating a mirrored data aggregate on each node

You must create a mirrored data aggregate on each node in the DR group.

- You should know what drives 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.
- Drives 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. See Disk and aggregate management.

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 that are to be added to the aggregate
- Number of drives to include



In the minimum-supported configuration, in which a limited number of drives are available, you must use the force-small-aggregate option to allow the creation of a three disk RAID-DP aggregate.

- · 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 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.

The following command creates a mirrored aggregate with 10 disks:

```
cluster_A::> storage aggregate create aggr1_node_A_1 -diskcount 10 -node
node_A_1 -mirror true
[Job 15] Job is queued: Create aggr1_node_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

Creating unmirrored data aggregates

You can optionally create unmirrored data aggregates for data that does not require the redundant mirroring provided by MetroCluster configurations.

Before you begin

- Verify that you know what drives will be used in the new aggregate.
- If you have multiple drive types in your system (heterogeneous storage), you should understand how you can verify that the correct drive type is selected.



In MetroCluster FC configurations, the unmirrored aggregates will only be online after a switchover if the remote disks in the aggregate are accessible. If the ISLs fail, the local node may be unable to access the data in the unmirrored remote disks. The failure of an aggregate can lead to a reboot of the local node.

• Drives 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.



The unmirrored aggregates must be local to the node owning them.

- Aggregate names should conform to the naming scheme you determined when you planned your MetroCluster configuration.
- Disks and aggregates management 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:

storage aggregate create

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 verify that the aggregate is created on a specific node, you should 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 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 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.

The following command creates a unmirrored aggregate with 10 disks:

```
controller_A_1::> storage aggregate create aggr1_controller_A_1
-diskcount 10 -node controller_A_1
[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

Related information

Disk and tier (aggregate) management

Implementing the MetroCluster configuration

You must run the metrocluster configure command to start data protection in a MetroCluster configuration.

Before you begin

• There should be at least two non-root mirrored data aggregates on each cluster.

Additional data aggregates can be either mirrored or unmirrored.

You can verify this with the storage aggregate show command.



If you want to use a single mirrored data aggregate, then see Step 1 for instructions.

• The ha-config state of the controllers and chassis must be "mcc".

About this task

You issue the metrocluster configure command once, on any of the nodes, to enable the MetroCluster configuration. You do not need to issue the command on each of the sites or nodes, and it does not matter which node or site you choose to issue the command on.

The metrocluster configure command automatically pairs the two nodes with the lowest system IDs in each of the two clusters as disaster recovery (DR) partners. In a four-node MetroCluster configuration, there are two DR partner pairs. The second DR pair is created from the two nodes with higher system IDs.



You must **not** configure Onboard Key Manager (OKM) or external key management before you run the command metrocluster configure.

Steps

1. Configure the MetroCluster in the following format:

If your MetroCluster configuration has	Then do this…
Multiple data aggregates	From any node's prompt, configure MetroCluster:
	metrocluster configure node-name
A single mirrored data aggregate	a. From any node's prompt, change to the advanced privilege level:
	set -privilege advanced
	You need to respond with $_{\rm Y}$ when you are prompted to continue into advanced mode and you see the advanced mode prompt (*>).
	b. Configure the MetroCluster with the -allow -with-one-aggregate true parameter:
	<pre>metrocluster configure -allow-with -one-aggregate true node-name</pre>
	c. Return to the admin privilege level:
	set -privilege admin



The best practice is to have multiple data aggregates. If the first DR group has only one aggregate, and you want to add a DR group with one aggregate, you must move the metadata volume off the single data aggregate. For more information on this procedure, see Moving a metadata volume in MetroCluster configurations.

The following command enables the MetroCluster configuration on all of the nodes in the DR group that contains controller_A_1:

```
cluster_A::*> metrocluster configure -node-name controller_A_1
```

[Job 121] Job succeeded: Configure is successful.

2. Verify the networking status on site A:

network port show

The following example shows the network port usage on a four-node MetroCluster configuration:

```
cluster A::> network port show
                                             Speed (Mbps)
Node Port IPspace Broadcast Domain Link MTU
                                             Admin/Oper
_____ _____
controller A 1
                                       9000 auto/1000
     e0a
             Cluster Cluster
                                 up
     e0b
             Cluster Cluster
                                       9000 auto/1000
                                 up
     e0c
            Default Default
                                       1500 auto/1000
                                 up
            Default Default
                                       1500 auto/1000
     e0d
                                  up
                                       1500 auto/1000
            Default Default
     e0e
                                  up
     eOf
            Default Default
                                       1500 auto/1000
                                  up
     e0g
            Default Default
                                  up
                                       1500 auto/1000
controller A 2
     e0a
             Cluster Cluster
                                        9000 auto/1000
                                  up
     e0b
             Cluster Cluster
                                       9000 auto/1000
                                  up
     e0c
            Default Default
                                  up
                                        1500 auto/1000
     e0d
            Default Default
                                        1500 auto/1000
                                  up
     e0e
            Default Default
                                       1500 auto/1000
                                  up
     eOf
            Default Default
                                       1500 auto/1000
                                  up
                                       1500 auto/1000
     e0g
             Default Default
                                  up
14 entries were displayed.
```

3. Verify the MetroCluster configuration from both sites in the MetroCluster configuration.

a. Verify the configuration from site A:

metrocluster show

cluster A::> metrocluster show Entry Name Cluster State _____ _____ ____ Local: cluster A Configuration state configured Mode normal AUSO Failure Domain auso-on-clusterdisaster Remote: cluster B Configuration state configured Mode normal AUSO Failure Domain auso-on-clusterdisaster

b. Verify the configuration from site B:

metrocluster show

cluster_B::> metrocluster Cluster	show Entry Name	State
Local: cluster_B	Configuration state Mode AUSO Failure Domain	configured normal auso-on-cluster-
disaster		
Remote: cluster_A	Configuration state Mode	configured normal
disaster	AUSO FALLULE DOMAIN	auso-on-cluster-

Configuring in-order delivery or out-of-order delivery of frames on ONTAP software

You must configure either in-order delivery (IOD) or out-of-order delivery (OOD) of frames according to the fibre channel (FC) switch configuration.

About this task

If the FC switch is configured for IOD, then the ONTAP software must be configured for IOD. Similarly, if the FC switch is configured for OOD, then ONTAP must be configured for OOD.



You must reboot the controller to change the configuration.

Step

- 1. Configure ONTAP to operate either IOD or OOD of frames.
 - By default, IOD of frames is enabled in ONTAP. To check the configuration details:
 - a. Enter advanced mode:

set advanced

b. Verify the settings:

metrocluster interconnect adapter show

<pre>mcc4-b12_siteB::*> metrocluster interconnect adapter show</pre>					
		Adapter	Link	Is OOD	
Node	Adapter Name	Туре	Status	Enabled?	IP Address
Port Number					
mcc4-b1	fcvi_device_0	FC-VI	Up	false	17.0.1.2
ба					
mcc4-b1	fcvi_device_1	FC-VI	Up	false	18.0.0.2
6b					
mcc4-b1	mlx4_0	IB	Down	false	192.0.5.193
ib2a					
mcc4-b1	mlx4_0	IB	Up	false	192.0.5.194
ib2b					
mcc4-b2	fcvi_device_0	FC-VI	Up	false	17.0.2.2
6a					
mcc4-b2	fcvi_device_1	FC-VI	Up	false	18.0.1.2
60	1 4 0		2	C 1	100 0 0 0
mcc4-b2	mlx4_0	TB	Down	false	192.0.2.9
1DZa	····]4_0	TD	T.T	6 - 1	100 0 0 10
mcc4−b∠	$III \times 4_0$	TR	Up	Laise	192.0.2.10
usu astriac va	re displayed				
o entries we	re ursprayed.				

- $\,\circ\,$ The following steps must be performed on each node to configure OOD of frames:
 - a. Enter advanced mode:

set advanced

b. Verify the MetroCluster configuration settings:

metrocluster interconnect adapter show

mcc4-b12 siteB::*> metrocluster interconnect adapter show Adapter Link Is OOD Adapter Name Type Status Enabled? IP Address Node Port Number ----- ------ ------ ------ ------_____ mcc4-b1 fcvi device 0 FC-VI Up false 17.0.1.2 6a fcvi device 1 FC-VI Up false 18.0.0.2 mcc4-b1 6b mcc4-b1 mlx4 0 IB Down false 192.0.5.193 ib2a mcc4-b1 mlx4 0 IB Up false 192.0.5.194 ib2b mcc4-b2 fcvi device 0 FC-VI Up false 17.0.2.2 6a mcc4-b2 fcvi device 1 FC-VI Up false 18.0.1.2 6b mcc4-b2 mlx4 O IB Down false 192.0.2.9 ib2a mcc4-b2 mlx4 O IB Up false 192.0.2.10 ib2b 8 entries were displayed.

c. Enable OOD on node "mcc4-b1" and node "mcc4-b2":

metrocluster interconnect adapter modify -node node_name -is-ood-enabled
true

mcc4-b12_siteB::*> metrocluster interconnect adapter modify -node mcc4-b1 -is-ood-enabled true mcc4-b12_siteB::*> metrocluster interconnect adapter modify -node mcc4-b2 -is-ood-enabled true

- d. Reboot the controller by performing a high-availability (HA) takeover in both directions.
- e. Verify the settings:

metrocluster interconnect adapter show

mcc4-b12 siteB::*> metrocluster interconnect adapter show Adapter Link Is OOD Node Adapter Name Туре Status Enabled? IP Address Port Number _____ ____ _____ _____ _____ mcc4-b1 fcvi device 0 FC-VI Up true 17.0.1.2 6a mcc4-b1 fcvi device 1 FC-VI Up 18.0.0.2 true 6b mcc4-b1 mlx4 0 192.0.5.193 IB Down false ib2a mlx4 0 192.0.5.194 mcc4-b1 IB Up false ib2b mcc4-b2 fcvi device 0 FC-VI 17.0.2.2 Up true 6a mcc4-b2 fcvi device 1 FC-VI Up 18.0.1.2 true 6b mcc4-b2 mlx4 O IB Down false 192.0.2.9 ib2a 192.0.2.10 mcc4-b2 mlx4 0 ΙB ЧU false ib2b 8 entries were displayed.

Configuring SNMPv3 in a MetroCluster configuration

Before you begin

The authentication and privacy protocols on the switches and on the ONTAP system must be the same.

About this task

ONTAP currently supports AES-128 encryption.

Steps

1. Create an SNMP user for each switch from the controller prompt:

```
security login create
```

```
Controller_A_1::> security login create -user-or-group-name snmpv3user
-application snmp -authentication-method usm -role none -remote-switch
-ipaddress 10.10.10.10
```

2. Respond to the following prompts as required at your site:



For EngineID, press **ENTER** to assign the default value.

```
Enter the authoritative entity's EngineID [remote EngineID]:
Which authentication protocol do you want to choose (none, md5, sha,
sha2-256) [none]: sha
Enter the authentication protocol password (minimum 8 characters long):
Enter the authentication protocol password again:
Which privacy protocol do you want to choose (none, des, aes128) [none]:
enter privacy protocol password (minimum 8 characters long):
Enter privacy protocol password again:
```



The same username can be added to different switches with different IP addresses.

3. Create an SNMP user for the rest of the switches.

The following example shows how to create a username for a switch with the IP address 10.10.10.11.

```
Controller_A_1::> security login create -user-or-group-name snmpv3user
-application snmp -authentication-method usm -role none -remote-switch
-ipaddress 10.
10.10.11
```

4. Check that there is one login entry for each switch:

security login show

```
Controller A 1::> security login show -user-or-group-name snmpv3user
-fields remote-switch-ipaddress
vserver user-or-group-name application authentication-method
remote-switch-ipaddress
_____
node A 1 SVM 1 snmpv3user snmp
                               usm
10.10.10.10
node A 1 SVM 2 snmpv3user snmp
                               usm
10.10.10.11
node_A_1_SVM_3_snmpv3user
                     snmp usm
10.10.10.12
node A 1 SVM 4 snmpv3user snmp
                               usm
10.10.10.13
4 entries were displayed.
```

5. Configure SNMPv3 on the switches from the switch prompt:

Brocade switches (FOS 9.0 and later)

```
snmpconfig --add snmpv3 -index <index> -user <user_name> -groupname <rw/ro>
-auth_proto <auth_protocol> -auth_passwd <auth_password> -priv_proto
<priv_protocol> -priv_passwd <priv_password>
```

Brocade switches (FOS 8.x and earlier)

snmpconfig --set snmpv3

The example shows how to configure a read-only user. You can adjust the RW users if needed. If you require RO access, after "User (ro):" specify the "snmpv3user".

```
Switch-A1:admin> snmpconfig --set snmpv3
SNMP Informs Enabled (true, t, false, f): [false] true
SNMPv3 user configuration(snmp user not configured in FOS user
database will have physical AD and admin role as the default):
User (rw): [snmpadmin1]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
Priv Protocol [DES(1)/noPriv(2)/AES128(3)/AES256(4)]): (2..2) [2]
Engine ID: [00:00:00:00:00:00:00]
User (ro): [snmpuser2] snmpv3user
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [2]
Priv Protocol [DES(1)/noPriv(2)/AES128(3)/AES256(4)]): (2..2) [3]
```

Cisco switches

```
snmp-server user <user_name> auth [md5/sha/sha-256] <auth_password> priv
(aes-128) <priv password>
```



You should also set passwords on unused accounts to secure them and use the best encryption available in your ONTAP release.

6. Configure encryption and passwords on the remaining switch users as required on your site.

Configuring MetroCluster components for health monitoring

You must perform some special configuration steps before monitoring the components in a MetroCluster configuration.



For improved security, NetApp recommends that you configure SNMPv2 or SNMPv3 to monitor the switch health.

About this task

These tasks apply only to systems with FC-to-SAS bridges.

Beginning in Fabric OS 9.0.1, SNMPv2 is not supported for health monitoring on Brocade switches, you must use SNMPv3 instead. If you are using SNMPv3, you must configure SNMPv3 in ONTAP before proceeding to the following section. For more details, see Configuring SNMPv3 in a MetroCluster configuration.



- You should place bridges and a node management LIF in a dedicated network to avoid interference from other sources.
- If you use a dedicated network for health monitoring, then each node must have a node management LIF in that dedicated network.

NetApp only supports the following tools to monitor the components in a MetroCluster FC configuration:

- Brocade Network Advisor (BNA)
- Brocade SANnav
- Active IQ Config Advisor
- NetApp Health Monitoring (ONTAP)
- MetroCluster data collector (MC_DC)

Configuring the MetroCluster FC switches for health monitoring

In a fabric-attached MetroCluster configuration, you must perform some additional configuration steps to monitor the FC switches.



Beginning with ONTAP 9.8, the storage switch command is replaced with system switch fibre-channel. The following steps show the storage switch command, but if you are running ONTAP 9.8 or later, the system switch fibre-channel command is preferred.

Steps

1. Add a switch with an IP address to each MetroCluster node:

The command you run depends on whether you are using SNMPv2 or SNMPv3.

Add a switch using SNMPv3:

```
storage switch add -address <ip_adddress> -snmp-version SNMPv3 -snmp
-community-or-username <SNMP_user_configured_on_the_switch>
```

Add a switch using SNMPv2:

storage switch add -address ipaddress

This command must be repeated on all four switches in the MetroCluster configuration.



Brocade 7840 FC switches and all alerts are supported in health monitoring, except NoISLPresent_Alert.

The following example shows the command to add a switch with IP address 10.10.10.10:

controller A 1::> storage switch add -address 10.10.10.10

2. Verify that all switches are properly configured:

```
storage switch show
```
It might take up to 15 minutes to reflect all data due to the 15-minute polling interval.

The following example shows the command given to verify that the MetroCluster 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 worldwide name (WWN) of the switch is shown, the ONTAP health monitor can contact and monitor the FC switch.

Related information

System administration

Configuring FC-to-SAS bridges for health monitoring

In systems running ONTAP versions prior to 9.8, you must perform some special configuration steps to monitor the FC-to-SAS bridges in the MetroCluster configuration.

About this task

- Third-party SNMP monitoring tools are not supported for FibreBridge bridges.
- Beginning with ONTAP 9.8, FC-to-SAS bridges are monitored via in-band connections by default, and additional configuration is not required.



Beginning with ONTAP 9.8, the storage bridge command is replaced with system bridge. The following steps show the storage bridge command, but if you are running ONTAP 9.8 or later, the system bridge command is preferred.

Steps

- 1. From the ONTAP cluster prompt, add the bridge to health monitoring:
 - a. Add the bridge, using the command for your version of ONTAP:

ONTAP version	Command
---------------	---------

9.5 and later	storage bridge add -address 0.0.0.0 -managed-by in-band -name bridge-name
9.4 and earlier	storage bridge add -address bridge-ip-address -name bridge-name

b. Verify that the bridge has been added and is properly configured:

storage bridge show

It might take as long as 15 minutes to reflect all data because of the polling interval. The ONTAP health monitor can contact and monitor the bridge if the value in the "Status" column is "ok", and other information, such as the worldwide name (WWN), is displayed.

The following example shows that the FC-to-SAS bridges are configured:

controller_A_1::> storage bridge show				
Bridge Symbolic Name Is Monitored Monitor Status Vendor Model Bridge WWN				Status
ATTO_10.10.20.10	atto01	true	ok	Atto
FibreBridge 7500N	2000	0010867038c0		
ATTO_10.10.20.11	atto02	true	ok	Atto
FibreBridge 7500N	2000	0010867033c0		
ATTO_10.10.20.12	atto03	true	ok	Atto
FibreBridge 7500N	2000	0010867030c0		
ATTO_10.10.20.13	atto04	true	ok	Atto
FibreBridge 7500N	2000	001086703b80		
4 entries were displayed				
<pre>controller_A_1::></pre>				

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. You should also do a check before a negotiated (planned) switchover or a switchback operation.

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, then will not show the expected output.

Steps

1. Check the configuration:

```
metrocluster check run
```

The command runs as a background job and might not be completed immediately.

```
cluster_A::> metrocluster check run
The operation has been started and is running in the background. Wait
for
it to complete and run "metrocluster check show" to view the results. To
check the status of the running metrocluster check operation, use the
command,
"metrocluster operation history show -job-id 2245"
```

```
cluster A::> metrocluster check show
Component
                 Result
----- -----
nodes
                 ok
lifs
                 ok
config-replication ok
aggregates
                ok
clusters
                ok
connections
                ok
volumes
                 ok
7 entries were displayed.
```

2. Display more detailed results from the most recent metrocluster check run command:

```
metrocluster check aggregate show
metrocluster check cluster 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 so that the information displayed is current.

The following example shows the metrocluster check aggregate show command output for a

cluster_A::> metroclu	aster check aggregate show	
Last Checked On: 8/5/	2014 00:42:58	
Node Result	Aggregate	Check
controller_A_1	controller_A_1_aggr0	mirroring-status
ok		disk-pool-allocation
ok		ownership-state
ok	controllor & 1 aggr1	
ok	controller_A_1_aggl1	mirroring-status
07		disk-pool-allocation
ok		ownership-state
ok	controller A 1 aggr2	
	>>	mirroring-status
ok		disk-pool-allocation
OK		ownership-state
ok		
controller A 2	controller A 2 aggr0	
		mirroring-status
ok		disk-pool-allocation
ok		ownorship-state
ok		OwnerShip-State
	controller_A_2_aggr1	mirroring-status
ok		disk-pool-allocation
ok		
		ownership-state

ok	controller A 2 aggr2	
	001101101_1_2_49912	mirroring-status
ok		disk-pool-allocation
ok		europahin atata
ok		ownersnip-state
18 entries were disp	layed.	

The following example shows the metrocluster check cluster show command output for a healthy four-node MetroCluster configuration. It indicates that the clusters are ready to perform a negotiated switchover if necessary.

Last Checked On: 9/13/2017 20:47:04				
Cluster	Check Result			
mccint-fas9000-0102				
	negotiated-switchover-ready	not-applicable		
	switchback-ready	not-applicable		
	job-schedules	ok		
	licenses	ok		
	periodic-check-enabled	ok		
mccint-fas9000-0304				
	negotiated-switchover-ready	not-applicable		
	switchback-ready	not-applicable		
	job-schedules	ok		
	licenses	ok		
	periodic-check-enabled	ok		
10 entries were displayed.				

Related information

Disk and aggregate management

Network and LIF management

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.



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

Steps

1. Go to the Config Advisor download page and download the tool.

NetApp Downloads: Config Advisor

2. Run Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

Verifying local HA operation

If you have a four-node MetroCluster configuration, you should verify the operation of the local HA pairs in the MetroCluster configuration. This is not required for two-node configurations.

About this task

Two-node MetroCluster configurations do not consist of local HA pairs and this task does not apply.

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

The output should indicate that takeover is possible for both nodes:

```
cluster_A::> storage failover show
Takeover
Node Partner 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. Take over 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

The output should indicate that controller_A_1 is in takeover state, meaning that it has taken over its HA partner:

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

The output should indicate that takeover is possible for both nodes:

- f. Repeat the previous substeps, this time taking over controller_A_1 from controller_A_2.
- 2. Repeat the preceding steps on cluster_B.

Related information

High-availability configuration

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 that are mentioned in the Recover from a disaster.

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

The Cluster Management with the CLI contains additional information under the section *Managing configuration backups*.

Considerations for using virtual IP and Border Gateway Protocol with a MetroCluster configuration

Beginning with ONTAP 9.5, ONTAP supports layer 3 connectivity using virtual IP (VIP) and Border Gateway Protocol (BGP). The combination VIP and BGP for redundancy in the front-end networking with the back-end MetroCluster redundancy provides a layer 3 disaster recovery solution.

Review the following guidelines and illustration when planning your layer 3 solution. For details on implementing VIP and BGP in ONTAP, refer to Configure virtual IP LIFs.



ONTAP limitations

ONTAP does not automatically verify that all nodes on both sites of the MetroCluster configuration are configured with BGP peering.

ONTAP does not perform route aggregation but announces all individual virtual LIF IPs as unique host routes at all times.

ONTAP does not support true AnyCast — only a single node in the cluster presents a specific virtual LIF IP (but is accepted by all physical interfaces, regardless of whether they are BGP LIFs, provided the physical port is part of the correct IPspace). Different LIFs can migrate independently of each other to different hosting nodes.

Guidelines for using this Layer 3 solution with a MetroCluster configuration

You must configure your BGP and VIP correctly to provide the required redundancy.

Simpler deployment scenarios are preferred over more complex architectures (for example, a BGP peering router is reachable across an intermediate, non-BGP router). However, ONTAP does not enforce network design or topology restrictions.

VIP LIFs only cover the frontend/data network.

Depending on your version of ONTAP, you must configure BGP peering LIFs in the node SVM, not the system or data SVM. In ONTAP 9.8, the BGP LIFs are visible in the cluster (system) SVM and the node SVMs are no longer present.

Each data SVM requires the configuration of all potential first hop gateway addresses (typically, the BGP router peering IP address), so that the return data path is available if a LIF migration or MetroCluster failover occurs.

BGP LIFs are node specific, similar to intercluster LIFs — each node has a unique configuration, which does not need to be replicated to DR site nodes.

Once configured, the existence of the v0a (v0b and so on) continuously validates the connectivity, guaranteeing that a LIF migrate or failover succeeds (unlike L2, where a broken configuration is only visible after the outage).

A major architectural difference is that clients should no longer share the same IP subnet as the VIP of data SVMs. An L3 router with appropriate enterprise grade resiliency and redundancy features enabled (for example, VRRP/HSRP) should be on the path between storage and clients for the VIP to operate correctly.

The reliable update process of BGP allows for smoother LIF migrations because they are marginally faster and have a lower chance of interruption to some clients

You can configure BGP to detect some classes of network or switch misbehaviors faster than LACP, if configured accordingly.

External BGP (EBGP) uses different AS numbers between ONTAP node(s) and peering routers and is the preferred deployment to ease route aggregation and redistribution on the routers. Internal BGP (IBGP) and the use of route reflectors is not impossible but outside the scope of a straightforward VIP setup.

After deployment, you must check that the data SVM is accessible when the associated virtual LIF is migrated between all nodes on each site (including MetroCluster switchover) to verify the correct configuration of the static routes to the same data SVM.

VIP works for most IP-based protocols (NFS, SMB, iSCSI).

Testing the MetroCluster configuration

You can test failure scenarios to confirm the correct operation of the MetroCluster

configuration.

Verifying negotiated switchover

You can test the negotiated (planned) switchover operation to confirm uninterrupted data availability.

About this task

This test validates that data availability is not affected (except for Microsoft Server Message Block (SMB) and Solaris Fibre Channel protocols) by switching the cluster over to the second data center.

This test should take about 30 minutes.

This procedure has the following expected results:

• The metrocluster switchover command will present a warning prompt.

If you respond yes to the prompt, the site the command is issued from will switch over the partner site.

For MetroCluster IP configurations:

- For ONTAP 9.4 and earlier:
 - Mirrored aggregates will become degraded after the negotiated switchover.
- For ONTAP 9.5 and later:
 - Mirrored aggregates will remain in normal state if the remote storage is accessible.
 - Mirrored aggregates will become degraded after the negotiated switchover if access to the remote storage is lost.
- For ONTAP 9.8 and later:
 - Unmirrored aggregates that are located at the disaster site will become unavailable if access to the remote storage is lost. This might lead to a controller outage.

Steps

1. Confirm that all nodes are in the configured state and normal mode:

metrocluster node show

```
      cluster_A::>
      metrocluster node show

      Cluster
      Configuration State
      Mode

      ------
      ------
      Mode

      ------
      ------
      ------

      Local: cluster_A
      configured
      normal

      Remote: cluster_B
      configured
      normal
```

2. Begin the switchover operation:

```
metrocluster switchover
```

```
cluster_A::> metrocluster switchover
Warning: negotiated switchover is about to start. It will stop all the
data Vservers on cluster "cluster_B" and
automatically re-start them on cluster "`cluster_A`". It will finally
gracefully shutdown cluster "cluster B".
```

3. Confirm that the local cluster is in the configured state and switchover mode:

metrocluster node show

4. Confirm that the switchover operation was successful:

metrocluster operation show

```
cluster_A::> metrocluster operation show
cluster_A::> metrocluster operation show
Operation: switchover
State: successful
Start Time: 2/6/2016 13:28:50
End Time: 2/6/2016 13:29:41
Errors: -
```

5. Use the vserver show and network interface show commands to verify that DR SVMs and LIFs have come online.

Verifying healing and manual switchback

You can test the healing and manual switchback operations to verify that data availability is not affected (except for SMB and Solaris FC configurations) by switching back the cluster to the original data center after a negotiated switchover.

About this task

This test should take about 30 minutes.

The expected result of this procedure is that services should be switched back to their home nodes.

Steps

1. Verify that healing is completed:

metrocluster node show

The following example shows the successful completion of the command:

2. Verify that all aggregates are mirrored:

storage aggregate show

The following example shows that all aggregates have a RAID Status of mirrored:

cluster A::> storage aggregate show cluster Aggregates: Aggregate Size Available Used% State #Vols Nodes RAID Status _____ data cluster 4.19TB 4.13TB 2% online 8 node_A_1 raid_dp, mirrored, normal root cluster 715.5GB 212.7GB 70% online 1 node A 1 raid4, mirrored, normal cluster B Switched Over Aggregates: Aggregate Size Available Used% State #Vols Nodes RAID Status _____ data cluster B 4.19TB 4.11TB 2% online 5 node_A_1 raid_dp, mirrored, normal root_cluster_B - - - unknown - node_A_1 -

- 3. Boot the nodes from the disaster site.
- 4. Check the status of switchback recovery:

metrocluster node show

5. Perform the switchback:

metrocluster switchback

```
cluster_A::> metrocluster switchback
[Job 938] Job succeeded: Switchback is successful.Verify switchback
```

6. Confirm the status of the nodes:

```
metrocluster node show
```

7. Confirm the status:

metrocluster operation show

The output should show a successful state.

```
cluster_A::> metrocluster operation show
Operation: switchback
State: successful
Start Time: 2/6/2016 13:54:25
End Time: 2/6/2016 13:56:15
Errors: -
```

Loss of a single FC-to-SAS bridge

You can test the failure of a single FC-to-SAS bridge to make sure there is no single point of failure.

About this task

This test should take about 15 minutes.

This procedure has the following expected results:

- Errors should be generated as the bridge is switched off.
- No failover or loss of service should occur.
- Only one path from the controller module to the drives behind the bridge is available.



Beginning with ONTAP 9.8, the storage bridge command is replaced with system bridge. The following steps show the storage bridge command, but if you are running ONTAP 9.8 or later, the system bridge command is preferred.

Steps

- 1. Turn off the power supplies of the bridge.
- 2. Confirm that the bridge monitoring indicates an error:

storage bridge show

3. Confirm that the drives behind the bridge are available with a single path:

storage disk error show

cluster A::> storage disk error show Error Type Error Text Disk -----_____ 1.0.0 onedomain 1.0.0 (5000cca057729118): All paths to this array LUN are connected to the same fault domain. This is a single point of failure. onedomain 1.0.1 (5000cca057727364): All paths 1.0.1 to this array LUN are connected to the same fault domain. This is a single point of failure. 1.0.2 onedomain 1.0.2 (5000cca05772e9d4): All paths to this array LUN are connected to the same fault domain. This is a single point of failure. . . . 1.0.23 onedomain 1.0.23 (5000cca05772e9d4): All paths to this array LUN are connected to the same fault domain. This is a single point of failure.

Verifying operation after power line disruption

You can test the MetroCluster configuration's response to the failure of a PDU.

About this task

The best practice is for each power supply unit (PSU) in a component to be connected to separate power supplies. If both PSUs are connected to the same power distribution unit (PDU) and an electrical disruption occurs, the site could down or a complete shelf might become unavailable. Failure of one power line is tested to confirm that there is no cabling mismatch that could cause a service disruption.

This test should take about 15 minutes.

This test requires turning off power to all left-hand PDUs and then all right-hand PDUs on all of the racks containing the MetroCluster components.

This procedure has the following expected results:

- Errors should be generated as the PDUs are disconnected.
- · No failover or loss of service should occur.

Steps

- 1. Turn off the power of the PDUs on the left-hand side of the rack containing the MetroCluster components.
- 2. Monitor the result on the console:

```
system environment sensors show -state fault
```

storage shelf show -errors

cluster A::> system environment sensors show -state fault Node Sensor State Value/Units Crit-Low Warn-Low Warn-Hi Crit-Hi ____ _____ _____ node A 1 PSU1 fault PSU OFF PSU1 Pwr In OK fault FAULT node A 2 PSU1 fault PSU OFF PSU1 Pwr In OK fault FAULT 4 entries were displayed. cluster A::> storage shelf show -errors Shelf Name: 1.1 Shelf UID: 50:0a:09:80:03:6c:44:d5 Serial Number: SHFHU1443000059 Error Type Description _____ Critical condition is detected in storage shelf Power power supply unit "1". The unit might fail.Reconnect PSU1

- 3. Turn the power back on to the left-hand PDUs.
- 4. Make sure that ONTAP clears the error condition.
- 5. Repeat the previous steps with the right-hand PDUs.

Verifying operation after a switch fabric failure

You can disable a switch fabric to show that data availability is not affected by the loss.

About this task

This test should take about 15 minutes.

The expected result of this procedure is that disabling a fabric results in all cluster interconnect and disk traffic flowing to the other fabric.

In the examples shown, switch fabric 1 is disabled. This fabric consists of two switches, one at each MetroCluster site:

• FC_switch_A_1 on cluster_A

• FC_switch_B_1 on cluster_B

Steps

- 1. Disable connectivity to one of the two switch fabrics in the MetroCluster configuration:
 - a. Disable the first switch in the fabric:

switchdisable

FC switch A 1::> switchdisable

b. Disable the second switch in the fabric:

```
switchdisable
```

```
FC_switch_B_1::> switchdisable
```

2. Monitor the result on the console of the controller modules.

You can use the following commands to check the cluster nodes to make sure that all data is still being served. The command output shows missing paths to disks. This is expected.

- vserver show
- network interface show
- aggr show
- · system node runnodename-command storage show disk -p
- storage disk error show
- 3. Reenable connectivity to one of the two switch fabrics in the MetroCluster configuration:
 - a. Reenable the first switch in the fabric:

switchenable

FC_switch_A_1::> switchenable

b. Reenable the second switch in the fabric:

switchenable

```
FC_switch_B_1::> switchenable
```

4. Wait at least 10 minutes and then repeat the above steps on the other switch fabric.

Verifying operation after loss of a single storage shelf

You can test the failure of a single storage shelf to verify that there is no single point of failure.

About this task

This procedure has the following expected results:

- An error message should be reported by the monitoring software.
- No failover or loss of service should occur.
- Mirror resynchronization starts automatically after the hardware failure is restored.

Steps

1. Check the storage failover status:

storage failover show

```
cluster_A::> storage failover show
Node Partner Possible State Description
------
node_A_1 node_A_2 true Connected to node_A_2
node_A_2 node_A_1 true Connected to node_A_1
2 entries were displayed.
```

2. Check the aggregate status:

storage aggregate show

```
cluster A::> storage aggregate show
cluster Aggregates:
Aggregate Size Available Used% State #Vols Nodes RAID
Status
_____
node_A_1data01_mirrored
        4.15TB 3.40TB 18% online 3 node_A_1
raid dp,
mirrored,
normal
node A 1root
       707.7GB 34.29GB 95% online 1 node_A_1
raid dp,
mirrored,
normal
node_A_2_data01_mirrored
        4.15TB 4.12TB 1% online 2 node_A_2
raid dp,
mirrored,
normal
node A 2 data02 unmirrored
        2.18TB 2.18TB 0% online 1 node_A_2
raid dp,
normal
node A 2 root
       707.7GB 34.27GB 95% online 1 node_A_2
raid dp,
mirrored,
normal
```

3. Verify that all data SVMs and data volumes are online and serving data:

```
vserver show -type data
```

```
network interface show -fields is-home false
```

```
volume show !vol0, !MDV*
```

cluster A::> vserver show -type data cluster A::> vserver show -type data Admin Operational Root State State Volum State Volume Vserver Type Subtype Aggregate _____ SVM1 data sync-source running SVM1_root node A 1 data01 mirrored SVM2 sync-source running SVM2 root data node_A_2_data01_mirrored cluster A::> network interface show -fields is-home false There are no entries matching your query. cluster A::> volume show !vol0,!MDV* Vserver Volume Aggregate State Type Size Available Used% _____ _ SVM1 SVM1 root node A 1data01 mirrored online RW 10GB 9.50GB 5% SVM1 SVM1 data vol node A 1data01 mirrored online RW 10GB 9.49GB 5% SVM2 SVM2 root node A 2 data01 mirrored online RW 10GB 9.49GB 5% SVM2 SVM2 data vol node A 2 data02 unmirrored online RW 1GB 972.6MB 5%

4. Identify a shelf in Pool 1 for node node_A_2 to power off to simulate a sudden hardware failure:

storage aggregate show -r -node node-name !*root

The shelf you select must contain drives that are part of a mirrored data aggregate.

In the following example, shelf ID 31 is selected to fail.

```
cluster A::> storage aggregate show -r -node node A 2 !*root
Owner Node: node A 2
Aggregate: node A 2 data01 mirrored (online, raid dp, mirrored) (block
checksums)
 Plex: /node A 2 data01 mirrored/plex0 (online, normal, active, pool0)
  RAID Group /node A 2 data01 mirrored/plex0/rg0 (normal, block
checksums)
                                                    Usable
Physical
   Position Disk
                                  Pool Type RPM Size
Size Status
    _____ _
_____ ___
   dparity 2.30.3
                                    0 BSAS
                                              7200 827.7GB
828.0GB (normal)
                                    0 BSAS 7200 827.7GB
    parity 2.30.4
828.0GB (normal)
   data 2.30.6
                                    0 BSAS 7200 827.7GB
828.0GB (normal)
   data 2.30.8
                                    0 BSAS 7200 827.7GB
828.0GB (normal)
    data 2.30.5
                                    0 BSAS 7200 827.7GB
828.0GB (normal)
 Plex: /node A 2 data01 mirrored/plex4 (online, normal, active, pool1)
  RAID Group /node A 2 data01 mirrored/plex4/rg0 (normal, block
checksums)
                                                    Usable
Physical
    Position Disk
                                  Pool Type RPM
                                                     Size
Size Status
    _____ ___
    dparity 1.31.7
                                   1 BSAS 7200 827.7GB
828.0GB (normal)
   parity 1.31.6
                                   1 BSAS 7200 827.7GB
828.0GB (normal)
    data 1.31.3
                                    1 BSAS 7200 827.7GB
828.0GB (normal)
```

data 1.31.4 BSAS 7200 827.7GB 1 828.0GB (normal) data 1.31.5 1 BSAS 7200 827.7GB 828.0GB (normal) Aggregate: node A 2 data02 unmirrored (online, raid dp) (block checksums) Plex: /node A 2 data02 unmirrored/plex0 (online, normal, active, pool0) RAID Group /node A 2 data02 unmirrored/plex0/rg0 (normal, block checksums) Usable Physical Position Disk Pool Type RPM Size Size Status _____ ____ _____ _____ _____ ___ dparity 2.30.12 0 BSAS 7200 827.7GB 828.0GB (normal) parity 2.30.22 0 BSAS 7200 827.7GB 828.0GB (normal) data 2.30.21 BSAS 7200 827.7GB 0 828.0GB (normal) 2.30.20 data 0 BSAS 7200 827.7GB 828.0GB (normal) 7200 827.7GB 2.30.14 data 0 BSAS 828.0GB (normal) 15 entries were displayed.

- 5. Physically power off the shelf that you selected.
- 6. Check the aggregate status again:

storage aggregate show

storage aggregate show -r -node node A 2 !*root

The aggregate with drives on the powered-off shelf should have a "degraded" RAID status, and drives on the affected plex should have a "failed" status, as shown in the following example:

cluster_A::> storage aggregate show Aggregate Size Available Used% State #Vols Nodes RAID Status -----node_A_1data01_mirrored 4.15TB 3.40TB 18% online 3 node_A_1 raid_dp, mirrored, normal node A 1root 707.7GB 34.29GB 95% online 1 node A 1 raid dp, mirrored, normal node A 2 data01 mirrored 4.15TB 4.12TB 1% online 2 node_A_2 raid_dp, mirror degraded node A 2 data02 unmirrored 2.18TB 2.18TB 0% online 1 node_A_2 raid_dp, normal node A 2 root 707.7GB 34.27GB 95% online 1 node_A_2 raid dp, mirror degraded cluster A::> storage aggregate show -r -node node A 2 !*root Owner Node: node A 2 Aggregate: node A 2 data01 mirrored (online, raid dp, mirror degraded) (block checksums) Plex: /node A 2 data01 mirrored/plex0 (online, normal, active, pool0) RAID Group /node A 2 data01 mirrored/plex0/rg0 (normal, block checksums) Usable Physical Pool Type RPM Size Position Disk Size Status ----- ----- ----- -----_____ ____ 0 BSAS 7200 827.7GB dparity 2.30.3 828.0GB (normal)

parity 2.30.4 0 BSAS 7200 827.7GB 828.0GB (normal) data 2.30.6 0 BSAS 7200 827.7GB 828.0GB (normal) data 2.30.8 0 BSAS 7200 827.7GB 828.0GB (normal) data 2.30.5 7200 827.7GB 0 BSAS 828.0GB (normal) Plex: /node A 2 data01 mirrored/plex4 (offline, failed, inactive, pooll) RAID Group /node A 2 data01 mirrored/plex4/rg0 (partial, none checksums) Usable Physical Position Disk Pool Type RPM Size Size Status _____ ____ dparity FAILED - 827.7GB _ - (failed) parity FAILED - 827.7GB - (failed) data FAILED - 827.7GB - (failed) - 827.7GB data FAILED - (failed) data FAILED - 827.7GB - -- (failed) Aggregate: node A 2 data02 unmirrored (online, raid dp) (block checksums) Plex: /node A 2_data02_unmirrored/plex0 (online, normal, active, pool0) RAID Group /node A 2 data02 unmirrored/plex0/rg0 (normal, block checksums) Usable Physical Position Disk Pool Type RPM Size Size Status _____ dparity 2.30.12 0 BSAS 7200 827.7GB 828.0GB (normal) parity 2.30.22 0 BSAS 7200 827.7GB 828.0GB (normal)

```
7200 827.7GB
    data 2.30.21
                                       0
                                           BSAS
828.0GB (normal)
   data
            2.30.20
                                       0
                                           BSAS
                                                  7200 827.7GB
828.0GB (normal)
   data 2.30.14
                                       0 BSAS
                                                  7200 827.7GB
828.0GB (normal)
15 entries were displayed.
```

7. Verify that the data is being served and that all volumes are still online:

vserver show -type data
network interface show -fields is-home false
volume show !vol0,!MDV*

cluster A::> vserver show -type data cluster A::> vserver show -type data Admin Operational Root Vserver Type Subtype State State Volume Aggregate _____ SVM1 data sync-source running SVM1_root node A 1 data01 mirrored SVM2 data sync-source running SVM2 root node A_1_data01_mirrored cluster A::> network interface show -fields is-home false There are no entries matching your query. cluster_A::> volume show !vol0,!MDV* Vserver Volume Aggregate State Type Size Available Used% _____ ____ _____ _ SVM1 SVM1 root node A 1data01 mirrored online RW 10GB 9.50GB 5% SVM1 SVM1 data vol node A 1data01 mirrored online RW 10GB 9.49GB 5% SVM2 SVM2 root node A 1data01 mirrored online RW 10GB 9.49GB 5% SVM2 SVM2 data vol node A 2 data02 unmirrored online RW 1GB 972.6MB 5%

8. Physically power on the shelf.

Resynchronization starts automatically.

9. Verify that resynchronization has started:

storage aggregate show

The affected aggregate should have a "resyncing" RAID status, as shown in the following example:

```
cluster A::> storage aggregate show
cluster Aggregates:
Aggregate Size Available Used% State #Vols Nodes RAID
Status
----- ----- ------ ----- ------ ------
_____
node A 1 data01 mirrored
        4.15TB 3.40TB 18% online 3 node_A_1
raid dp,
mirrored,
normal
node A 1 root
        707.7GB 34.29GB 95% online 1 node A 1
raid dp,
mirrored,
normal
node_A_2_data01_mirrored
        4.15TB 4.12TB 1% online 2 node_A_2
raid dp,
resyncing
node A 2_data02_unmirrored
         2.18TB 2.18TB 0% online 1 node_A_2
raid dp,
normal
node A 2 root
        707.7GB 34.27GB 95% online 1 node_A_2
raid dp,
resyncing
```

10. Monitor the aggregate to confirm that resynchronization is complete:

storage aggregate show

The affected aggregate should have a "normal" RAID status, as shown in the following example:

cluster A::> storage aggregate show cluster Aggregates: Aggregate Size Available Used% State #Vols Nodes RAID Status _____ ____ _____ node A 1data01 mirrored 4.15TB 3.40TB 18% online 3 node_A_1 raid dp, mirrored, normal node A 1root 707.7GB 34.29GB 95% online 1 node_A_1 raid dp, mirrored, normal node_A_2_data01_mirrored 4.15TB 4.12TB 1% online 2 node A 2 raid dp, normal node A 2 data02 unmirrored 2.18TB 2.18TB 0% online 1 node A 2 raid dp, normal node A 2 root 707.7GB 34.27GB 95% online 1 node A 2 raid dp, resyncing

Remove MetroCluster configurations

If you need to remove the MetroCluster configuration, contact technical support.

Contact NetApp technical support and reference the appropriate guide for your configuration from How to remove nodes from a MetroCluster configuration - Resolution Guide.



You cannot reverse the MetroCluster unconfiguration. This process should only be done with the assistance of technical support. After removing the MetroCluster configuration, all disk connectivity and interconnects should be adjusted to be in a supported state.

How to use the Active IQ Unified Manager and ONTAP System Manager for further configuration and monitoring

Use the Active IQ Unified Manager and ONTAP System Manager for further configuration and monitoring

The Active IQ Unified Manager and ONTAP System Manager can be used for GUI management of the clusters and monitoring the configuration.

Each node has ONTAP 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 Active IQ Unified Manager to monitor the MetroCluster configuration.

Related information

Active IQ Unified Manager Documentation

Synchronize the system time using NTP

Each cluster needs its own Network Time Protocol (NTP) server to synchronize the time between the nodes and their clients.

About this task

- You cannot modify the time zone settings for a failed node or the partner node after a takeover occurs.
- Each cluster in the MetroCluster FC configuration should have its own separate NTP server or servers used by the nodes, FC switches, and FC-to-SAS bridges at that MetroCluster site.
- If you are using the MetroCluster Tiebreaker software, it should also have its own separate NTP server.

Depending on your ONTAP version, you can configure the NTP from the **Cluster** or **Insights** tab in the System Manager UI.

Cluster

In System Manager, you can configure the NTP from the **Cluster** tab using two different options, depending on your ONTAP version:

ONTAP 9.8 or later:

Use the following steps to synchronize the NTP from the **Cluster** tab in ONTAP 9.8 or later.

Steps

- 1. Go to Cluster > Overview
- 2. Then select the **More** option and select **Edit**.
- 3. In the Edit Cluster Details window, select the +Add option below NTP Servers.
- 4. Add the name, location, and specify the IP address of the time server.
- 5. Then, select Save.
- 6. Repeat the steps for any additional time servers.

ONTAP 9.11.1 or later:

Use the following steps to synchronize the NTP from the **Insights** window in the **Cluster** tab in ONTAP 9.11.1 or later.

Steps

- 1. Go to Cluster > Overview
- 2. Scroll down to the **Insights** window on the page, locate **Too few NTP servers are configured**, and then select **Fix It**.
- 3. Specify the IP address of the time server, and then select **Save**.
- 4. Repeat the previous step for any additional time servers.

Insights

In ONTAP 9.11.1 or later, you can also configure the NTP by using the **Insights** tab in System Manager:

Steps

- 1. Go to the Insights tab in the System Manager UI.
- 2. Scroll down to Too few NTP servers are configured and select Fix It.
- 3. Specify the IP address of the time server, and then select **Save**.
- 4. Repeat the previous step for any additional time servers.

Considerations when using ONTAP in a MetroCluster configuration

When using ONTAP in a MetroCluster configuration, you should be aware of certain considerations for licensing, peering to clusters outside the MetroCluster configuration, performing volume operations, NVFAIL operations, and other ONTAP operations.

Licensing considerations

- Both sites should be licensed for the same site-licensed features.
- All nodes should be licensed for the same node-locked features.

SnapMirror consideration

 SnapMirror SVM disaster recovery is only supported on MetroCluster configurations running versions of ONTAP 9.5 or later.

FlexCache support in a MetroCluster configuration

Beginning with ONTAP 9.7, FlexCache volumes are supported on MetroCluster configurations. You should be aware of requirements for manual repeering after switchover or switchback operations.

SVM repeering after switchover when FlexCache origin and cache are within the same MetroCluster site

After a negotiated or unplanned switchover, any SVM FlexCache peering relationship within the cluster must be manually configured.

For example, SVMs "vs1" (cache) and "vs2" (origin) are on site_A. These SVMs are peered.

After switchover, SVMs "vs1-mc" and "vs2-mc" are activated at the partner site (site_B). They must be manually repeered for FlexCache to work using the vserver peer repeer command.

SVM repeering after switchover or switchback when a FlexCache destination is on a third cluster and in disconnected mode

For FlexCache relationships to a cluster outside of the MetroCluster configuration, the peering must always be manually reconfigured after a switchover if the involved clusters are in disconnected mode during switchover.

For example:

- One end of the FlexCache (cache_1 on vs1) resides on MetroCluster site_A.
- The other end of the FlexCache (origin_1 on vs2) resides on site_C (not in the MetroCluster configuration).

When switchover is triggered, and if site_A and site_C are not connected, you must manually repeer the SVMs on site_B (the switchover cluster) and site_C using the vserver peer repeer command after the switchover.

When switchback is performed, you must again repeer the SVMs on site_A (the original cluster) and site_C.

Related information

FlexCache volumes management with the CLI

FabricPool support in MetroCluster configurations

Beginning with ONTAP 9.7, MetroCluster configurations support FabricPool storage tiers.

For general information on using FabricPools, see Disk and tier (aggregate) management.

Considerations when using FabricPools

• The clusters must have FabricPool licenses with matching capacity limits.

• The clusters must have IPspaces with matching names.

This can be the default IPspace, or an IPspace that an administer has created. This IPspace will be used for FabricPool object store configuration setups.

- For the selected IPspace, each cluster must have an intercluster LIF defined that can reach the external object store.
- SVM migration isn't supported with FabricPool when the source or destination is a MetroCluster cluster.

Learn more about SVM data mobility.

Configuring an aggregate for use in a mirrored FabricPool



Before you configure the aggregate, you must set up object stores as described in Set up object stores for FabricPool in a MetroCluster configuration.

Steps

To configure an aggregate for use in a FabricPool:

- 1. Create the aggregate or select an existing aggregate.
- 2. Mirror the aggregate as a typical mirrored aggregate within the MetroCluster configuration.
- 3. Create the FabricPool mirror with the aggregate, as described in the Disks and aggregates management
 - a. Attach a primary object store.

This object store is physically closer to the cluster.

b. Add a mirror object store.

This object store is physically further distant to the cluster than the primary object store.



It's recommended you maintain at least 20% free space for mirrored aggregates for optimal storage performance and availability. Although the recommendation is 10% for non-mirrored aggregates, the additional 10% of space may be used by the filesystem to absorb incremental changes. Incremental changes increase space utilization for mirrored aggregates due to ONTAP's copy-on-write Snapshot-based architecture. Failure to adhere to these best practices may have a negative impact on performance.

FlexGroup support in MetroCluster configurations

Beginning with ONTAP 9.6 MetroCluster configurations support FlexGroup volumes.

Consistency group support in MetroCluster configurations

Beginning with ONTAP 9.11.1, consistency groups are supported in MetroCluster configurations.

Job schedules in a MetroCluster configuration

In ONTAP 9.3 and later, user-created job schedules are automatically replicated between clusters in a MetroCluster configuration. If you create, modify, or delete a job schedule on a cluster, the same schedule is automatically created on the partner cluster, using Configuration Replication Service (CRS).



System-created schedules are not replicated and you must manually perform the same operation on the partner cluster so that job schedules on both clusters are identical.

Cluster peering from the MetroCluster site to a third cluster

Because the 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 is so that peering can be maintained if a switchover occurs.

The non-MetroCluster cluster must be running ONTAP 8.3 or later. If not, peering is lost if a switchover occurs even if the peering has been configured on both MetroCluster partners.

LDAP client configuration replication in a MetroCluster configuration

An LDAP client configuration created on a storage virtual machine (SVM) on a local cluster is replicated to its partner data SVM on the remote cluster. For example, if the LDAP client configuration is created on the admin SVM on the local cluster, then it is replicated to all the admin data SVMs on the remote cluster. This MetroCluster feature is intentional so that the LDAP client configuration is active on all the partner SVMs on the remote cluster.

Networking and LIF creation guidelines for MetroCluster configurations

You should be aware of how LIFs are created and replicated in a MetroCluster configuration. You must also know about the requirement for consistency so that you can make proper decisions when configuring your network.

Related information

- Network and LIF management
- You should be aware of the requirements for replicating IPspace objects to the partner cluster and for configuring subnets and IPv6 in a MetroCluster configuration.

IPspace object replication and subnet configuration requirements

• You should be aware of the requirements for creating LIFs when configuring your network in a MetroCluster configuration.

Requirements for LIF creation in a MetroCluster configuration

You should be aware of the LIF replication requirements in a MetroCluster configuration. You should also
know how a replicated LIF is placed on a partner cluster, and you should be aware of the issues that occur
when LIF replication or LIF placement fails.

LIF replication and placement requirements and issues

IPspace object replication and subnet configuration requirements

You should be aware of the requirements for replicating IPspace objects to the partner cluster and for configuring subnets and IPv6 in a MetroCluster configuration.

IPspace replication

You must consider the following guidelines while replicating IPspace objects to the partner cluster:

- The IPspace names of the two sites must match.
- IPspace objects must be manually replicated to the partner cluster.

Any storage virtual machines (SVMs) that are created and assigned to an IPspace before the IPspace is replicated will not be replicated to the partner cluster.

Subnet configuration

You must consider the following guidelines while configuring subnets in a MetroCluster configuration:

- Both clusters of the MetroCluster configuration must have a subnet in the same IPspace with the same subnet name, subnet, broadcast domain, and gateway.
- The IP ranges of the two clusters must be different.

In the following example, the IP ranges are different:

cluster_A::> network subnet show						
IPspace: Subnet	Default	Broadcast	<u>Cabaaaaa</u>	Avail/	Denne	
Name	Subnet	Domain	Gateway	TOTAL	Ranges	
subnet1 192.168.2	192.168.2.0/24	Default	192.168.2.1	10/10		
cluster_B	cluster B::> network subnet show					
IPspace:	Default					
Subnet		Broadcast		Avail/		
Name	Subnet	Domain	Gateway	Total	Ranges	
subnet1 192.168.2	192.168.2.0/24 .21-192.168.2.30	Default	192.168.2.1	10/10		

IPv6 configuration

If IPv6 is configured on one site, IPv6 must be configured on the other site as well.

Related information

• You should be aware of the requirements for creating LIFs when configuring your network in a MetroCluster configuration.

Requirements for LIF creation in a MetroCluster configuration

You should be aware of the LIF replication requirements in a MetroCluster configuration. You should also
know how a replicated LIF is placed on a partner cluster, and you should be aware of the issues that occur
when LIF replication or LIF placement fails.
Requirements for LIF creation in a MetroCluster configuration

You should be aware of the requirements for creating LIFs when configuring your network in a MetroCluster configuration.

You must consider the following guidelines when creating LIFs:

- Fibre Channel: You must use stretched VSAN or stretched fabrics
- IP/iSCSI: You must use layer 2 stretched network
- ARP broadcasts: You must enable ARP broadcasts between the two clusters
- Duplicate LIFs: You must not create multiple LIFs with the same IP address (duplicate LIFs) in an IPspace
- NFS and SAN configurations: You must use different storage virtual machines (SVMs) for both the unmirrored and mirrored aggregates
- You should create a subnet object before you create a LIF. A subnet object enables ONTAP to determine failover targets on the destination cluster because it has an associated broadcast domain.

Verify LIF creation

You can confirm the successful creation of a LIF in a MetroCluster configuration by running the metrocluster check lif show command. If you encounter any issues while creating the LIF, you can use the metrocluster check lif repair-placement command to fix the issues.

Related information

• You should be aware of the requirements for replicating IPspace objects to the partner cluster and for configuring subnets and IPv6 in a MetroCluster configuration.

IPspace object replication and subnet configuration requirements

You should be aware of the LIF replication requirements in a MetroCluster configuration. You should also
know how a replicated LIF is placed on a partner cluster, and you should be aware of the issues that occur
when LIF replication or LIF placement fails.

LIF replication and placement requirements and issues

LIF replication and placement requirements and issues

You should be aware of the LIF replication requirements in a MetroCluster configuration. You should also know how a replicated LIF is placed on a partner cluster, and you should be aware of the issues that occur when LIF replication or LIF placement fails.

Replication of LIFs to the partner cluster

When you create a LIF on a cluster in a MetroCluster configuration, the LIF is replicated on the partner cluster. LIFs are not placed on a one-to-one name basis. For availability of LIFs after a switchover operation, the LIF placement process verifies that the ports are able to host the LIF based on reachability and port attribute checks.

The system must meet the following conditions to place the replicated LIFs on the partner cluster:

Condition	LIF type: FC	LIF type: IP/iSCSI
Node identification	ONTAP attempts to place the replicated LIF on the disaster recovery (DR) partner of the node on which it was created. If the DR partner is unavailable, the DR auxiliary partner is used for placement.	ONTAP attempts to place the replicated LIF on the DR partner of the node on which it was created. If the DR partner is unavailable, the DR auxiliary partner is used for placement.
Port identification	ONTAP identifies the connected FC target ports on the DR cluster.	The ports on the DR cluster that are in the same IPspace as the source LIF are selected for a reachability check. If there are no ports in the DR cluster in the same IPspace, the LIF cannot be placed. All of the ports in the DR cluster that are already hosting a LIF in the same IPspace and subnet are automatically marked as reachable; and can be used for placement. These ports are not included in the reachability check.
Reachability check	Reachability is determined by checking for the connectivity of the source fabric WWN on the ports in the DR cluster. If the same fabric is not present at the DR site, the LIF is placed on a random port on the DR partner.	Reachability is determined by the response to an Address Resolution Protocol (ARP) broadcast from each previously identified port on the DR cluster to the source IP address of the LIF to be placed. For reachability checks to succeed, ARP broadcasts must be allowed between the two clusters. Each port that receives a response from the source LIF will be marked as possible for placement.

Port selection	ONTAP categorizes the ports based on attributes such as adapter type and speed, and then selects the ports with matching attributes. If no ports with matching attributes are found, the LIF is placed on a random connected port on the DR partner.	From the ports that are marked as reachable during the reachability check, ONTAP prefers ports that are in the broadcast domain that is associated with the subnet of the LIF. If there are no network ports available on the DR cluster that are in the broadcast domain that is associated with the subnet of the LIF, then ONTAP selects ports that have reachability to the source LIF. If there are no ports with reachability to the source LIF, a port is selected from the broadcast domain that is associated with the subnet of the source LIF, and if no such broadcast domain exists, a random port is selected. ONTAP categorizes the ports based on attributes such as adapter type, interface type, and speed, and then selects the ports with matching attributes.
LIF placement	From the reachable ports, ONTAP selects the least loaded port for placement.	From the selected ports, ONTAP selects the least loaded port for placement.

Placement of replicated LIFs when the DR partner node is down

When an iSCSI or FC 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 operation, the LIFs are not automatically moved to the DR partner. This can lead to LIFs being concentrated on a single node in the partner cluster. During a MetroCluster switchover operation, subsequent attempts to map LUNs belonging to the storage virtual machine (SVM) fail.

You should run the metrocluster check lif show command after a takeover operation or giveback operation to verify that the LIF placement is correct. If errors exist, you can run the metrocluster check lif repair-placement command to resolve the issues.

LIF placement errors

LIF placement errors that are displayed by the metrocluster check lif show command are retained after a switchover operation. If the network interface modify, network interface rename, or network interface delete command is issued for a LIF with a placement error, the error is removed and does not appear in the output of the metrocluster check lif show command.

LIF replication failure

You can also check whether LIF replication was successful by using the metrocluster check lif show command. An EMS message is displayed if LIF replication fails.

You can correct a replication failure by running the metrocluster check lif repair-placement command for any LIF that fails to find a correct port. You should resolve any LIF replication failures as soon as possible to verify the availability of LIF during a MetroCluster switchover operation.



Even if the source SVM is down, LIF placement might proceed normally if there is a LIF belonging to a different SVM in a port with the same IPspace and network in the destination SVM.

LIFs inaccessible after a switchover

If any change is made in the FC switch fabric to which the FC target ports of the source and DR nodes are connected, then the FC LIFs that are placed at the DR partner might become inaccessible to the hosts after a switchover operation.

You should run the metrocluster check lif repair-placement command on the source as well as the DR nodes after a change is made in the FC switch fabric to verify the host connectivity of LIFs. The changes in the switch fabric might result in LIFs getting placed in different target FC ports at the DR partner node.

Related information

• You should be aware of the requirements for replicating IPspace objects to the partner cluster and for configuring subnets and IPv6 in a MetroCluster configuration.

IPspace object replication and subnet configuration requirements

• You should be aware of the requirements for creating LIFs when configuring your network in a MetroCluster configuration.

Requirements for LIF creation in a MetroCluster configuration

Volume creation on a root aggregate

The system does not allow the creation of new volumes on the root aggregate (an aggregate with an HA policy of CFO) of a node in a MetroCluster configuration.

Because of this restriction, root aggregates cannot be added to an SVM using the vserver addaggregates command.

SVM disaster recovery in a MetroCluster configuration

Beginning with ONTAP 9.5, active storage virtual machines (SVMs) in a MetroCluster configuration can be used as sources with the SnapMirror SVM disaster recovery feature. The destination SVM must be on the third cluster outside of the MetroCluster configuration.

Beginning with ONTAP 9.11.1, both sites within a MetroCluster configuration can be the source for an SVM DR relationship with a FAS or AFF destination cluster as shown in the following image.



You should be aware of the following requirements and limitations of using SVMs with SnapMirror disaster recovery:

• Only an active SVM within a MetroCluster configuration can be the source of an SVM disaster recovery relationship.

A source can be a sync-source SVM before switchover or a sync-destination SVM after switchover.

• When a MetroCluster configuration is in a steady state, the MetroCluster sync-destination SVM cannot be the source of an SVM disaster recovery relationship, since the volumes are not online.

The following image shows the SVM disaster recovery behavior in a steady state:



• When the sync-source SVM is the source of an SVM DR relationship, the source SVM DR relationship information is replicated to the MetroCluster partner.

This enables the SVM DR updates to continue after a switchover as shown in the following image:



• During the switchover and switchback processes, replication to the SVM DR destination might fail.

However, after the switchover or switchback process completes, the next SVM DR scheduled updates will succeed.

See the section "Replicating the SVM configuration" in the Data Protection with the CLI for details on configuring an SVM DR relationship.

SVM resynchronization at a disaster recovery site

During resynchronization, the storage virtual machines (SVMs) disaster recovery (DR) source on the MetroCluster configuration is restored from the destination SVM on the non-MetroCluster site.

During resynchronization, the source SVM (cluster_A) temporarily acts as a destination SVM as shown in the

following image:



If an unplanned switchover occurs during resynchronization

Unplanned switchovers that occur during the resynchronization will halt the resynchronization transfer. If an unplanned switchover occurs, the following conditions are true:

- The destination SVM on the MetroCluster site (which was a source SVM prior to resynchronization) remains as a destination SVM. The SVM at the partner cluster will continue to retain its subtype and remain inactive.
- The SnapMirror relationship must be re-created manually with the sync-destination SVM as the destination.
- The SnapMirror relationship does not appear in the SnapMirror show output after a switchover at the survivor site unless a SnapMirror create operation is executed.

Performing switchback after an unplanned switchover during resynchronization

To successfully perform the switchback process, the resynchronization relationship must be broken and deleted. Switchback is not permitted if there are any SnapMirror DR destination SVMs in the MetroCluster configuration or if the cluster has an SVM of subtype "dp-destination".

Output for the "storage aggregate plex show" command is indeterminate after a MetroCluster switchover

When you run the storage aggregate plex show command after a MetroCluster switchover, the status of plex0 of the switched over root aggregate is indeterminate and is displayed as "failed". During this time, the switched over root is not updated. The actual status of this plex can only be determined after the MetroCluster healing phase.

Modifying volumes to set the NVFAIL flag in case of switchover

You can modify a volume so that the NVFAIL flag is set on the volume in the event of a MetroCluster switchover. 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.

About this task



In ONTAP versions earlier than 9.0, the NVFAIL flag is used for each switchover. In ONTAP 9.0 and later versions, the unplanned switchover (USO) is used.

Step

1. Enable MetroCluster configuration to trigger NVFAIL on switchover by setting the vol -dr-force -nvfail parameter to "on":

vol modify -vserver vserver-name -volume volume-name -dr-force-nvfail on

Where to find additional information

You can learn more about MetroCluster configuration and operation.

MetroCluster and miscellaneous information

Information	Subject
ONTAP 9 Documentation	All MetroCluster information
NetApp MetroCluster Solution Architecture and Design, TR-4705	 A technical overview of the MetroCluster FC configuration and operation. Best practices for MetroCluster FC configuration.
MetroCluster IP Solution Architecture and Design, TR- 4689	 A technical overview of the MetroCluster IP configuration and operation. Best practices for a MetroCluster IP configuration.
Stretch MetroCluster installation and configuration	 Stretch MetroCluster architecture Cabling the configuration Configuring the FC-to-SAS bridges Configuring the MetroCluster in ONTAP
MetroCluster IP installation and configuration: Differences among the ONTAP MetroCluster configurations	 MetroCluster IP architecture Cabling the configuration Configuring the MetroCluster in ONTAP
MetroCluster management and disaster recovery	 Understanding the MetroCluster configuration Switchover, healing and switchback Disaster recovery

Maintain the MetroCluster components	 Guidelines for maintenance in a MetroCluster FC configuration
	 Hardware replacement or upgrade and firmware upgrade procedures for FC-to-SAS bridges and FC switches
	 Hot-adding a disk shelf in a fabric-attached or stretch MetroCluster FC configuration
	 Hot-removing a disk shelf in a fabric-attached or stretch MetroCluster FC configuration
	 Replacing hardware at a disaster site in a fabric- attached or stretch MetroCluster FC configuration
	 Expanding a two-node fabric-attached or stretch MetroCluster FC configuration to a four-node MetroCluster configuration.
	 Expanding a four-node fabric-attached or stretch MetroCluster FC configuration to an eight-node MetroCluster FC configuration.
Transition from MetroCluster FC to MetroCluster IP	 Upgrading or refreshing a MetroCluster configuration
MetroCluster Upgrade and Expansion Guide	 Transitioning from a MetroCluster FC configuration to a MetroCluster IP configuration
	 Expanding a MetroCluster configuration by adding additional nodes
MetroCluster Tiebreaker Software installation and configuration	 Monitoring the MetroCluster configuration with the MetroCluster Tiebreaker software
Active IQ Digital Advisor documentation	 Monitoring the MetroCluster configuration and performance
NetApp Documentation: Product Guides and Resources	
Copy-based transition	 Transitioning data from 7-Mode storage systems to clustered storage systems
ONTAP concepts	 How mirrored aggregates work

Copyright information

Copyright © 2025 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.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at http://www.netapp.com/TM are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.