



Plan and install a MetroCluster configuration with array LUNs

ONTAP MetroCluster

NetApp
May 24, 2022

Table of Contents

- Plan and install a MetroCluster configuration with array LUNs 1
 - Planning for a MetroCluster configuration with array LUNs 1
 - Supported MetroCluster configuration with array LUNs 1
 - Requirements for a MetroCluster configuration with array LUNs 1
- Install and cable the MetroCluster components in a configuration with array LUNs 3
- Switch zoning in a MetroCluster configuration with array LUNs 32
- Set up ONTAP in a MetroCluster configuration with array LUNs 39
- Implement a MetroCluster configuration with both disks and array LUNs 53

Plan and install a MetroCluster configuration with array LUNs

Planning for a MetroCluster configuration with array LUNs

Creating a detailed plan for your MetroCluster configuration helps you understand the unique requirements for a MetroCluster configuration that uses LUNs on storage arrays. Installing a MetroCluster configuration involves connecting and configuring a number of devices, which might be done by different people. Therefore, the plan also helps you communicate with other people involved in the installation.

Supported MetroCluster configuration with array LUNs

You can set up either a MetroCluster configuration with array LUNs. Both stretch and fabric-attached configurations are supported. AFF systems are not supported with array LUNs.

The features supported on the MetroCluster configurations vary with the configuration types. The following table lists the features supported on the different types of MetroCluster configurations with array LUNs:

Feature	Fabric-attached configurations			Stretch configurations
	Eight-node	Four-node	Two-node	Two-node
Number of controllers	Eight	Four	Two	Two
Uses an FC switch storage fabric	Yes	Yes	Yes	Yes
Uses FC-to-SAS bridges	Yes	Yes	Yes	Yes
Supports local HA	Yes	Yes	No	No
Supports automatic switchover	Yes	Yes	Yes	Yes

Related information

[Differences between the ONTAP MetroCluster configurations](#)

Requirements for a MetroCluster configuration with array LUNs

The ONTAP systems, storage arrays, and FC switches used in MetroCluster configurations must meet the requirements for such types of configurations. In addition, you must also consider the SyncMirror requirements for MetroCluster configurations with

array LUNs.

Requirements for ONTAP systems

- The ONTAP systems must be identified as supported for MetroCluster configurations.

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.



You must refer to the alert details associated with any configuration that you select in the Interoperability Matrix.

- All the ONTAP systems in a MetroCluster configuration must be of the same model.
- FC-VI adapters must be installed in the appropriate slots for each ONTAP system, depending on the model.

[NetApp Hardware Universe](#)

Requirements for storage arrays

- The storage arrays must be identified as supported for MetroCluster configurations.

[NetApp Interoperability Matrix Tool](#)

- The storage arrays in the MetroCluster configuration must be symmetric:
 - The two storage arrays must be from the same supported vendor family and have the same firmware version installed.

[FlexArray virtualization implementation for NetApp E-Series storage](#)

[FlexArray virtualization implementation for third-party storage](#)

- Disk types (for example, SATA, SSD, or SAS) used for mirrored storage must be the same on both storage arrays.
- The parameters for configuring storage arrays, such as RAID type and tiering, must be the same across both sites.

Requirements for FC switches

- The switches and switch firmware must be identified as supported for MetroCluster configurations.

[NetApp Interoperability Matrix Tool](#)

- Each fabric must have two FC switches.
- Each ONTAP system must be connected to storage using redundant components so that there is redundancy in case of device and path failures.
- AFF A900 and FAS9000 storage systems support up to eight ISLs per fabric. Other storage system models support up to four ISLs per fabric.

The switches must use the MetroCluster basic switch configuration, ISL settings, and FC-VI configurations.

[Configuring the Cisco or Brocade FC switches manually](#)

SyncMirror requirements

- SyncMirror is required for a MetroCluster configuration.
- Two separate storage arrays, one at each site, are required for the mirrored storage.
- Two sets of array LUNs are required.

One set is required for the aggregate on the local storage array (pool0) and another set is required at the remote storage array for the mirror of the aggregate (the other plex of the aggregate, pool1).

The array LUNs must be of the same size for mirroring the aggregate.

- Unmirrored aggregates are also supported in the MetroCluster configuration.

They are not protected in the event of a site disaster.

Install and cable the MetroCluster components in a configuration with array LUNs

Racking the hardware components in a MetroCluster configuration with array LUNs

You must ensure that the hardware components required to set up a MetroCluster configuration with array LUNs are properly racked.

About this task

You must perform this task at both the MetroCluster sites.

Steps

1. Plan the positioning of the MetroCluster components.

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

2. Properly ground yourself.
3. Install the storage controllers in the rack or cabinet.



AFF systems are not supported with array LUNs.

[Installation procedures for your AFF or FAS system](#)

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

Preparing a storage array for use with ONTAP systems

Before you can begin setting up ONTAP systems in a MetroCluster configuration with array LUNs, the storage array administrator must prepare the storage for use with

ONTAP.

Before you begin

The storage arrays, firmware, and switches that you plan to use in the configuration must be supported by the specific ONTAP version.

- [NetApp Interoperability \(IMT\)](#)

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.

- [NetApp Hardware Universe](#)

About this task

You must coordinate with the storage array administrator to perform this task on the storage array.

Steps

1. Create LUNs on the storage array depending on the number of nodes in the MetroCluster configuration.

Each node in the MetroCluster configuration requires array LUNs for the root aggregate, data aggregate, and spares.
2. Configure parameters on the storage array that are required to work with ONTAP.
 - [FlexArray virtualization implementation for third-party storage](#)
 - [FlexArray virtualization implementation for NetApp E-Series storage](#)

Switch ports required for a MetroCluster configuration with array LUNs

When you are connecting ONTAP systems to FC switches for setting up a MetroCluster configuration with array LUNs, you must connect FC-VI and HBA ports from each controller to specific switch ports.

If you are using both array LUNs and disks in the MetroCluster configuration, you must ensure that the controller ports are connected to the switch ports recommended for configuration with disks, and then use the remaining ports for configuration with array LUNs.

The following table lists the specific FC switch ports to which you must connect the different controller ports in an eight-node MetroCluster configuration with array LUNs.

Overall cabling guidelines with array LUNs

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.
- FAS8200 storage systems can be ordered 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.
- 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.

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.

Brocade port usage for controllers in a MetroCluster configuration

The following tables show port usage on Brocade switches. The tables show the maximum supported configuration, with eight controller modules in two DR groups. For smaller configurations, ignore the rows for the additional controller modules. Note that eight ISLs are supported on the Brocade 6510 and G620 switches.



Port usage for the Brocade 6505 switch 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 following table shows the cabling for the first DR group:

		Brocade 6520, 6510, 6505, G620, G610, or 7840 switch	
Component	Port	Switch 1	Switch 2
controller_x_1	FC-VI port a	0	
	FC-VI port b	-	0
	FC-VI port c	1	-
	FC-VI port d	-	1
	HBA port a	2	-
	HBA port b	-	2
	HBA port c	3	-
	HBA port d	-	3

controller_x_2	FC-VI port a	4	-
	FC-VI port b	-	4
	FC-VI port c	5	-
	FC-VI port d	-	5
	HBA port a	6	-
	HBA port b	-	6
	HBA port c	7	-
	HBA port d	-	7

The following table shows the cabling for the second DR group:

		Brocade 6510		Brocade 6520		Brocade G620	
Component	Port	Switch 1	Switch 2	Switch 1	Switch 2	Switch 1	Switch 2
controller_x_3	FC-VI port a	24	-	48	-	18	-
	FC-VI port b	-	24	-	48	-	18
	FC-VI port c	25	-	49	-	19	-
	FC-VI port d	-	25	-	49	-	19
	HBA port a	26	-	50	-	24	-
	HBA port b	-	26	-	50	-	24
	HBA port c	27	-	51	-	25	-
	HBA port d	-	27	-	51	-	25

controller_x_4	FC-VI port a	28	-	52	-	22	-
	FC-VI port b	-	28	-	52	-	22
	FC-VI port c	29	-	53	-	23	-
	FC-VI port d	-	29	-	53	-	23
	HBA port a	30	-	54	-	28	-
	HBA port b	-	30	-	54	-	28
	HBA port c	31	-	55	-	29	-
	HBA port d	-	31	-	55	-	29
ISLs							
ISL 1		40	40	23	23	40	40
ISL 2		41	41	47	47	41	41
ISL 3		42	42	71	71	42	42
ISL 4		43	43	95	95	43	43
ISL 5		44	44	Not supported		44	44
ISL 6		45	45			45	45
ISL 7		46	46			46	46
ISL 8		47	47			47	47

Cisco port usage for controllers in a MetroCluster configuration running ONTAP 9.4 or later

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

Cisco 9396S port usage

Cisco 9396S			
Component	Port	Switch 1	Switch 2

controller_x_1	FC-VI port a	1	-
	FC-VI port b	-	1
	FC-VI port c	2	-
	FC-VI port d	-	2
	HBA port a	3	-
	HBA port b	-	3
	HBA port c	4	-
	HBA port d	-	4
controller_x_2	FC-VI port a	5	-
	FC-VI port b	-	5
	FC-VI port c	6	-
	FC-VI port d	-	6
	HBA port a	7	-
	HBA port b	-	7
	HBA port c	8	-
	HBA port d	-	8

controller_x_3	FC-VI port a	49	
	FC-VI port b	-	49
	FC-VI port c	50	
	FC-VI port d	-	50
	HBA port a	51	
	HBA port b	-	51
	HBA port c	52	
	HBA port d	-	52
controller_x_4	FC-VI port a	53	-
	FC-VI port b	-	53
	FC-VI port c	54	-
	FC-VI port d	-	54
	HBA port a	55	-
	HBA port b	-	55
	HBA port c	56	-
	HBA port d	-	56

Cisco 9148S port usage

Cisco 9148S			
Component	Port	Switch 1	Switch 2

controller_x_1	FC-VI port a	1	-
	FC-VI port b	-	1
	FC-VI port c	2	-
	FC-VI port d	-	2
	HBA port a	3	-
	HBA port b	-	3
	HBA port c	4	-
	HBA port d	-	4
controller_x_2	FC-VI port a	5	-
	FC-VI port b	-	5
	FC-VI port c	6	-
	FC-VI port d	-	6
	HBA port a	7	-
	HBA port b	-	7
	HBA port c	8	-
	HBA port d	-	8

controller_x_3	FC-VI port a	25	
	FC-VI port b	-	25
	FC-VI port c	26	-
	FC-VI port d	-	26
	HBA port a	27	-
	HBA port b	-	27
	HBA port c	28	-
	HBA port d	-	28
controller_x_4	FC-VI port a	29	-
	FC-VI port b	-	29
	FC-VI port c	30	-
	FC-VI port d	-	30
	HBA port a	31	-
	HBA port b	-	31
	HBA port c	32	-
	HBA port d	-	32

Cisco 9132T port usage

Cisco 9132T			
MDS module 1			
Component	Port	Switch 1	Switch 2

controller_x_1	FC-VI port a	1	-
	FC-VI port b	-	1
	FC-VI port c	2	-
	FC-VI port d	-	2
	HBA port a	3	-
	HBA port b	-	3
	HBA port c	4	-
	HBA port d	-	4
controller_x_2	FC-VI port a	5	-
	FC-VI port b	-	5
	FC-VI port c	6	-
	FC-VI port d	-	6
	HBA port a	7	-
	HBA port b	-	7
	HBA port c	8	-
	HBA port d	-	8
MDS module 2			
Component	Port	Switch 1	Switch 2

controller_x_3	FC-VI port a	1	-
	FC-VI port b	-	1
	FC-VI port c	2	-
	FC-VI port d	-	2
	HBA port a	3	-
	HBA port b	-	3
	HBA port c	4	-
	HBA port d	-	4
controller_x_4	FC-VI port a	5	-
	FC-VI port b	-	5
	FC-VI port c	6	-
	FC-VI port d	-	6
	HBA port a	7	-
	HBA port b	-	7
	HBA port c	8	-
	HBA port d	-	8

Cisco 9250 port usage



The following table shows systems with two FC-VI ports. AFF A700 and FAS9000 systems have four FC-VI ports (a, b, c, and d). If using an AFF A700 or FAS9000 system, the port assignments move along by one position. For example, FC-VI ports c and d go to switch port 2 and HBA ports a and b go to switch port 3.

Cisco 9250i			
The Cisco 9250i switch is not supported for eight-node MetroCluster configurations.			
Component	Port	Switch 1	Switch 2

controller_x_1	FC-VI port a	1	-
	FC-VI port b	-	1
	HBA port a	2	-
	HBA port b	-	2
	HBA port c	3	-
	HBA port d	-	3
controller_x_2	FC-VI port a	4	-
	FC-VI port b	-	4
	HBA port a	5	-
	HBA port b	-	5
	HBA port c	6	-
	HBA port d	-	6
controller_x_3	FC-VI port a	7	-
	FC-VI port b	-	7
	HBA port a	8	-
	HBA port b	-	8
	HBA port c	9	-
	HBA port d	-	9

controller_x_4	FC-VI port a	10	-
	FC-VI port b	-	10
	HBA port a	11	-
	HBA port b	-	11
	HBA port c	13	-
	HBA port d	-	13

Shared initiator and shared target support for MetroCluster configuration with array LUNs

Being able to share a given FC initiator port or target ports is useful for organizations that want to minimize the number of initiator or target ports used. For example, an organization that expects low I/O usage over an FC initiator port or target ports might prefer to share FC initiator port or target ports instead of dedicating each FC initiator port to a single target port.

However sharing of initiator or target ports can adversely affect performance.

[How to support Shared Initiator and Shared Target configuration with Array LUNs in a MetroCluster environment](#)

Cable the FC-VI and HBA ports in a MetroCluster configuration with array LUNs

Cabling the FC-VI and HBA ports in a two-node fabric-attached MetroCluster configuration with array LUNs

If you are setting up a two-node fabric-attached MetroCluster configuration with array LUNs, you must cable the FC-VI ports and the HBA ports to the switch ports.

About this task

- You must repeat this task for each controller at both of the MetroCluster sites.
- If you plan to use disks in addition to array LUNs in your MetroCluster configuration, you must use the HBA ports and switch ports specified for configuration with disks.
 - [Port assignments for FC switches when using ONTAP 9.1 and later](#)
 - [Port assignments for FC switches when using ONTAP 9.0](#)


Steps

1. Cable the FC-VI ports from the controller to alternate switch ports.
2. Perform the controller-to-switch cabling at both of the MetroCluster sites.

You must ensure redundancy in connections from the controller to the switches. Therefore, for each controller at a site, you must ensure that both of the HBA ports in the same port pair are connected to alternate FC switches.

The following example shows the connections between the HBA ports on Controller A and ports on

FC_switch_A_1 and FC_switch_A_2:



The following table lists the connections between the HBA ports and the FC switch ports in the illustration:

HBA ports	Switch ports
Port pair	
Port a	FC_switch_A_1, Port 2
Port d	FC_switch_A_2, Port 3
Port pair	
Port b	FC_switch_A_2, Port 2
Port c	FC_switch_A_1, Port 3

After you finish

You should cable the ISLs between the FC switches across the MetroCluster sites.

Cabling the FC-VI and HBA ports in a four-node fabric-attached MetroCluster configuration with array LUNs

If you are setting up a four-node fabric-attached MetroCluster configuration with array LUNs, you must cable the FC-VI ports and the HBA ports to the switch ports.

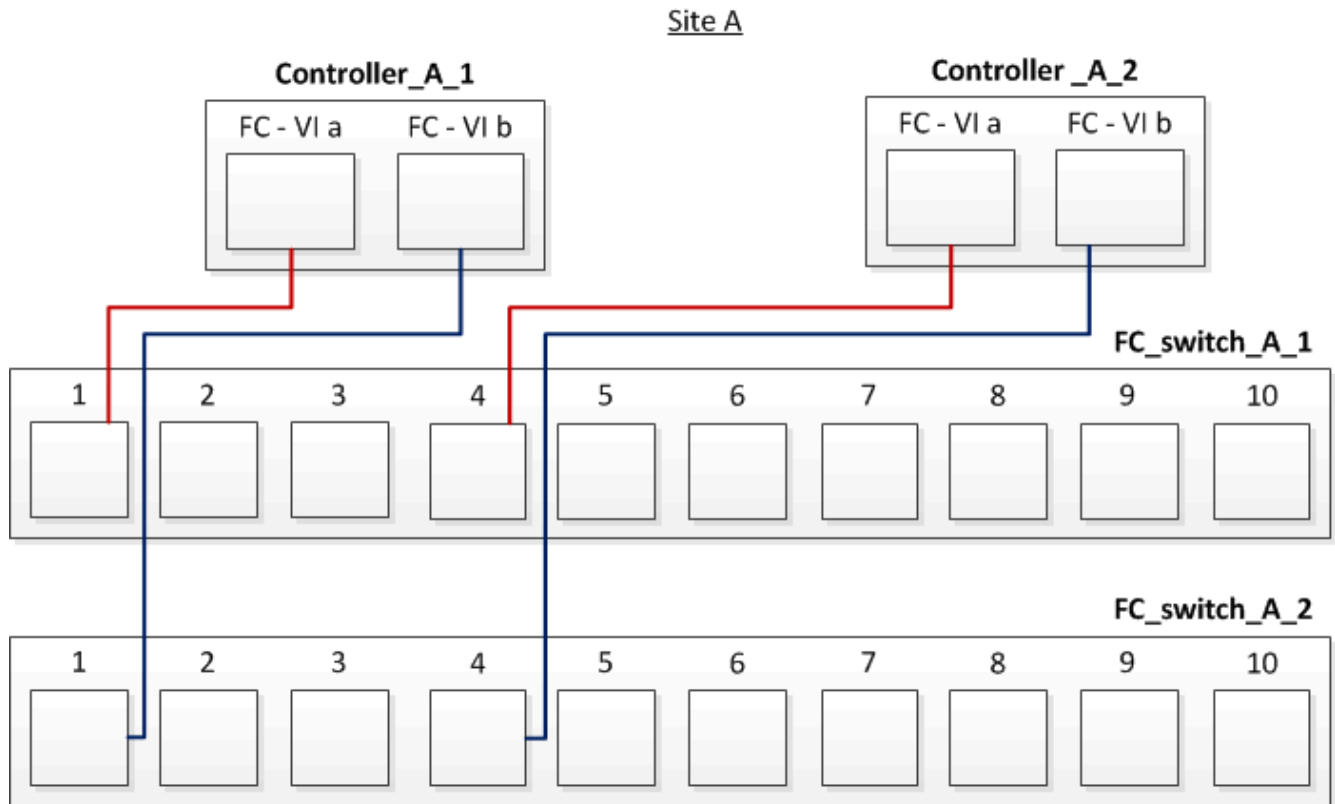
About this task

- You must repeat this task for each controller at both of the MetroCluster sites.
- If you plan to use disks in addition to array LUNs in your MetroCluster configuration, you must use the HBA ports and switch ports specified for configuration with disks.
 - [Port assignments for FC switches when using ONTAP 9.1 and later](#)
 - [Port assignments for FC switches when using ONTAP 9.0](#)

Steps

1. Cable the FC-VI ports from each controller to the ports on alternate FC switches.

The following example shows the connections between the FC-VI ports and switch ports at Site A:




2. Perform the controller-to-switch cabling at both of the MetroCluster sites.

You must ensure redundancy in connections from the controller to the switches. Therefore, for each controller at a site, you must ensure that both of the HBA ports in the same port pair are connected to alternate FC switches.

The following example shows the connections between the HBA ports and switch ports at Site A:

Site A



The following table lists the connections between the HBA ports on controller_A_1 and the FC switch ports in the illustration:

HBA ports	Switch ports
Port pair	
Port a	FC_switch_A_1, Port 2
Port d	FC_switch_A_2, Port 3
Port pair	
Port b	FC_switch_A_2, Port 2
Port c	FC_switch_A_1, Port 3

The following table lists the connections between the HBA ports on controller_A_2 and the FC switch ports in the illustration:

HBA ports	Switch ports
Port pair	
Port a	FC_switch_A_1, Port 5
Port d	FC_switch_A_2, Port 6

Port pair	
Port b	FC_switch_A_2, Port 5
Port c	FC_switch_A_1, Port 6

After you finish

You should cable the ISLs between the FC switches across the MetroCluster sites.

Related information

When you are connecting ONTAP systems to FC switches for setting up a MetroCluster configuration with array LUNs, you must connect FC-VI and HBA ports from each controller to specific switch ports.

[Switch ports required for a MetroCluster configuration with array LUNs](#)

Cabling the FC-VI and HBA ports in an eight-node fabric-attached MetroCluster configuration with array LUNs

If you are setting up an eight-node fabric-attached MetroCluster configuration with array LUNs, you must cable the FC-VI ports and the HBA ports to the switch ports.

About this task

- You must repeat this task for each controller at both of the MetroCluster sites.
- If you plan to use disks in addition to array LUNs in your MetroCluster configuration, you must use the HBA ports and switch ports specified for configuration with disks.
 - [Port assignments for FC switches when using ONTAP 9.1 and later](#)
 - [Port assignments for FC switches when using ONTAP 9.0](#)

Step

1. Cable the FC-VI ports and HBA ports from each controller to the ports on alternate FC switches. Refer to the following tables:

Cabling configurations for FibreBridge 7500N or 7600N using both FC ports

Configurations using FibreBridge 7500N or 7600N using both FC ports (FC1 and FC2)				
MetroCluster 1 or DR Group 1				
Component	Port	Brocade switch models 6505, 6510, 6520, 7810, 7840, G610, G620, G620-1, G630, G630-1, and DCX 8510-8		
		Connects to FC_switch...	Connects to switch port...	Connects to switch port...

controller_x_1		FC-VI port a	1	0	0
		FC-VI port b	2	0	0
		FC-VI port c	1	1	1
		FC-VI port d	2	1	1
		HBA port a	1	2	8
		HBA port b	2	2	8
		HBA port c	1	3	9
		HBA port d	2	3	9
controller_x_2		FC-VI port a	1	4	4
		FC-VI port b	2	4	4
		FC-VI port c	1	5	5
		FC-VI port d	2	5	5
		HBA port a	1	6	12
		HBA port b	2	6	12
		HBA port c	1	7	13
		HBA port d	2	7	13
Stack 1	bridge_x_1a	FC1	1	8	10
		FC2	2	8	10
	bridge_x_1B	FC1	1	9	11
		FC2	2	9	11

Stack 2	bridge_x_2a	FC1	1	10	14
		FC2	2	10	14
	bridge_x_2B	FC1	1	11	15
		FC2	2	11	15
Stack 3	bridge_x_3a	FC1	1	12*	16
		FC2	2	12*	16
	bridge_x_3B	FC1	1	13*	17
		FC2	2	13*	17
Stack y	bridge_x_ya	FC1	1	14*	20
		FC2	2	14*	20
	bridge_x_yb	FC1	1	15*	21
		FC2	2	15*	21

*** Ports 12 through 15 are reserved for the second MetroCluster or DR group on the Brocade 7840 switch.**

Note: Additional bridges can be cabled to ports 16, 17, 20 and 21 in G620, G630, G620-1 and G630-1 switches.

After you finish

You should cable the ISLs between the FC switches across the MetroCluster sites.

Cabling configurations for Cisco 9250i

Cisco 9250i*			
Component	Port	Switch 1	Switch 2

controller_x_1	FC-VI port a	1	-
	FC-VI port b	-	1
	HBA port a	2	-
	HBA port b	-	2
	HBA port c	3	-
	HBA port d	-	3
controller_x_2	FC-VI port a	4	-
	FC-VI port b	-	4
	HBA port a	5	-
	HBA port b	-	5
	HBA port c	6	-
	HBA port d	-	6
controller_x_3	FC-VI port a	7	-
	FC-VI port b	-	7
	HBA port a	8	-
	HBA port b	-	8
	HBA port c	9	-
	HBA port d	-	9

controller_x_4	FC-VI port a	10	-
	FC-VI port b	-	10
	HBA port a	11	-
	HBA port b	-	11
	HBA port c	13	-
	HBA port d	-	13

After you finish

You should cable the ISLs between the FC switches across the MetroCluster sites.

Cabling the ISLs in a MetroCluster configuration with array LUNs

You must connect the FC switches across the sites through Inter-Switch Links (ISLs) to form switch fabrics in your MetroCluster configuration with array LUNs.

Steps

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

The switch port numbers that you can use for the FC ISLs are as follows:

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 8510-8	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
	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 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 1	ge0 (40-Gbps) or ge2 (10-Gbps)
	ISL port 2	ge1 (40-Gbps) or ge3 (10-Gbps)
	ISL port 3	ge10 (10-Gbps)
	ISL port 4	ge11 (10-Gbps)
Brocade G610	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
	ISL port 8	47
Switch mode I	ISL port	Switch port
Cisco 9396S	ISL 1	44
	ISL 2	48
	ISL 3	92
	ISL 4	96
Cisco 9250i with 24 port license	ISL 1	12
	ISL 2	16
	ISL 3	20
	ISL 4	24
Cisco 9148S	ISL 1	20
	ISL 2	24
	ISL 3	44
	ISL 4	48

Cisco 9132T	ISL 1	MDS module 1 port 13
	ISL 2	MDS module 1 port 14
	ISL 3	MDS module 1 port 15
	ISL 4	MDS module 1 port 16
* The Cisco 9250i switch uses the FCIP ports for the ISL. There are certain limitations and procedures for using the FCIP ports.		
Ports 40 through 48 are 10 GbE ports and are not used in the MetroCluster configuration.		

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

In eight- 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

[AFF and FAS Documentation Center](#)

[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.

[Considerations for configuring 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.

[AFF and FAS Documentation Center](#)

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.

[AFF and FAS Documentation Center](#)

Cable storage arrays to FC switches in a MetroCluster configuration

Cabling storage arrays to FC switches in a MetroCluster configuration

You must connect storage arrays to FC switches so that the ONTAP systems in the MetroCluster configuration can access a specific array LUN through at least two paths.

Before you begin

- The storage arrays must be set up to present array LUNs to ONTAP.
- The ONTAP controllers must be connected to the FC switches.
- The ISLs must be cabled between the FC switches across the MetroCluster sites.
- You must repeat this task for each storage array at both of the MetroCluster sites.
- You must connect the controllers in a MetroCluster configuration to the storage arrays through FC switches.

Steps

1. Connect the storage array ports to FC switch ports.

At each site, connect the redundant port pairs in the storage array to FC switches on alternate fabrics. This provides redundancy in the paths for accessing the array LUNs.

Related information

- [Configuring switch zoning enables you to define which array LUNs can be viewed by a specific ONTAP system in the MetroCluster configuration.](#)

[Switch zoning in a MetroCluster configuration with array LUNs](#)

- In a MetroCluster configuration with array LUNs, you must connect the storage array ports that form a redundant port pair to alternate FC switches.

[Example of cabling storage array ports to FC switches in a two-node MetroCluster configuration](#)


[Example of cabling storage array ports to FC switches in a four-node MetroCluster configuration](#)

[Example of cabling storage array ports to FC switches in an eight-node MetroCluster configuration](#)

Example of cabling storage array ports to FC switches in a two-node MetroCluster configuration

In a MetroCluster configuration with array LUNs, you must connect the storage array ports that form a redundant port pair to alternate FC switches.

The following illustration shows the connections between storage arrays and FC switches in a two-node fabric-attached MetroCluster configuration with array LUNs:



The connections between storage array ports and FC switch ports are similar for both stretch and fabric-attached variants of two-node MetroCluster configurations with array LUNs.



If you plan to use disks in addition to array LUNs in your MetroCluster configuration, you must use the switch ports specified for the configuration with disks.

[Port assignments for FC switches when using ONTAP 9.1 and later](#)

In the illustration, the redundant array port pairs for both the sites are as follows:

- Storage array at Site A:
 - Ports 1A and 2A
 - Ports 1B and 2B
- Storage array at Site B:
 - Ports 1A' and 2A'
 - Ports 1B' and 2B'

FC_switch_A_1 at Site A and FC_switch_B_1 at Site B are connected to form fabric_1. Similarly, FC_switch_A_2 at Site A and FC_switch_B_2 are connected to form fabric_2.

The following table lists the connections between the storage array ports and the FC switches for the example MetroCluster illustration:

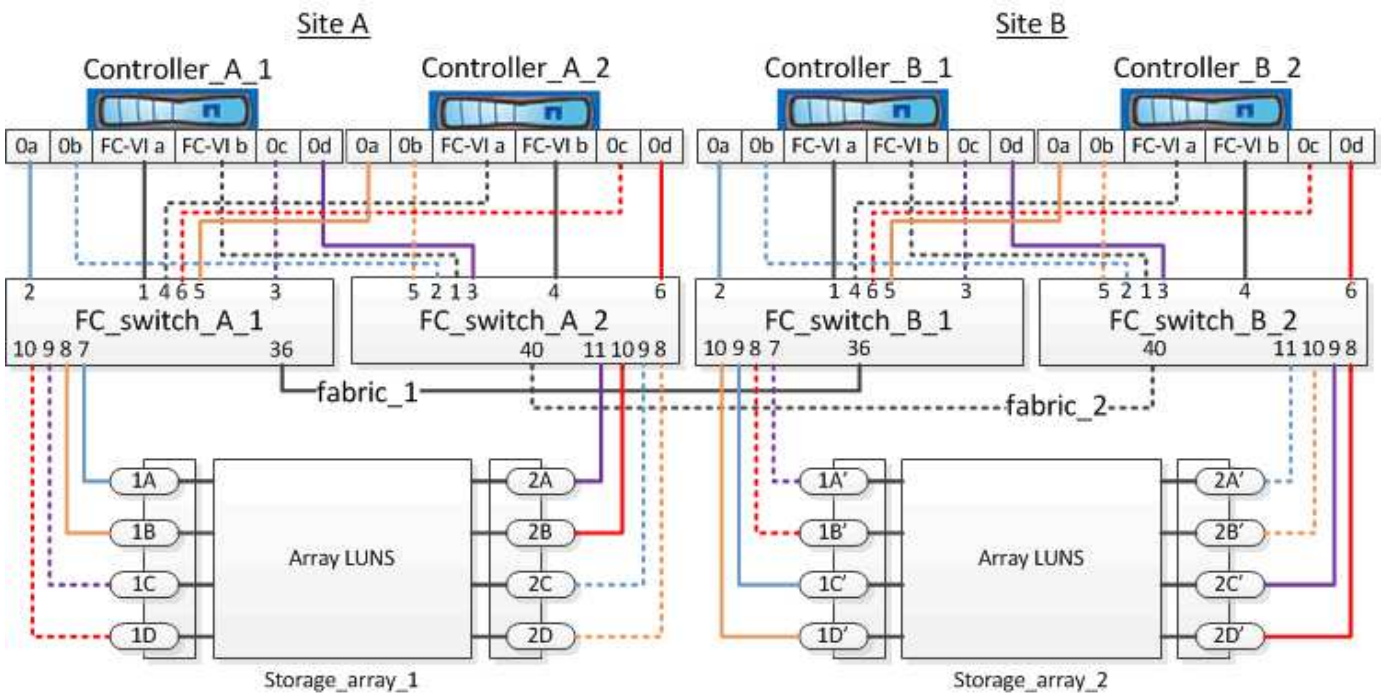
Array LUN ports	FC switch ports	Switch fabrics
Site A		
1A	FC_switch_A_1, Port 9	fabric_1
2A	FC_switch_A_2, Port 10	fabric_2
1B	FC_switch_A_1, Port 10	fabric_1

2B	FC_switch_A_2, Port 9	fabric_2
Site B		
1A'	FC_switch_B_1, Port 9	fabric_1
2A'	FC_switch_B_2, Port 10	fabric_2
1B'	FC_switch_B_1, Port 10	fabric_1
2B'	FC_switch_B_2, Port 9	fabric_2

Example of cabling storage array ports to FC switches in a four-node MetroCluster configuration

In a MetroCluster configuration with array LUNs, you must connect the storage array ports that form a redundant port pair to alternate FC switches.

The following reference illustration shows the connections between storage arrays and FC switches in a four-node MetroCluster configuration with array LUNs:



If you plan to use disks in addition to array LUNs in your MetroCluster configuration, you must use the switch ports specified for the configuration with disks.

[Port assignments for FC switches when using ONTAP 9.1 and later](#)

In the illustration, the redundant array port pairs for both the sites are as follows:

- Storage array at Site A:
 - Ports 1A and 2A

- Ports 1B and 2B
- Ports 1C and 2C
- Ports 1D and 2D
- Storage array at Site B:
 - Ports 1A' and 2A'
 - Ports 1B' and 2B'
 - Ports 1C' and 2C'
 - Ports 1D' and 2D'

FC_switch_A_1 at Site A and FC_switch_B_1 at Site B are connected to form fabric_1. Similarly, FC_switch_A_2 at Site A and FC_switch_B_2 are connected to form fabric_2.

The following table lists the connections between the storage array ports and the FC switches for the MetroCluster illustration:

Array LUN ports	FC switch ports	Switch fabrics
Site A		
1A	FC_switch_A_1, Port 7	fabric_1
2A	FC_switch_A_2, Port 11	fabric_2
1B	FC_switch_A_1, Port 8	fabric_1
2B	FC_switch_A_2, Port 10	fabric_2
1C	FC_switch_A_1, Port 9	fabric_1
2C	FC_switch_A_2, Port 9	fabric_2
1D	FC_switch_A_1, Port 10	fabric_1
2D	FC_switch_A_2, Port 8	fabric_2
Site B		
1A'	FC_switch_B_1, Port 7	fabric_1
2A'	FC_switch_B_2, Port 11	fabric_2
1B'	FC_switch_B_1, Port 8	fabric_1
2B'	FC_switch_B_2, Port 10	fabric_2
1C'	FC_switch_B_1, Port 9	fabric_1

2C'	FC_switch_B_2, Port 9	fabric_2
1D'	FC_switch_B_1, Port 10	fabric_1
2D'	FC_switch_B_2, Port 8	fabric_2

Example of cabling storage array ports to FC switches in an eight-node MetroCluster configuration

In a MetroCluster configuration with array LUNs, you must connect the storage array ports that form a redundant port pair to alternate FC switches.

An eight-node MetroCluster configuration consists of two four-node DR groups. The first DR group consists of the following nodes:

- controller_A_1
- controller_A_2
- controller_B_1
- controller_B_2

The second DR group consists of the following nodes:

- controller_A_3
- controller_A_4
- controller_B_3
- controller_B_4

To cable the array ports for the first DR group, you can use the cabling examples for a four-node MetroCluster configuration for the first DR group.

[Example of cabling storage array ports to FC switches in a four-node MetroCluster configuration](#)

To cable the array ports for the second DR group, follow the same examples and extrapolate for the FC-VI ports and FC initiator ports belonging to the controllers in the second DR group.

Switch zoning in a MetroCluster configuration with array LUNs

Requirements for switch zoning in a MetroCluster configuration with array LUNs

When using switch zoning in a MetroCluster configuration with array LUNs, you must ensure that certain basic requirements are followed.

The requirements for switch zoning in a MetroCluster configuration with array LUNs are as follows:

- The MetroCluster configuration must follow the single-initiator to single-target zoning scheme.

Single-initiator to single-target zoning limits each zone to a single FC initiator port and a single target port.

- The FC-VI ports must be zoned end-to-end across the fabric.
- Sharing of multiple initiator ports with a single target port can cause performance issues.

Similarly, sharing of multiple target ports with a single initiator port can cause performance issues.

- You must have performed a basic configuration of the FC switches used in the MetroCluster configuration.

[Configuring the Cisco or Brocade FC switches manually](#)

Shared initiator and shared target support for MetroCluster configuration with array LUNs

Being able to share a given FC initiator port or target ports is useful for organizations that want to minimize the number of initiator or target ports used. For example, an organization that expects low I/O usage over an FC initiator port or target ports might prefer to share FC initiator port or target ports instead of dedicating each FC initiator port to a single target port.

However sharing of initiator or target ports can adversely affect performance.

Related information

[How to support Shared Initiator and Shared Target configuration with Array LUNs in a MetroCluster environment](#)

- Switch zoning defines paths between connected nodes. Configuring the zoning enables you to define which array LUNs can be viewed by specific ONTAP systems.

[Example of switch zoning in a two-node MetroCluster configuration with array LUNs](#)

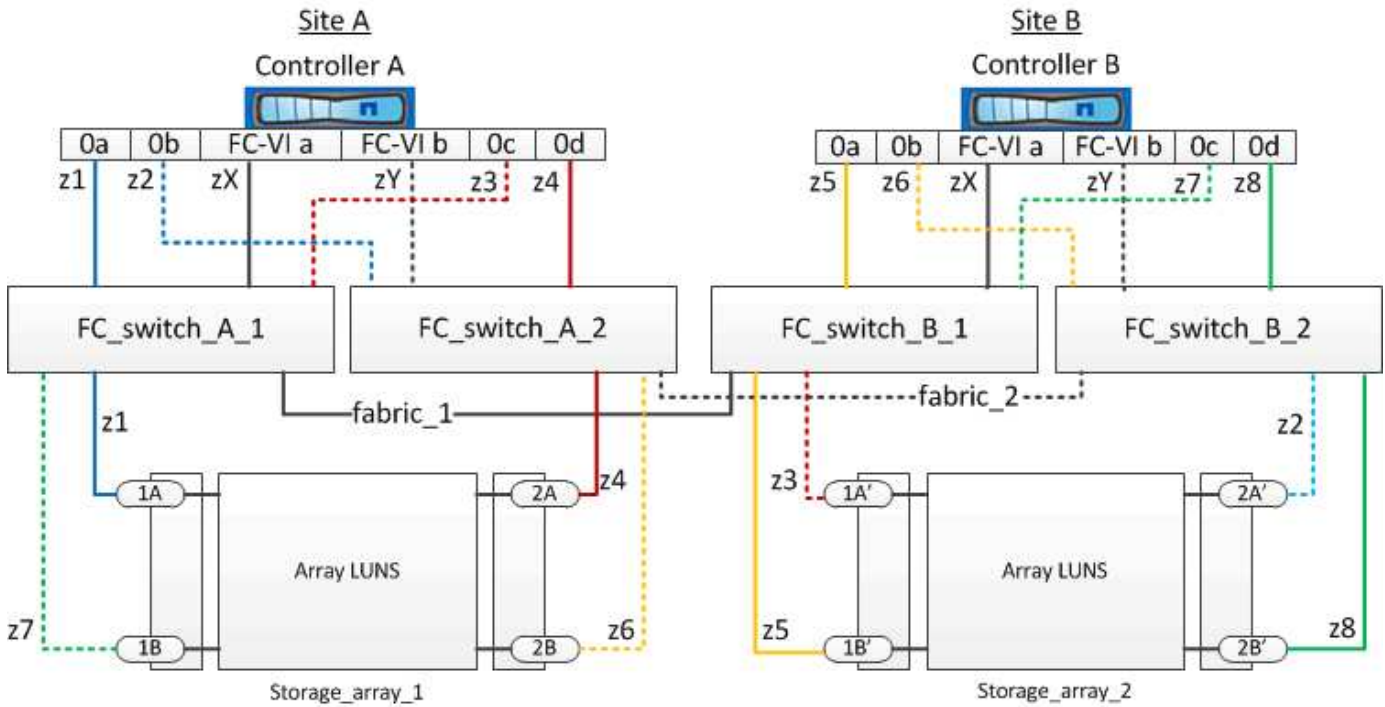
[Example of switch zoning in a four-node MetroCluster configuration with array LUNs](#)

[Example of switch zoning in an eight-node MetroCluster configuration with array LUNs](#)

Example of switch zoning in a two-node MetroCluster configuration with array LUNs

Switch zoning defines paths between connected nodes. Configuring the zoning enables you to define which array LUNs can be viewed by specific ONTAP systems.

You can use the following example as a reference when determining zoning for a two-node fabric-attached MetroCluster configuration with array LUNs:



The example shows single-initiator to single-target zoning for the MetroCluster configurations. The lines in the example represent zones rather than connections; each line is labeled with its zone number.

In the example, array LUNs are allocated on each storage array. LUNs of equal size are provisioned on the storage arrays at both sites, which is a SyncMirror requirement. Each ONTAP system has two paths to array LUNs. The ports on the storage array are redundant.

The redundant array port pairs for both the sites are as follows:

- Storage array at Site A:
 - Ports 1A and 2A
 - Ports 1B and 2B
- Storage array at Site B:
 - Ports 1A' and 2A'
 - Ports 1B' and 2B'

The redundant port pairs on each storage array form alternate paths. Therefore, both the ports of the port pairs can access the LUNs on the respective storage arrays.

The following table shows the zones for the illustrations:

Zone	ONTAP controller and initiator port	Storage array port
FC_switch_A_1		
z1	Controller A: Port 0a	Port 1A
z3	Controller A: Port 0c	Port 1A'

FC_switch_A_2		
z2	Controller A: Port 0b	Port 2A'
z4	Controller A: Port 0d	Port 2A
FC_switch_B_1		
z5	Controller B: Port 0a	Port 1B'
z7	Controller B: Port 0c	Port 1B
FC_switch_B_2		
z6	Controller B: Port 0b	Port 2B
z8	Controller B: Port 0d	Port 2B'

The following table shows the zones for the FC-VI connections:

Zone	ONTAP controller and initiator port	Switch
Site A		
zX	Controller A: Port FC-VI a	FC_switch_A_1
zY	Controller A: Port FC-VI b	FC_switch_A_2
Site B		
zX	Controller B: Port FC-VI a	FC_switch_B_1
zY	Controller B: Port FC-VI b	FC_switch_B_2

Related information

- Switch zoning defines paths between connected nodes. Configuring the zoning enables you to define which array LUNs can be viewed by a specific ONTAP system.

[Requirements for switch zoning in a MetroCluster configuration with array LUNs](#)

[Example of switch zoning in a four-node MetroCluster configuration with array LUNs](#)


- When using switch zoning in a MetroCluster configuration with array LUNs, you must ensure that certain basic requirements are followed.

[Example of switch zoning in an eight-node MetroCluster configuration with array LUNs](#)

Example of switch zoning in a four-node MetroCluster configuration with array LUNs

Switch zoning defines paths between connected nodes. Configuring the zoning enables you to define which array LUNs can be viewed by a specific ONTAP systems.

You can use the following example as a reference when determining zoning for a four-node MetroCluster configuration with array LUNs. The example shows single-initiator to single-target zoning for a MetroCluster configuration. The lines in the following example represent zones rather than connections; each line is labeled with its zone number:



In the illustration, array LUNs are allocated on each storage array for the MetroCluster configuration. LUNs of equal size are provisioned on the storage arrays at both sites, which is a SyncMirror requirement. Each ONTAP system has two paths to array LUNs. The ports on the storage array are redundant.

In the illustration, the redundant array port pairs for both the sites are as follows:

- Storage array at Site A:
 - Ports 1A and 2A
 - Ports 1B and 2B
 - Ports 1C and 2C
 - Ports 1D and 2D
- Storage array at Site B:
 - Ports 1A' and 2A'
 - Ports 1B' and 2B'
 - Ports 1C' and 2C'
 - Ports 1D' and 2D'

The redundant port pairs on each storage array form alternate paths. Therefore, both the ports of the port pairs

can access the LUNs on the respective storage arrays.

The following tables show the zones for this example:

Zones for FC_switch_A_1

Zone	ONTAP controller and initiator port	Storage array port
z1	Controller_A_1: Port 0a	Port 1A
z3	Controller_A_1: Port 0c	Port 1A'
z5	Controller_A_2: Port 0a	Port 1B
z7	Controller_A_2: Port 0c	Port 1B'

Zones for FC_switch_A_2

Zone	ONTAP controller and initiator port	Storage array port
z2	Controller_A_1: Port 0b	Port 2A'
z4	Controller_A_1: Port 0d	Port 2A
z6	Controller_A_2: Port 0b	Port 2B'
z8	Controller_A_2: Port 0d	Port 2B

Zones for FC_switch_B_1

Zone	ONTAP controller and initiator port	Storage array port
z9	Controller_B_1: Port 0a	Port 1C'
z11	Controller_B_1: Port 0c	Port 1C
z13	Controller_B_2: Port 0a	Port 1D'
z15	Controller_B_2: Port 0c	Port 1D

Zones for FC_switch_B_2

Zone	ONTAP controller and initiator port	Storage array port
------	-------------------------------------	--------------------

z10	Controller_B_1: Port 0b	Port 2C
z12	Controller_B_1: Port 0d	Port 2C'
z14	Controller_B_2: Port 0b	Port 2D
z16	Controller_B_2: Port 0d	Port 2D'

Zones for the FC-VI connections at Site A

Zone	ONTAP controller and FC initiator port	Switch
zX	Controller_A_1: Port FC-VI a	FC_switch_A_1
zY	Controller_A_1: Port FC-VI b	FC_switch_A_2
zX	Controller_A_2: Port FC-VI a	FC_switch_A_1
zY	Controller_A_2: Port FC-VI b	FC_switch_A_2

Zones for the FC-VI connections at Site B

Zone	ONTAP controller and FC initiator port	Switch
zX	Controller_B_1: Port FC-VI a	FC_switch_B_1
zY	Controller_B_1: Port FC-VI b	FC_switch_B_2
zX	Controller_B_2: Port FC-VI a	FC_switch_B_1
zY	Controller_B_2: Port FC-VI b	FC_switch_B_2

Related information

- Switch zoning defines paths between connected nodes. Configuring the zoning enables you to define which array LUNs can be viewed by specific ONTAP systems.

[Example of switch zoning in a two-node MetroCluster configuration with array LUNs](#)

[Example of switch zoning in an eight-node MetroCluster configuration with array LUNs](#)

- When using switch zoning in a MetroCluster configuration with array LUNs, you must ensure that certain basic requirements are followed.

[Requirements for switch zoning in a MetroCluster configuration with array LUNs](#)

Example of switch zoning in an eight-node MetroCluster configuration with array LUNs

Switch zoning defines paths between connected nodes. Configuring the zoning enables you to define which array LUNs can be viewed by specific ONTAP systems.

An eight-node MetroCluster configuration consists of two four-node DR groups. The first DR group consists of the following nodes:

- controller_A_1
- controller_A_2
- controller_B_1
- controller_B_2

The second DR group consists of the following nodes:

- controller_A_3
- controller_A_4
- controller_B_3
- controller_B_4

To configure the switch zoning, you can use the zoning examples for a four-node MetroCluster configuration for the first DR group.

[Example of switch zoning in a four-node MetroCluster configuration with array LUNs](#)

To configure zoning for the second DR group, follow the same examples and requirements for the FC initiator ports and array LUNs belonging to the controllers in the second DR group.

Related information

- Switch zoning defines paths between connected nodes. Configuring the zoning enables you to define which array LUNs can be viewed by specific ONTAP systems.

[Example of switch zoning in a two-node MetroCluster configuration with array LUNs](#)

[Example of switch zoning in a four-node MetroCluster configuration with array LUNs](#)

- When using switch zoning in a MetroCluster configuration with array LUNs, you must ensure that certain basic requirements are followed.

[Requirements for switch zoning in a MetroCluster configuration with array LUNs](#)

Set up ONTAP in a MetroCluster configuration with array LUNs

Verifying and configuring the HA state of components in Maintenance mode

When configuring a storage system in a MetroCluster configuration, you must make sure that the high-availability (HA) state of the controller module and chassis components is

"mcc" or "mcc-2n" so that these components boot properly.

Before you begin

The system must be in Maintenance mode.

About this task

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

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.

Number of controllers in the MetroCluster configuration	HA state for all components should be...
Eight- or four-node MetroCluster FC configuration	mcc
Two-node MetroCluster FC configuration	mcc-2n
MetroCluster IP configuration	mccip

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

Number of controllers in the MetroCluster configuration	Command
Eight- or four-node MetroCluster FC configuration	ha-config modify controller mcc
Two-node MetroCluster FC configuration	ha-config modify controller mcc-2n
MetroCluster IP configuration	ha-config modify controller mccip

3. If the displayed system state of the chassis is not correct, set the HA state for the chassis:

Number of controllers in the MetroCluster configuration	Command
Eight- or four-node MetroCluster FC configuration	ha-config modify chassis mcc
Two-node MetroCluster FC configuration	ha-config modify chassis mcc-2n
MetroCluster IP configuration	ha-config modify chassis mccip

4. Boot the node to ONTAP:

```
boot_ontap
```

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

Configuring ONTAP on a system that uses only array LUNs

If you want to configure ONTAP for use with array LUNs, you must configure the root aggregate and root volume, reserve space for diagnostics and recovery operations, and set up the cluster.

Before you begin

- The ONTAP system must be connected to the storage array.
- The storage array administrator must have created LUNs and presented them to ONTAP.
- The storage array administrator must have configured the LUN security.

About this task

You must configure each node that you want to use with array LUNs. If the node is in an HA pair, then you must complete the configuration process on one node before proceeding with the configuration on the partner node.

Steps

1. Power on the primary node and interrupt the boot process by pressing Ctrl-C when you see the following message on the console:

```
Press CTRL-C for special boot menu.
```

2. Select option **4 (Clean configuration and initialize all disks)** on the boot menu.

The list of array LUNs made available to ONTAP is displayed. In addition, the array LUN size required for root volume creation is also specified. The size required for root volume creation differs from one ONTAP system to another.

- If no array LUNs were previously assigned, ONTAP detects and displays the available array LUNs, as shown in the following example:

```

mcc8040-ams1::> disk show NET-1.6 -instance
      Disk: NET-1.6
      Container Type: aggregate
      Owner/Home: mcc8040-ams1-01 / mcc8040-ams1-01
      DR Home: -
      Stack ID/Shelf/Bay: - / - / -
      LUN: 0
      Array: NETAPP_INF_1
      Vendor: NETAPP
      Model: INF-01-00
      Serial Number: 60080E50004317B4000003B158E35974
      UID:
60080E50:004317B4:000003B1:58E35974:00000000:00000000:00000000:000000
00:00000000:00000000
      BPS: 512
      Physical Size: 87.50GB
      Position: data
      Checksum Compatibility: block
      Aggregate: eseries
      Plex: plex0

Paths:

      LUN  Initiator Side      Target
Side                               Link
Controller      Initiator  ID  Switch Port      Switch
Port            Acc Use  Target Port      TPGN  Speed
I/O KB/s            IOPS
-----
-----
-----
mcc8040-ams1-01    2c                0  mccb6505-ams1:16  mccb6505-
ams1:18          AO  INU  20330080e54317b4  1  4 Gb/S
0                0
mcc8040-ams1-01    2a                0  mccb6505-ams1:17  mccb6505-
ams1:19          ANO RDY 20320080e54317b4  0  4 Gb/S
0                0

Errors:
-
```

- If array LUNs were previously assigned, for example, through the maintenance mode, they are either marked local or partner in the list of the available array LUNs, depending on whether the array LUNs were selected from the node on which you are installing ONTAP or its HA partner:

In this example, array LUNs with index numbers 3 and 6 are marked "local" because they had been previously assigned from this particular node:

*
* No disks are owned by this node, but array LUNs are assigned.
*
* You can use the following information to verify connectivity from
*
* HBAs to switch ports. If the connectivity of HBAs to switch ports
*
* does not match your expectations, configure your SAN and rescan.
*
* You can rescan by entering 'r' at the prompt for selecting
*
* array LUNs below.

WPN	HBA	HBA WWPN	Switch port	Switch port
-----	---	-----	-----	-----
20190027f88948dd	0e	500a098001baf8e0	vgbr6510s203:25	
2011547feead680	0f	500a098101baf8e0	vgci9710s202:1-17	
201b0027f88948dd	0g	500a098201baf8e0	vgbr6510s203:27	
2012547feead680	0h	500a098301baf8e0	vgci9710s202:1-18	

No native disks were detected, but array LUNs were detected.
You will need to select an array LUN to be used to create the root aggregate and root volume.

The array LUNs visible to the system are listed below. Select one array LUN to be used to create the root aggregate and root volume. **The root volume requires 350.0 GB of space.**

Warning: The contents of the array LUN you select will be erased by ONTAP prior to their use.

Index	Array LUN Name	Model	Vendor	Size	Owner
Checksum	Serial Number				
-----	-----	-----	-----	-----	-----
-----	-----	-----	-----	-----	-----

0	vgci9710s202:2-24.0L19	RAID5	DGC	217.3	GB	
Block	6006016083402B0048E576D7					
1	vgbr6510s203:30.126L20	RAID5	DGC	217.3	GB	
Block	6006016083402B0049E576D7					
2	vgci9710s202:2-24.0L21	RAID5	DGC	217.3	GB	
Block	6006016083402B004AE576D7					
3	vgbr6510s203:30.126L22	RAID5	DGC	405.4	GB	local
Block	6006016083402B004BE576D7					
4	vgci9710s202:2-24.0L23	RAID5	DGC	217.3	GB	
Block	6006016083402B004CE576D7					
5	vgbr6510s203:30.126L24	RAID5	DGC	217.3	GB	
Block	6006016083402B004DE576D7					
6	vgbr6510s203:30.126L25	RAID5	DGC	423.5	GB	local
Block	6006016083402B003CF93694					
7	vgci9710s202:2-24.0L26	RAID5	DGC	423.5	GB	
Block	6006016083402B003DF93694					

3. Select the index number corresponding to the array LUN you want to assign as the root volume.

The array LUN must be of sufficient size to create the root volume.

The array LUN selected for root volume creation is marked "local (root)".

In the following example, the array LUN with index number 3 is marked for root volume creation:

The root volume will be created on switch 0:5.183L33.

ONTAP requires that 11.0 GB of space be reserved for use in diagnostic and recovery operations. Select one array LUN to be used as spare for diagnostic and recovery operations.

Index	Array LUN Name	Model	Vendor	Size	Owner
Checksum	Serial Number				
0	switch0:5.183L1	SYMMETRIX	EMC	266.1 GB	
Block	600604803436313734316631				
1	switch0:5.183L3	SYMMETRIX	EMC	266.1 GB	
Block	600604803436316333353837				
2	switch0:5.183L31	SYMMETRIX	EMC	266.1 GB	
Block	600604803436313237643666				
3	switch0:5.183L33	SYMMETRIX	EMC	658.3 GB	local (root)
Block	600604803436316263613066				
4	switch0:7.183L0	SYMMETRIX	EMC	173.6 GB	
Block	600604803436313261356235				
5	switch0:7.183L2	SYMMETRIX	EMC	173.6 GB	
Block	600604803436313438396431				
6	switch0:7.183L4	SYMMETRIX	EMC	658.3 GB	
Block	600604803436313161663031				
7	switch0:7.183L30	SYMMETRIX	EMC	173.6 GB	
Block	600604803436316538353834				
8	switch0:7.183L32	SYMMETRIX	EMC	266.1 GB	
Block	600604803436313237353738				
9	switch0:7.183L34	SYMMETRIX	EMC	658.3 GB	
Block	600604803436313737333662				

4. Select the index number corresponding to the array LUN you want to assign for use in diagnostic and recovery options.

The array LUN must be of sufficient size for use in diagnostic and recovery options. If required, you can also select multiple array LUNs with a combined size greater than or equal to the specified size. To select multiple entries, you must enter the comma-separated values of all of the index numbers corresponding to the array LUNs you want to select for diagnostic and recovery options.

The following example shows a list of array LUNs selected for root volume creation and for diagnostic and recovery options:

Here is a list of the selected array LUNs

Index	Array LUN Name	Model	Vendor	Size	Owner
Checksum	Serial Number				
-----	-----	-----	-----	-----	-----
2	switch0:5.183L31	SYMMETRIX	EMC	266.1 GB	local
Block	600604803436313237643666				
3	switch0:5.183L33	SYMMETRIX	EMC	658.3 GB	local (root)
Block	600604803436316263613066				
4	switch0:7.183L0	SYMMETRIX	EMC	173.6 GB	local
Block	600604803436313261356235				
5	switch0:7.183L2	SYMMETRIX	EMC	173.6 GB	local
Block	600604803436313438396431				

Do you want to continue (yes|no)?



Selecting “no” clears the LUN selection.

5. Enter **y** when prompted by the system to continue with the installation process.

The root aggregate and the root volume are created and the rest of the installation process continues.

6. Enter the required details to create the node management interface.

The following example shows the node management interface screen with a message confirming the creation of the node management interface:


```
Welcome to node setup.
```

```
You can enter the following commands at any time:
```

```
"help" or "?" - if you want to have a question clarified,  
"back" - if you want to change previously answered questions, and  
"exit" or "quit" - if you want to quit the setup wizard.  
Any changes you made before quitting will be saved.
```

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

```
Enter the node management interface port [e0M]:
```

```
Enter the node management interface IP address: 192.0.2.66
```

```
Enter the node management interface netmask: 255.255.255.192
```

```
Enter the node management interface default gateway: 192.0.2.7
```

```
A node management interface on port e0M with IP address 192.0.2.66 has  
been created.
```

```
This node has its management address assigned and is ready for cluster  
setup.
```

After you finish

After configuring ONTAP on all of the nodes that you want to use with array LUNs, you should complete the [Cluster setup process](#)

Related information

[FlexArray virtualization installation requirements and reference](#)

Setting up the cluster

Setting up the cluster involves setting up each node, creating the cluster on the first node, and joining any remaining nodes to the cluster.

Related information

[Software setup](#)

Installing the license for using array LUNs in a MetroCluster configuration

You must install the V_StorageAttach license on each MetroCluster node that you want to use with array LUNs. You cannot use array LUNs in an aggregate until the license is installed.

Before you begin

- The cluster must be installed.
- You must have the license key for the V_StorageAttach license.

About this task

You must use a separate license key for each node on which you want to install the V_StorageAttach license.

Steps

1. Install the V_StorageAttach license.

```
system license add
```

Repeat this step for each cluster node on which you want to install the license.

2. Verify that the V_StorageAttach license is installed on all required nodes in a cluster.

```
system license show
```

The following sample output shows that the V_StorageAttach license is installed on the nodes of cluster_A:

```
cluster_A::> system license show
Serial Number: nnnnnnnn
Owner: controller_A_1
Package          Type      Description          Expiration
-----
V_StorageAttach  license  Virtual Attached Storage

Serial Number: 11111111
Owner: controller_A_2
Package          Type      Description          Expiration
-----
V_StorageAttach  license  Virtual Attached Storage
```

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:fc.adapter.offlining:info]: Offlining
Fibre Channel adapter 1a.
Host adapter 1a disable succeeded
Jun 03 02:17:57 [controller_B_1:fc.adapter.offline:info]: Fibre Channel
adapter 1a is now offline.
*> storage disable adapter 1b
Jun 03 02:18:43 [controller_B_1:fc.adapter.offlining:info]: Offlining
Fibre Channel adapter 1b.
Host adapter 1b disable succeeded
Jun 03 02:18:43 [controller_B_1:fc.adapter.offline:info]: Fibre Channel
adapter 1b is now offline.
*>
```

2. Verify that the ports are disabled:

```
ucadmin show
```

```
*> ucadmin show
```

Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin 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:

```
ucadmin show
```

```
*> ucadmin show
      Current  Current  Pending  Pending  Admin
Adapter Mode    Type     Mode     Type     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    Type     Mode     Type     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 displayed.
```

Assigning ownership of array LUNs

Array LUNs must be owned by a node before they can be added to an aggregate to be used as storage.

Before you begin

- Back-end configuration testing (testing of the connectivity and configuration of devices behind the ONTAP systems) must be completed.
- Array LUNs that you want to assign must be presented to the ONTAP systems.

About this task

You can assign ownership of array LUNs that have the following characteristics:

- They are unowned.
- They have no storage array configuration errors, such as the following:
 - The array LUN is smaller than or larger than the size that ONTAP supports.
 - The LDEV is mapped on only one port.
 - The LDEV has inconsistent LUN IDs assigned to it.
 - The LUN is available on only one path.

ONTAP issues an error message if you try to assign ownership of an array LUN with back-end configuration errors that would interfere with the ONTAP system and the storage array operating together. You must fix such errors before you can proceed with array LUN assignment.

ONTAP alerts you if you try to assign an array LUN with a redundancy error: for example, all paths to this array LUN are connected to the same controller or only one path to the array LUN. You can fix a redundancy error before or after assigning ownership of the LUN.

Steps

1. View the array LUNs that have not yet been assigned to a node:

```
storage disk show -container-type unassigned
```

2. Assign an array LUN to this node:

```
storage disk assign -disk array_LUN_name -owner nodename
```

If you want to fix a redundancy error after disk assignment instead of before, you must use the `-force` parameter with the `storage disk assign` command.

Related information

[FlexArray virtualization installation requirements and reference](#)

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.

Steps

1. Configure intercluster LIFs using the procedure in:

[Configuring intercluster LIFs](#)

2. Create a cluster peer relationship using the procedure in:

[Peering the clusters](#)

Mirroring the root aggregates

You must mirror the root aggregates in your MetroCluster configuration to ensure data protection.

Before you begin

You must have ensured that the SyncMirror requirements for the MetroCluster configuration with array LUNs are satisfied. Refer to [Requirements for a MetroCluster configuration with array LUNs](#).

About this task

You must repeat this task for each controller in the MetroCluster configuration.

Step

1. Mirror the unmirrored root aggregate:

```
storage aggregate mirror
```

The following command mirrors the root aggregate for controller_A_1:

```
controller_A_1::> storage aggregate mirror aggr0_controller_A_1
```

The root aggregate is mirrored with array LUNs from pool1.

Creating data aggregates on, implementing, and verifying the MetroCluster configuration

You must create data aggregates on each node, implement, and verify the MetroCluster configuration.

Steps

1. Create data aggregates on each node:
 - a. Create a mirrored data aggregate on each node:
[Mirror the root aggregates.](#)
 - b. If needed, create unmirrored data aggregates:
[Create a mirrored data aggregate on each node.](#)

2. [Implement the MetroCluster configuration.](#)

3. [Configure the MetroCluster FC switches for health monitoring.](#)
4. Check and verify the configuration:
 - a. [Check the MetroCluster configuration.](#)
 - b. [Check for MetroCluster configuration errors with Config Advisor.](#)
 - c. [Verify switchover, healing, and switchback.](#)
5. Install and configure the MetroCluster Tiebreaker software:
 - a. [Install the Tiebreaker Software.](#)
 - b. [Configure the Tiebreaker Software.](#)
6. Set the destination for configuration backup files:

[Protect configuration backup files.](#)

Implement a MetroCluster configuration with both disks and array LUNs

Implementing a MetroCluster configuration with both disks and array LUNs

To implement a MetroCluster configuration with native disks and array LUNs, you must ensure that the ONTAP systems used in the configuration can attach to storage arrays.

A MetroCluster configuration with disks and array LUNs can have either two or four nodes. Although the four-node MetroCluster configuration must be fabric-attached, the two-node configuration can either be stretch or fabric-attached.

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.

Related information

For setting up a two-node fabric-attached MetroCluster configuration or a four-node MetroCluster configuration with native disks and array LUNs, you must use FC-to-SAS bridges to connect the ONTAP systems with the disk shelves through the FC switches. You can connect array LUNs through the FC switches to the ONTAP systems.

[Example of a two-node fabric-attached MetroCluster configuration with disks and array LUNs](#)

[Example of a four-node MetroCluster configuration with disks and array LUNs](#)

Considerations when implementing a MetroCluster configuration with disks and array LUNs

When planning your MetroCluster configuration for use with disks and array LUNs, you must consider various factors, such as the order of setting up access to storage, root aggregate location, and the usage of FC initiator ports, switches, and FC-to-SAS bridges.

Consider the information in the following table when planning your configuration:

Consideration	Guideline
Order of setting up access to the storage	You can set up access to either disks or array LUNs first. You must complete all setup for that type of storage and verify that it is set up correctly before setting up the other type of storage.
Location of the root aggregate	<ul style="list-style-type: none"> • If you are setting up a <i>new</i> MetroCluster deployment with both disks and array LUNs, you must create the root aggregate on native disks. <p>When doing this, ensure that <i>at least one</i> disk shelf (with 24 disk drives) is set up at each of the sites.</p> <ul style="list-style-type: none"> • If you are adding native disks to an <i>existing</i> MetroCluster configuration that uses array LUNs, the root aggregate can remain on an array LUN.
Using switches and FC-to-SAS bridges	<p>FC-to-SAS bridges are required in four-node configurations and two-node fabric-attached configurations to connect the ONTAP systems to the disk shelves through the switches.</p> <p>You must use the same switches to connect to the storage arrays and the FC-to-SAS bridges.</p>
Using FC initiator ports	<p>The initiator ports used to connect to an FC-to-SAS bridge must be different from the ports used to connect to the switches, which connect to the storage arrays.</p> <p>A minimum of eight initiator ports is required to connect an ONTAP system to both disks and array LUNs.</p>

Related information

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

[Configuring the Brocade FC switches manually](#)

[Configuring the Cisco FC switches manually](#)

- You install and cable ATTO FibreBridge bridges and SAS disk shelves when adding new storage to the configuration.

[Installing FC-to-SAS bridges and SAS disk shelves](#)

- Switch zoning defines paths between connected nodes. Configuring the zoning enables you to define which array LUNs can be viewed by a specific ONTAP system.

[Example of switch zoning in a four-node MetroCluster configuration with array LUNs](#)

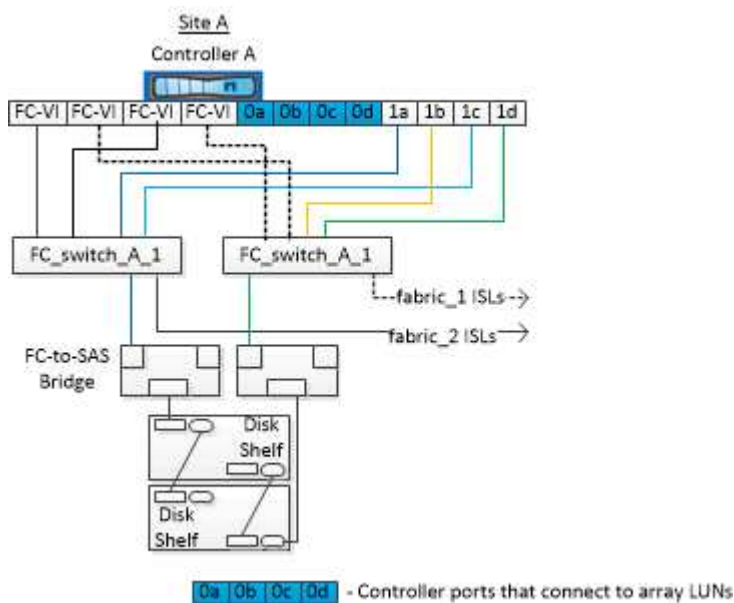
- [NetApp Hardware Universe](#)

Example of a two-node fabric-attached MetroCluster configuration with disks and array LUNs


For setting up a two-node fabric-attached MetroCluster configuration with native disks and array LUNs, you must use FC-to-SAS bridges to connect the ONTAP systems with the disk shelves through the FC switches. You can connect array LUNs through the FC switches to the ONTAP systems.

The following illustrations represent examples of a two-node fabric-attached MetroCluster configuration with disks and array LUNs. They both represent the same MetroCluster configuration; the representations for disks and array LUNs are separated only for simplification.

In the following illustration showing the connectivity between ONTAP systems and disks, the HBA ports 1a through 1d are used for connectivity with disks through the FC-to-SAS bridges:



In the following illustration showing the connectivity between ONTAP systems and array LUNs, the HBA ports 0a through 0d are used for connectivity with array LUNs because ports 1a through 1d are used for connectivity with disks:



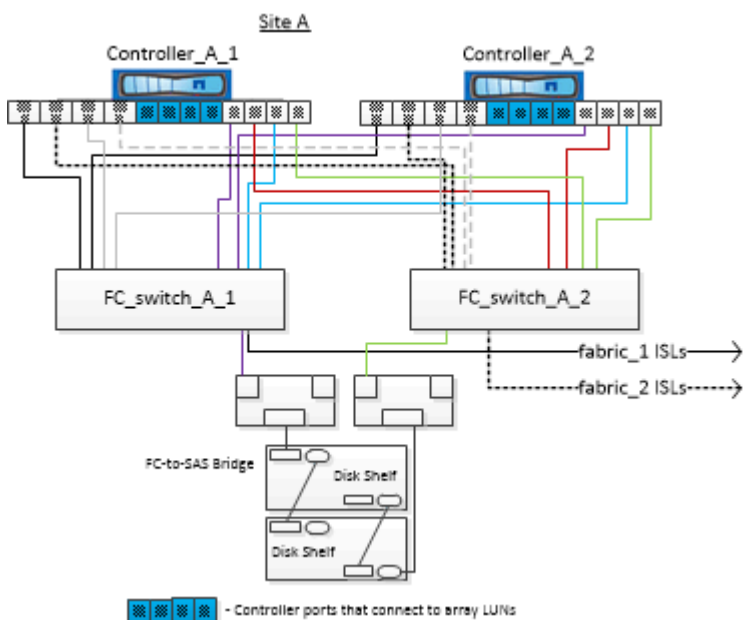
Example of a four-node MetroCluster configuration with disks and array LUNs

For setting up a four-node MetroCluster configuration with native disks and array LUNs, you must use FC-to-SAS bridges to connect the ONTAP systems with the disk shelves through the FC switches. You can connect array LUNs through the FC switches to the ONTAP systems.

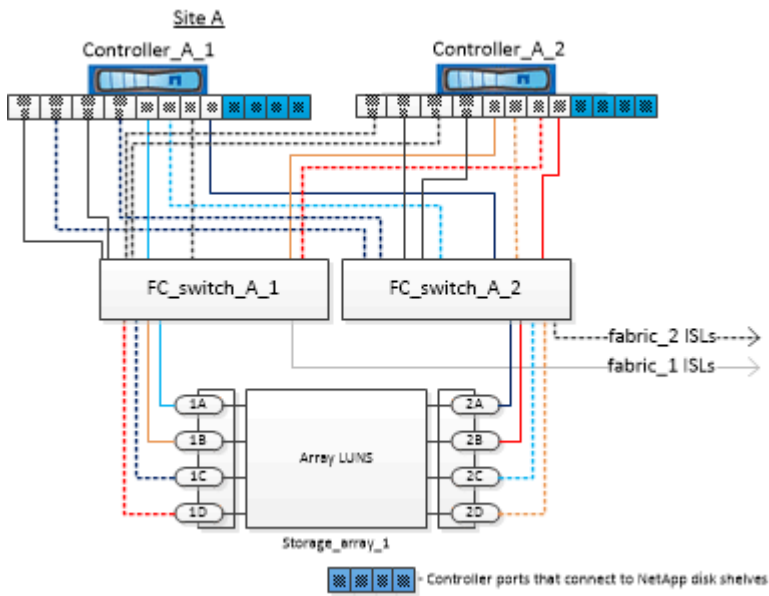
A minimum of eight initiator ports is required for an ONTAP system to connect to both native disks and array LUNs.

The following illustrations represent examples of a MetroCluster configuration with disks and array LUNs. They both represent the same MetroCluster configuration; the representations for disks and array LUNs are separated only for simplification.

In the following illustration that shows the connectivity between ONTAP systems and disks, the HBA ports 1a through 1d are used for connectivity with disks through the FC-to-SAS bridges:



In the following illustration that shows the connectivity between ONTAP systems and array LUNs, the HBA ports 0a through 0d are used for connectivity with array LUNs because ports 1a through 1d are used for connectivity with disks:



Copyright Information

Copyright © 2022 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means-graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system- without prior written permission of the copyright owner.

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

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

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NETAPP, the NETAPP logo, 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.