



Configure the MetroCluster hardware components

ONTAP MetroCluster

NetApp
March 06, 2026

This PDF was generated from https://docs.netapp.com/us-en/ontap-metrocluster/install-ip/concept_parts_of_an_ip_mcc_configuration_mcc_ip.html on March 06, 2026. Always check docs.netapp.com for the latest.

Table of Contents

- Configure the MetroCluster hardware components 1
 - Learn about hardware component interconnections in a MetroCluster IP configuration 1
 - Key hardware elements 1
 - Disaster Recovery (DR) groups 2
 - Illustration of the local HA pairs in a MetroCluster configuration 2
 - Illustration of the MetroCluster IP and cluster interconnect network 3
 - Illustration of the cluster peering network 4
- Required MetroCluster IP configuration components and naming conventions 5
 - Supported software and hardware 5
 - Hardware redundancy requirements in a MetroCluster IP configuration 5
 - ONTAP cluster requirements in a MetroCluster IP configuration 6
 - IP switch requirements in a MetroCluster IP configuration 6
 - Controller module requirements in a MetroCluster IP configuration 6
 - Gigabit Ethernet adapter requirements in a MetroCluster IP configuration 7
 - Pool and drive requirements (minimum supported) 8
 - Drive location considerations for partially populated shelves 8
 - Drive location considerations for AFF A800 internal drives 9
- Rack the MetroCluster IP configuration hardware components 9
- Cable the MetroCluster IP switches 10
 - How to use the port tables with multiple MetroCluster IP configurations 10
 - Platform port assignments for Cisco 3132Q-V switches in a MetroCluster IP configuration 11
 - Platform port assignments for Cisco 3232C or 36-port Cisco 9336C switches in a MetroCluster IP configuration 14
 - Platform port assignments for 12-port Cisco 9336C-FX2 switches in a MetroCluster IP configuration 25
 - Platform port assignments for a 36-port Cisco 9336C-FX2 switch connecting NS224 storage in a MetroCluster IP configuration 33
 - Platform port assignments for Broadcom supported BES-53248 IP switches in a MetroCluster IP configuration 48
 - Platform port assignments for NVIDIA supported SN2100 IP switches in a MetroCluster IP configuration 55
- Cable the ONTAP controller module ports in a MetroCluster IP configuration 61
- Configure the MetroCluster IP switches 62
 - Choose the correct MetroCluster IP switch configuration procedure 62
 - Configure Broadcom IP switches for cluster interconnect and backend MetroCluster IP connectivity 62
 - Configure Cisco IP switches 81
 - Configure NVIDIA IP switches 103
- Monitor MetroCluster IP switch health 119
 - Learn about switch health monitoring in a MetroCluster IP configuration 119
 - Configure SNMPv3 to monitor the health of MetroCluster IP switches 119
 - Configure log collection on a MetroCluster IP switch 137
 - Manage the monitoring of Ethernet switches in a MetroCluster IP configuration 144
 - Verify Ethernet switch monitoring in a MetroCluster IP configuration 145

Configure the MetroCluster hardware components

Learn about hardware component interconnections in a MetroCluster IP configuration

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

Key hardware elements

A MetroCluster IP configuration includes the following key hardware elements:

- Storage controllers

The storage controllers are configured as two two-node clusters.

- IP network

This back-end IP network provides connectivity for two distinct uses:

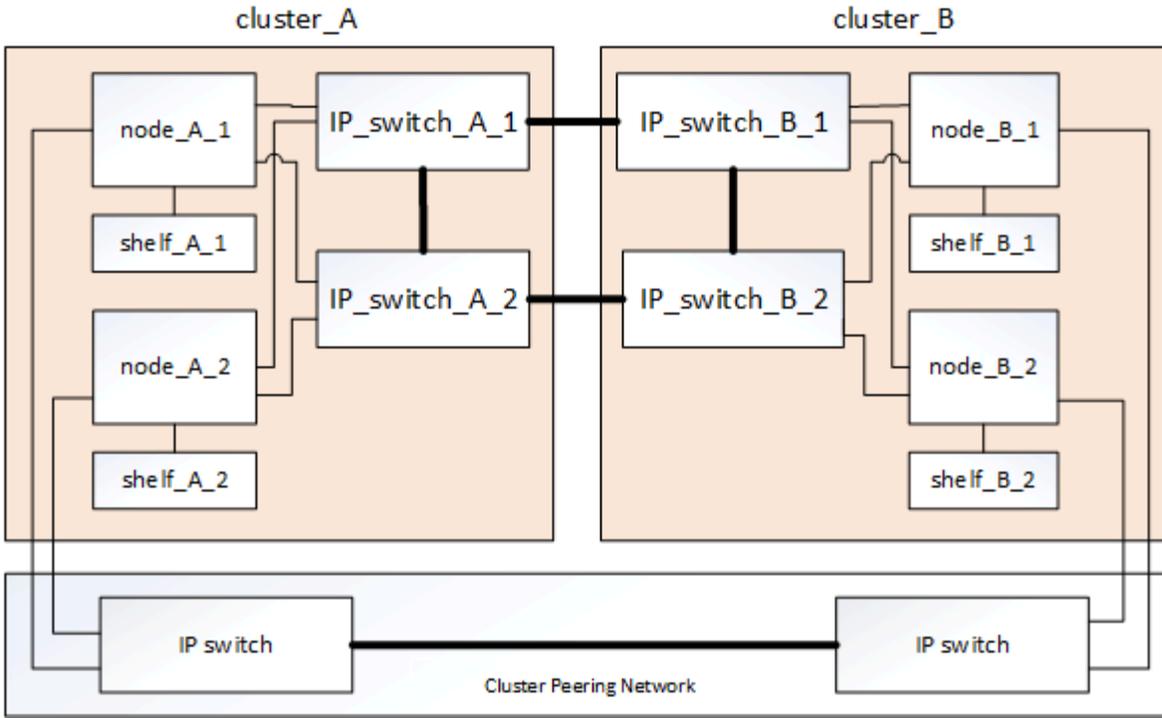
- Standard cluster connectivity for intra-cluster communications.

This is the same cluster switch functionality used in non-MetroCluster switched ONTAP clusters.

- MetroCluster back-end connectivity for replication of storage data and non-volatile cache.

- 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.



Disaster Recovery (DR) groups

A MetroCluster IP configuration consists of one DR group of four nodes.

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

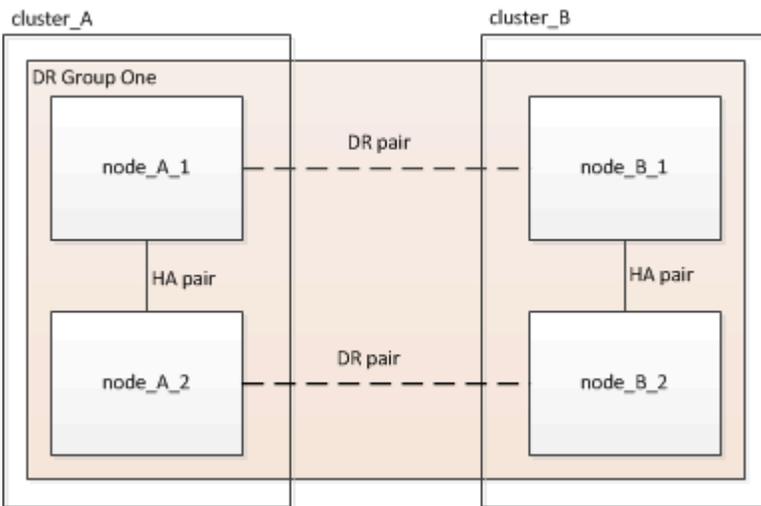
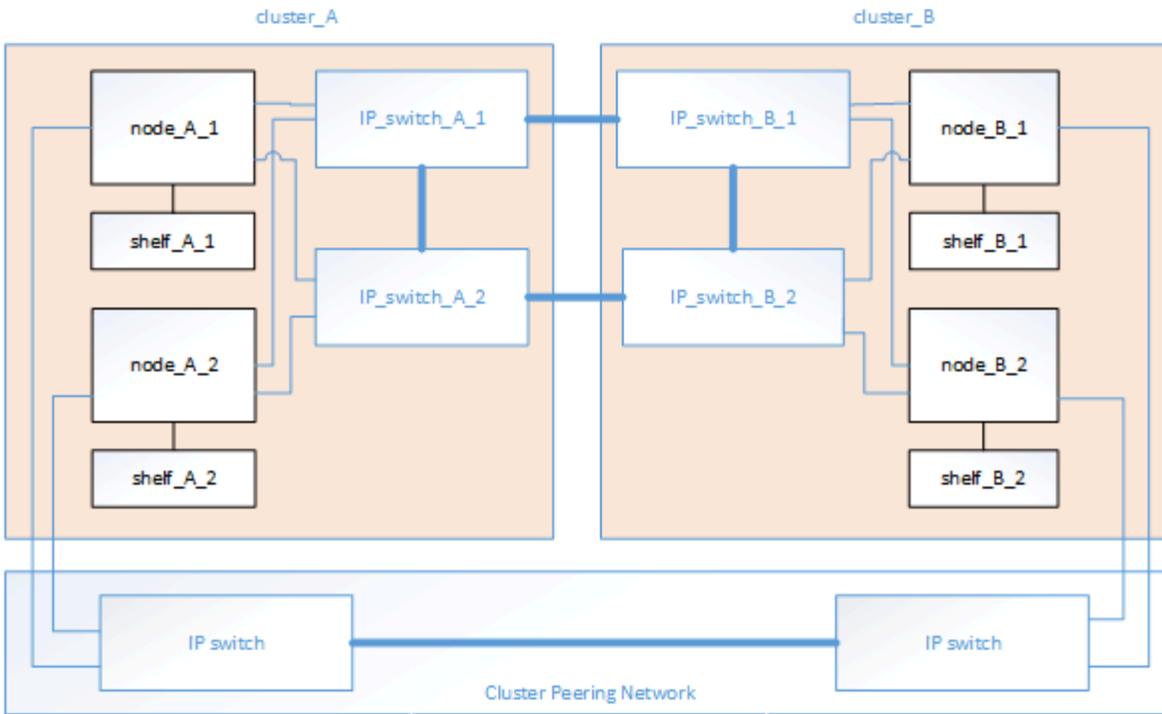


Illustration of the local HA pairs in a MetroCluster configuration

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

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

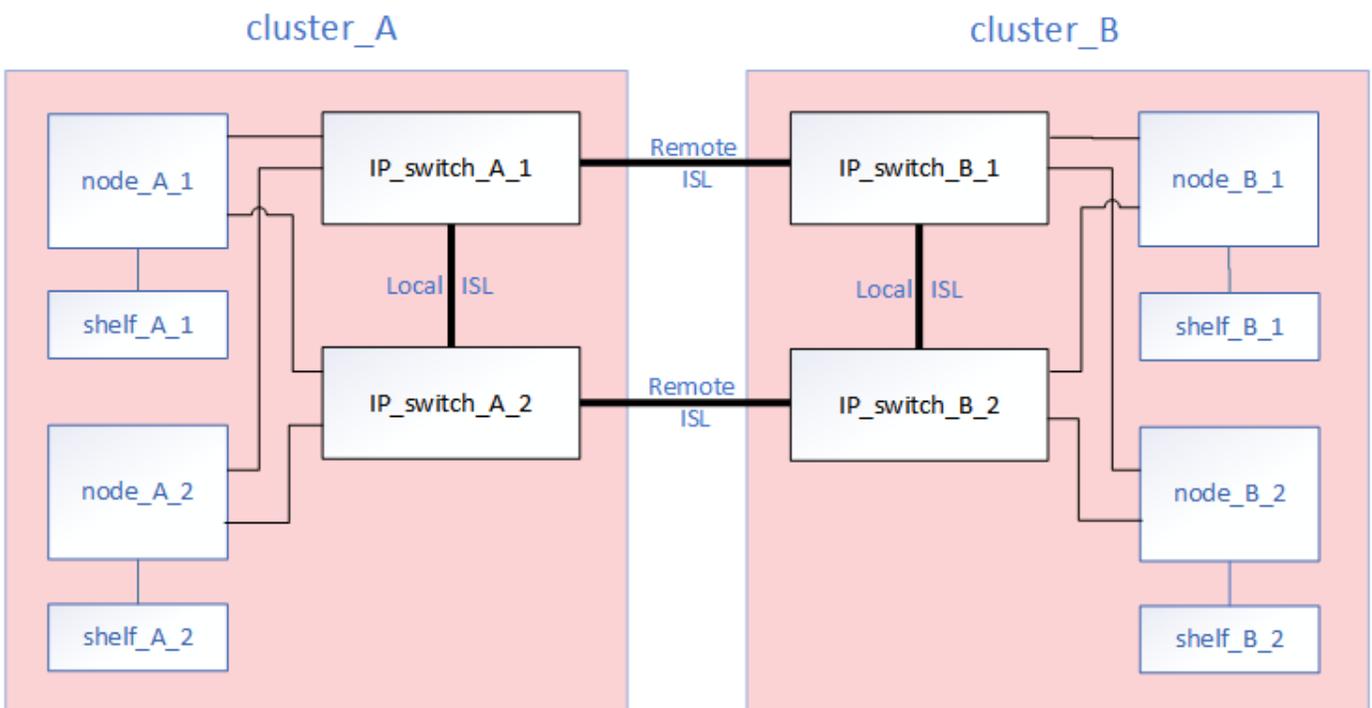


Related information

[ONTAP concepts](#)

Illustration of the MetroCluster IP and cluster interconnect network

ONTAP clusters typically include a cluster interconnect network for traffic between the nodes in the cluster. In MetroCluster IP configurations, this network is also used for carrying data replication traffic between the MetroCluster sites.

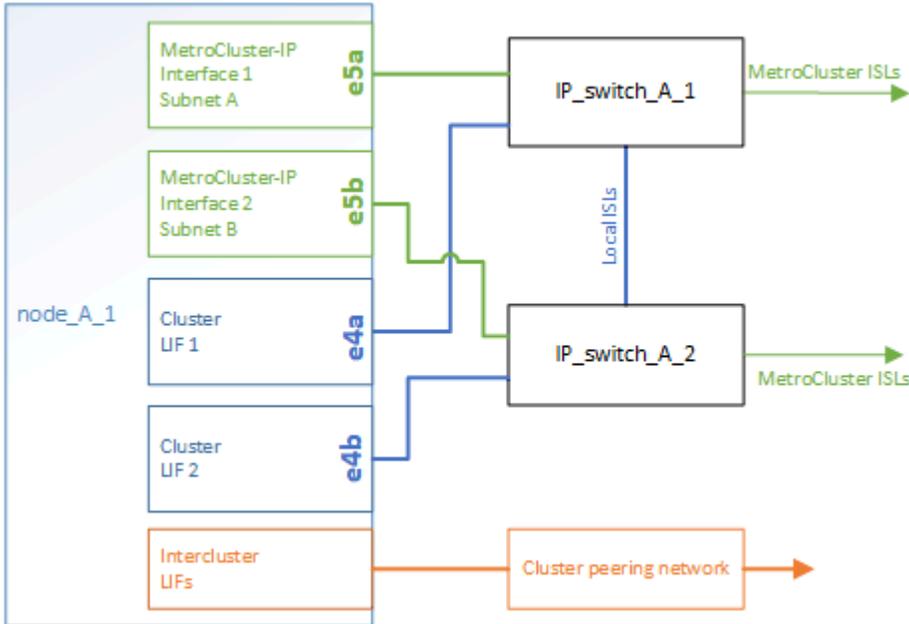


Each node in the MetroCluster IP configuration has dedicated interfaces for connection to the back-end IP

network:

- Two MetroCluster IP interfaces
- Two local cluster interfaces

The following illustration shows these interfaces. The port usage shown is for an AFF A700 or FAS9000 system.



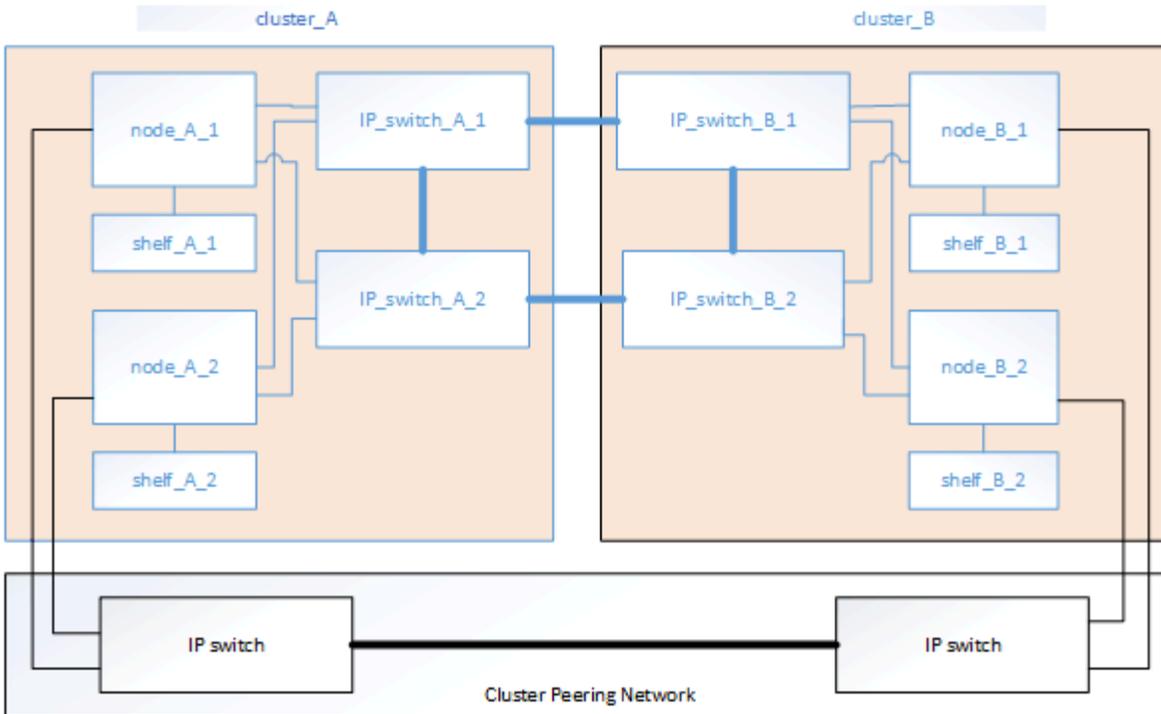
Related information

[Considerations for MetroCluster IP configurations](#)

Illustration of the cluster peering network

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

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



Related information

[Cluster and SVM peering express configuration](#)

[Considerations for configuring cluster peering](#)

[Cabling the cluster peering connections](#)

[Peering the clusters](#)

Required MetroCluster IP configuration components and naming conventions

Identify the required and supported hardware and software components for a MetroCluster IP configuration. Review the naming conventions that documentation examples use for the components.

Supported software and hardware

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

[NetApp Hardware Universe](#)

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

Hardware redundancy requirements in a MetroCluster IP configuration

Because of the hardware redundancy in the MetroCluster IP 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.

ONTAP cluster requirements in a MetroCluster IP configuration

MetroCluster IP configurations require 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

IP switch requirements in a MetroCluster IP configuration

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

The IP switches also provide intracluster communication among the controller modules in each cluster.

Naming must be unique within the MetroCluster configuration.

Example names:

- Site A: cluster_A
 - IP_switch_A_1
 - IP_switch_A_2
- Site B: cluster_B
 - IP_switch_B_1
 - IP_switch_B_2

Controller module requirements in a MetroCluster IP configuration

MetroCluster IP configurations require four or eight controller modules.

The controller modules at each site form an HA pair. Each controller module has a DR partner at the other site.

Each controller module must be running the same ONTAP version. Supported platform models depend on the ONTAP version:

- New MetroCluster IP installations on FAS systems are not supported in ONTAP 9.4.
Existing MetroCluster IP configurations on FAS systems can be upgraded to ONTAP 9.4.
- Beginning with ONTAP 9.5, new MetroCluster IP installations on FAS systems are supported.
- Beginning with ONTAP 9.4, controller modules configured for ADP are supported.

Example names

The following example names are used in the documentation:

- Site A: cluster_A
 - controller_A_1

- controller_A_2
- Site B: cluster_B
 - controller_B_1
 - controller_B_2

Gigabit Ethernet adapter requirements in a MetroCluster IP configuration

MetroCluster IP configurations use a 40/100 Gbps or 10/25 Gbps Ethernet adapter for the IP interfaces to the IP switches used for the MetroCluster IP fabric.



Onboard ports are built into the controller hardware (Slot 0) and can't be replaced, so the required slot for adapter is not applicable.

Platform model	Required Gigabit Ethernet adapter	Required slot for adapter	Ports
AFF A900, ASA A900, and FAS9500	X91146A	Slot 5, Slot 7	e5b, e7b Note: Ports e5a and e7a can only be used for intercluster LIFs. You cannot use these ports for a data LIF.
AFF A700 and FAS9000	X91146A-C	Slot 5	e5a, e5b
AFF A800, AFF C800, ASA A800, and ASA C800	X1146A/onboard ports	Slot 1/Not applicable for onboard ports	e0b, e1b
FAS8300, AFF A400, ASA A400, ASA C400, AFF C400	X1146A	Slot 1	e1a, e1b
AFF A300, FAS8200	X1116A	Slot 1	e1a, e1b
FAS2750, AFF A150, ASAA150, AFF A220	Onboard ports	Not applicable	e0a, e0b
FAS500f, AFF A250, ASA A250, ASA C250, AFF C250	Onboard ports	Not applicable	e0c, e0d
AFF A320	Onboard ports	Not applicable	e0g, e0h
AFF A70, FAS70, AFF C80	X50132A	Slot 2	e2a, e2b

AFF A90, AFF A1K, FAS90	X50132A	Slot 2, Slot 3	e2b, e3b Note: Ports e2a and e3a must remain unused. Using these ports for front-end networks or peering is not supported.
AFF A50	X60134A	Slot 2	e2a, e2b
AFF A30, AFF C30, AFF C60, FAS50	X60134A	Slot 2	e2a, e2b
AFF A20	X60132A	Slot 4, Slot 2	e2b, e4b

[Learn about automatic drive assignment and ADP systems in MetroCluster IP configurations.](#)

Pool and drive requirements (minimum supported)

A four-node MetroCluster IP configuration requires the minimum configuration at each site:

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



Aggregate names must be unique across the MetroCluster sites. This means that you cannot have two different aggregates with the same name on site A and site B.

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, RAID4, and RAID-TEC.



Beginning with ONTAP 9.4, MetroCluster IP configurations support new installations using automatic disk assignment and ADP (Advanced Drive Partitioning). Refer to [Considerations for automatic drive assignment and ADP systems](#) for more information.

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.

Drive location considerations for AFF A800 internal drives

For correct implementation of the ADP feature, the AFF A800 system disk slots must be divided into quarters and the disks must be located symmetrically in the quarters.

An AFF A800 system has 48 drive bays. The bays can be divided into quarters:

- Quarter one:
 - Bays 0 - 5
 - Bays 24 - 29
- Quarter two:
 - Bays 6 - 11
 - Bays 30 - 35
- Quarter three:
 - Bays 12 - 17
 - Bays 36 - 41
- Quarter four:
 - Bays 18 - 23
 - Bays 42 - 47

If this system is populated with 16 drives, they must be symmetrically distributed among the four quarters:

- Four drives in the first quarter: 0, 1, 2, 3
- Four drives in the second quarter: 6, 7, 8, 9
- Four drives in the third quarter: 12, 13, 14, 15
- Four drives in the fourth quarter: 18, 19, 20, 21

Rack the MetroCluster IP configuration 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.

Follow the steps to *Install Hardware* under the *Install and Setup* instructions for your platform model in the [ONTAP hardware systems documentation](#).

4. Install the IP 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.
- Unique shelf IDs are highly recommended for each SAS disk shelf within each MetroCluster DR group to aid troubleshooting.



Do not cable disk shelves intended to contain unmirrored aggregates at this time. You must wait to deploy shelves intended for unmirrored aggregates until after the MetroCluster configuration is complete and only deploy them after using the `metrocluster modify -enable-unmirrored-aggr-deployment true` command.

Cable the MetroCluster IP switches

How to use the port tables with multiple MetroCluster IP configurations

You must understand how to use the information in the port tables to correctly generate your RCF files.

Before you begin

Review these considerations before using the tables:

- The following tables show the port usage for site A. The same cabling is used for site B.
- You cannot configure the switches with ports of different speeds (for example, a mix of 100 Gbps ports and 40 Gbps ports).
- Keep track of the MetroCluster port group (MetroCluster 1, MetroCluster 2, etc.). You'll need this information when using the RcfFileGenerator tool as described later in this configuration procedure.
- You should cable all of the nodes in the same way. If there are different port combination options available to cable the nodes, all nodes should use the same port combinations. For example, e1a on node1 and e1a on node2 should be attached to one switch. Similarly, the second port from both nodes should be attached to the second switch.
- The [RcfFileGenerator for MetroCluster IP](#) also provides a per-port cabling overview for each switch. Use this cabling overview to verify your cabling.

Cabling two MetroCluster configurations to the switches

When cabling more than one MetroCluster configuration to a switch, you cable each MetroCluster according to the appropriate table. For example, if you are cabling a FAS2750 and an AFF A700 to the same switch, you cable the FAS2750 as per "MetroCluster 1" in Table 1, and the AFF A700 as per "MetroCluster 2" or "MetroCluster 3" in Table 2. You cannot physically cable both the FAS2750 and the AFF A700 as "MetroCluster 1".

Cabling eight-node MetroCluster configurations

For MetroCluster configuration running ONTAP 9.8 and earlier, some procedures that are performed to transition an upgrade require the addition of a second four-node DR group to the configuration to create a temporary eight-node configuration. Beginning with ONTAP 9.9.1, permanent eight-node MetroCluster configurations are supported.

About this task

For eight-node configurations, you use the same method as described above. Instead of a second MetroCluster, you are cabling an additional four-node DR group.

For example, your configuration includes the following:

- Cisco 3132Q-V switches
- MetroCluster 1: FAS2750 platforms
- MetroCluster 2: AFF A700 platforms (these platforms are being added as a second four-node DR group)

Steps

1. For MetroCluster 1, cable the Cisco 3132Q-V switches using the table for the FAS2750 platform and the rows for MetroCluster 1 interfaces.
2. For MetroCluster 2 (the second DR group), cable the Cisco 3132Q-V switches using the table for the AFF A700 platform and the rows for MetroCluster 2 interfaces.

Platform port assignments for Cisco 3132Q-V switches in a MetroCluster IP configuration

The port usage in a MetroCluster IP configuration depends on the switch model and platform type.

Review these guidelines before using the tables:

- If you configure the switch for MetroCluster FC to IP transition, port 5, port 6, port 13, or port 14 can be used to connect the local cluster interfaces of the MetroCluster FC node. Refer to the [RcfFileGenerator](#) and the generated cabling files for more details on cabling this configuration. For all other connections, you can use the port usage assignments listed in the tables.

Choose the correct cabling table for your configuration

Use the following table to determine which cabling table you should follow.

If your system is...	Use this cabling table...
FAS2750, AFF A220	Cisco 3132Q-V platform port assignments (group 1)
FAS9000, AFF A700	Cisco 3132Q-V platform port assignments (group 2)
AFF A800, ASA A800	Cisco 3132Q-V platform port assignments (group 3)

Cisco 3132Q-V platform port assignments (group 1)

Review the platform port assignments to cable a FAS2750 or AFF A220 system to a Cisco 3132Q-V switch:

Switch Port	Port use	FAS2750 AFF A220	
		IP_Switch_x_1	IP_Switch_x_2
1 - 6	Unused	disabled	
7	ISL, Local Cluster native speed / 40G / 100G	ISL, Local Cluster	
8			
9/1	MetroCluster 1, Shared Cluster and MetroCluster interface	e0a	e0b
9/2-4		disabled	
10/1		e0a	e0b
10/2-4		disabled	
11/1	MetroCluster 2, Shared Cluster and MetroCluster interface	e0a	e0b
11/2-4		disabled	
12/1		e0a	e0b
12/2-4		disabled	
13/1	MetroCluster 3, Shared Cluster and MetroCluster interface	e0a	e0b
13/2-4		disabled	
14/1		e0a	e0b
14/2-4		disabled	
15	ISL, MetroCluster native speed 40G	ISL, MetroCluster	
16			
17			
18			
19			
20			
21/1-4	ISL, MetroCluster breakout mode 10G	ISL, MetroCluster	
22/1-4			
23/1-4			
24/1-4			
25 - 32	Unused	disabled	

Cisco 3132Q-V platform port assignments (group 2)

Review the platform port assignments to cable a FAS9000 or AFF A700 system to a Cisco 3132Q-V switch:

Switch Port	Port use	FAS9000 AFF A700	
		IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4e / e8a
2			
3	MetroCluster 2, Local Cluster interface	e4a	e4e / e8a
4			
5	MetroCluster 3, Local Cluster interface	e4a	e4e / e8a
6			
7	ISL, Local Cluster native speed 40G	ISL, Local Cluster	
8			
9	MetroCluster 1, MetroCluster interface	e5a	e5b
10			
11	MetroCluster 2, MetroCluster interface	e5a	e5b
12			
13	MetroCluster 3, MetroCluster interface	e5a	e5b
14			
15	ISL, MetroCluster native speed 40G	ISL, MetroCluster	
16			
17			
18			
19			
20			
21/1-4	ISL, MetroCluster breakout mode 10G	ISL, MetroCluster	
22/1-4			
23/1-4			
24/1-4			
25 - 32	Unused	disabled	

Cisco 3132Q-V platform port assignments (group 3)

Review the platform port assignments to cable an AFF A800 or ASA A800 system to a Cisco 3132Q-V switch:

Switch Port	Port use	AFF A800 ASA A800	
		IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e0a	e1a
2			
3	MetroCluster 2, Local Cluster interface	e0a	e1a
4			
5	MetroCluster 3, Local Cluster interface	e0a	e1a
6			
7	ISL, Local Cluster native speed 40G	ISL, Local Cluster	
8			
9	MetroCluster 1, MetroCluster interface	e0b	e1b
10			
11	MetroCluster 2, MetroCluster interface	e0b	e1b
12			
13	MetroCluster 3, MetroCluster interface	e0b	e1b
14			
15	ISL, MetroCluster native speed 40G	ISL, MetroCluster	
16			
17			
18			
19			
20			
21/1-4	ISL, MetroCluster breakout mode 10G	ISL, MetroCluster	
22/1-4			
23/1-4			
24/1-4			
25 - 32	Unused	disabled	

Platform port assignments for Cisco 3232C or 36-port Cisco 9336C switches in a MetroCluster IP configuration

The port usage in a MetroCluster IP configuration depends on the switch model and platform type.

Review the following considerations before using the configuration tables:

- The tables in this section are for Cisco 3232C switches or 36-port Cisco 9336C-FX2 switches that aren't connecting NS224 storage.

If you have a 12-port Cisco 9336C-FX2 switch, use the tables in [Platform port assignments for 12-port Cisco 9336C-FX2 switches](#).

If you have a 36-port Cisco 9336C-FX2 switch and at least one MetroCluster configuration or DR group is connecting NS224 shelves to the MetroCluster switch, use the tables in [Platform port assignments for a 36-port Cisco 9336C-FX2 switch connecting NS224 storage](#).

- The following tables show the port usage for site A. The same cabling is used for site B.

- You cannot configure the switches with ports of different speeds (for example, a mix of 100 Gbps ports and 40 Gbps ports).
- If you are configuring a single MetroCluster with the switches, use the **MetroCluster 1** port group.

Keep track of the MetroCluster port group (MetroCluster 1, MetroCluster 2, MetroCluster 3, or MetroCluster 4). You will need it when using the RcfFileGenerator tool as described later in this configuration procedure.

- The RcfFileGenerator for MetroCluster IP also provides a per-port cabling overview for each switch.

Use this cabling overview to verify your cabling.

- RCF file version v2.10 or later is required for 25G breakout mode for MetroCluster ISLs.
- ONTAP 9.13.1 or later and RCF file version 2.00 are required to use a platform other than FAS8200 or AFF A300 in the "MetroCluster 4" group.



The RCF file version is different to the version of the RCFfilegenerator tool used to generate the file. For example, you can generate an RCF file version 2.00 using RCFfilegenerator v1.6c.

Choose the correct cabling table for your configuration

Use the following table to determine which cabling table you should follow.

If your system is...	Use this cabling table...
AFF A150, ASA A150 FAS2750, AFF A220 FAS500f, AFF C250, ASA C250 AFF A250, ASA A250	Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 1)
AFF A20	Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 2)
AFF A30, AFF C30 FAS50 AFF C60	The table you follow depends on whether you are using a 25G (group 3a) or 100G (group 3b) Ethernet card. <ul style="list-style-type: none"> • Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 3a - 25G) • Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 3b - 100G)
FAS8200, AFF A300	Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 4)
AFF A320 FAS8300, AFF C400, ASA C400, FAS8700 AFF A400, ASA A400	Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 5)
AFF A50	Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 6)

If your system is...	Use this cabling table...
FAS9000, AFF A700 AFF C800, ASA C800, AFF A800, ASA A800 FAS9500, AFF A900, ASA A900	Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 7)
FAS70, AFF A70 AFF C80 FAS90, AFF A90 AFF A1K	Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 8)

Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 1)

Review the platform port assignments to cable an AFF A150, ASA A150, FAS2750, AFF A220, FAS500f, AFF C250, ASA C250, AFF A250, or ASA A250 system to a Cisco 3232C or 9336C-FX2 switch:

Switch Port	Port use	AFF A150 ASA A150 FAS2750 AFF A220		FAS500f AFF C250 ASA C250 AFF A250 ASA A250	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1 - 6	Unused	disabled		disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster	
8					
9/1	MetroCluster 1, Shared Cluster and MetroCluster interface	e0a	e0b	e0c	e0d
9/2-4		disabled		disabled	
10/1		e0a	e0b	e0c	e0d
10/2-4		disabled		disabled	
11/1	MetroCluster 2, Shared Cluster and MetroCluster interface	e0a	e0b	e0c	e0d
11/2-4		disabled		disabled	
12/1		e0a	e0b	e0c	e0d
12/2-4		disabled		disabled	
13/1	MetroCluster 3, Shared Cluster and MetroCluster interface	e0a	e0b	e0c	e0d
13/2-4		disabled		disabled	
14/1		e0a	e0b	e0c	e0d
14/2-4		disabled		disabled	
15	ISL, MetroCluster native speed 40G / 100G				
16					
17		ISL, MetroCluster		ISL, MetroCluster	
18					
19					
20					
21/1-4	ISL, MetroCluster breakout mode 10G / 25G				
22/1-4		ISL, MetroCluster		ISL, MetroCluster	
23/1-4					
24/1-4					
25/1	MetroCluster 1, Shared Cluster and MetroCluster interface	e0a	e0b	e0c	e0d
25/2-4		disabled		disabled	
26/1		e0a	e0b	e0c	e0d
26/2-4		disabled		disabled	
27 - 32	Unused	disabled		disabled	
33 - 36	Unused (Cisco 9336C-FX2 only)	disabled		disabled	

Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 2)

Review the platform port assignments to cable an AFF A20 system to a Cisco 3232C or 9336C-FX2 switch:

Switch Port	Port use	AFF A20	
		IP_Switch_x_1	IP_Switch_x_2
1/1	MetroCluster 1, Local Cluster interface	e2a	e4a
1/2-4		disabled	
2/1		e2a	e4a
2/2-4		disabled	
3/1	MetroCluster 2, Local Cluster interface	e2a	e4a
3/2-4		disabled	
4/1		e2a	e4a
4/2-4		disabled	
5/1	MetroCluster 3, Local Cluster interface	e2a	e4a
5/2-4		disabled	
6/1		e2a	e4a
6/2-4		disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
8			
9/1	MetroCluster 1, MetroCluster interface	e2b	e4b
9/2-4		disabled	
10/1		e2b	e4b
10/2-4		disabled	
11/1	MetroCluster 2, MetroCluster interface	e2b	e4b
11/2-4		disabled	
12/1		e2b	e4b
12/2-4		disabled	
13/1	MetroCluster 3, MetroCluster interface	e2b	e4b
13/2-4		disabled	
14/1		e2b	e4b
14/2-4		disabled	
15	ISL, MetroCluster native speed 40G / 100G	ISL, MetroCluster	
16			
17			
18			
19			
20			
21/1-4	ISL, MetroCluster breakout mode 10G / 25G	ISL, MetroCluster	
22/1-4			
23/1-4			
24/1-4			
25/1	MetroCluster 4, MetroCluster interface	e2b	e4b
25/2-4		disabled	
26/1		e2b	e4b
26/2-4		disabled	
27 - 28	Unused	disabled	
29/1	MetroCluster 4, Local Cluster interface	e2a	e4a
29/2-4		disabled	
30/1		e2a	e4a
30/2-4		disabled	
25 - 32	Unused	disabled	
33 - 36	Unused (Cisco 9336C-FX2 only)	disabled	

Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 3a)

Review the platform port assignments to cable an AFF A30, AFF C30, AFF C60, or FAS50 system to a Cisco 3232C or 9336C-FX2 switch using a four-port 25G Ethernet card.



This configuration requires a four-port 25G Ethernet card in slot 4 to connect the local cluster and HA interfaces.

Switch Port	Port use	AFF C30 (25G Cluster/HA) AFF A30 (25G Cluster/HA)		FAS50 (25G Cluster/HA)		AFF C60 (25G Cluster/HA)	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1/1	MetroCluster 1, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
1/2-4		disabled		disabled		disabled	
2/1		e4a	e4b	e4a	e4b	e4a	e4b
2/2-4		disabled		disabled		disabled	
3/1	MetroCluster 2, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
3/2-4		disabled		disabled		disabled	
4/1		e4a	e4b	e4a	e4b	e4a	e4b
4/2-4		disabled		disabled		disabled	
5/1	MetroCluster 3, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
5/2-4		disabled		disabled		disabled	
6/1		e4a	e4b	e4a	e4b	e4a	e4b
6/2-4		disabled		disabled		disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
8		ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
9	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
10		e2a	e2b	e2a	e2b	e2a	e2b
11	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
12		e2a	e2b	e2a	e2b	e2a	e2b
13	MetroCluster 3, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
14		e2a	e2b	e2a	e2b	e2a	e2b
15	ISL, MetroCluster native speed 40G / 100G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
16		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
17		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
18		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
19		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
20		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
21/1-4	ISL, MetroCluster breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
22/1-4		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
23/1-4		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
24/1-4		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
25	MetroCluster 4, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
26		e2a	e2b	e2a	e2b	e2a	e2b
27 - 28	Unused	disabled		disabled		disabled	
29/1	MetroCluster 4, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
29/2-4		disabled		disabled		disabled	
30/1		e4a	e4b	e4a	e4b	e4a	e4b
30/2-4		disabled		disabled		disabled	
25 - 32	Unused	disabled		disabled		disabled	
33 - 36	Unused (Cisco 9336C-FX2 only)	disabled		disabled		disabled	

Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 3b)

Review the platform port assignments to cable an AFF A30, AFF C30, AFF C60, or FAS50 system to a Cisco 3232C or 9336C-FX2 switch using a two-port 100G Ethernet card.



This configuration requires a two-port 100G Ethernet card in slot 4 to connect the local cluster and HA interfaces.

Switch Port	Port use	AFF C30 (100G Cluster/HA) AFF A30 (100G Cluster/HA)		FAS50 (100G Cluster/HA)		AFF C60 (100G Cluster/HA)	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
		1	MetroCluster 1, Local Cluster interface	e4a	e4b	e4a	e4b
2		e4a	e4b	e4a	e4b	e4a	e4b
3	MetroCluster 2, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
4		e4a	e4b	e4a	e4b	e4a	e4b
5	MetroCluster 3, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
6		e4a	e4b	e4a	e4b	e4a	e4b
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
8							
9	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
10		e2a	e2b	e2a	e2b	e2a	e2b
11	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
12		e2a	e2b	e2a	e2b	e2a	e2b
13	MetroCluster 3, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
14		e2a	e2b	e2a	e2b	e2a	e2b
15	ISL, MetroCluster native speed 40G / 100G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
16							
17							
18							
19							
20							
21/1-4	ISL, MetroCluster breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
22/1-4							
23/1-4							
24/1-4							
25	MetroCluster 4, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
26		e2a	e2b	e2a	e2b	e2a	e2b
27 - 28	Unused	disabled		disabled		disabled	
29	MetroCluster 4, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
30		e4a	e4b	e4a	e4b	e4a	e4b
25 - 32	Unused	disabled		disabled		disabled	
33 - 36	Unused (Cisco 9336C-FX2 only)	disabled		disabled		disabled	

Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 4)

Review the platform port assignments to cable a FAS8200 or AFF A300 system to a Cisco 3232C or 9336C-FX2 switch:

Switch Port	Port use	FAS8200 AFF A300	
		IP_Switch_x_1	IP_Switch_x_2
1/1	MetroCluster 1, Local Cluster interface	e0a	e0b
1/2-4		disabled	
2/1		e0a	e0b
2/2-4		disabled	
3/1	MetroCluster 2, Local Cluster interface	e0a	e0b
3/2-4		disabled	
4/1		e0a	e0b
4/2-4		disabled	
5/1	MetroCluster 3, Local Cluster interface	e0a	e0b
5/2-4		disabled	
6/1		e0a	e0b
6/2-4		disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
8			
9/1	MetroCluster 1, MetroCluster interface	e1a	e1b
9/2-4		disabled	
10/1		e1a	e1b
10/2-4		disabled	
11/1	MetroCluster 2, MetroCluster interface	e1a	e1b
11/2-4		disabled	
12/1		e1a	e1b
12/2-4		disabled	
13/1	MetroCluster 3, MetroCluster interface	e1a	e1b
13/2-4		disabled	
14/1		e1a	e1b
14/2-4		disabled	
15	ISL, MetroCluster native speed 40G / 100G	ISL, MetroCluster	
16			
17			
18			
19			
20			
21/1-4	ISL, MetroCluster breakout mode 10G / 25G	ISL, MetroCluster	
22/1-4			
23/1-4			
24/1-4			
25/1	MetroCluster 4, MetroCluster interface	e1a	e1b
25/2-4		disabled	
26/1		e1a	e1b
26/2-4		disabled	
27 - 28	Unused	disabled	
29/1	MetroCluster 4, Local Cluster interface	e0a	e0b
29/2-4		disabled	
30/1		e0a	e0b
30/2-4		disabled	
25 - 32	Unused	disabled	
33 - 36	Unused (Cisco 9336C-FX2 only)	disabled	

If you are upgrading from older RCF files, the cabling configuration might be using ports in the "MetroCluster 4" group (ports 25/26 and 29/30).

Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 5)

Review the platform port assignments to cable an AFF A320, FAS8300, AFF C400, ASA C400, FAS8700, AFF A400, or ASA A400 system to a Cisco 3232C or 9336C-FX2 switch:

Switch Port	Port use	AFF A320		FAS8300 AFF C400 ASA C400 FAS8700		AFF A400 ASA A400	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e0a	e0d	e0c	e0d	e3a	e3b
2							
3	MetroCluster 2, Local Cluster interface	e0a	e0d	e0c	e0d	e3a	e3b
4							
5	MetroCluster 3, Local Cluster interface	e0a	e0d	e0c	e0d	e3a	e3b
6							
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
8							
9	MetroCluster 1, MetroCluster interface	e0g	e0h	e1a	e1b	e1a	e1b
10							
11	MetroCluster 2, MetroCluster interface	e0g	e0h	e1a	e1b	e1a	e1b
12							
13	MetroCluster 3, MetroCluster interface	e0g	e0h	e1a	e1b	e1a	e1b
14							
15	ISL, MetroCluster native speed 40G / 100G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
16							
17							
18							
19							
20							
21/1-4	ISL, MetroCluster breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
22/1-4							
23/1-4							
24/1-4							
25	MetroCluster 4, MetroCluster interface	e0g	e0h	e1a	e1b	e1a	e1b
26							
27 - 28	Unused	disabled		disabled		disabled	
29	MetroCluster 4, Local Cluster interface	e0a	e0d	e0c	e0d	e3a	e3b
30							
31 - 32	Unused	disabled		disabled		disabled	
33 - 34	Unused (Cisco 9336C-FX2 only)	disabled		disabled		disabled	



Using ports in the "MetroCluster 4" group requires ONTAP 9.13.1 or later.

Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 6)

Review the platform port assignments to cable an AFF A50 system to a Cisco 3232C or 9336C-FX2 switch:

Switch Port	Port use	AFF A50	
		IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4b
2		e4a	e4b
3	MetroCluster 2, Local Cluster interface	e4a	e4b
4		e4a	e4b
5	MetroCluster 3, Local Cluster interface	e4a	e4b
6		e4a	e4b
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
8			
9	MetroCluster 1, MetroCluster interface	e2a	e2b
10		e2a	e2b
11	MetroCluster 2, MetroCluster interface	e2a	e2b
12		e2a	e2b
13	MetroCluster 3, MetroCluster interface	e2a	e2b
14		e2a	e2b
15	ISL, MetroCluster native speed 40G / 100G	ISL, MetroCluster	
16			
17			
18			
19			
20			
21/1-4	ISL, MetroCluster breakout mode 10G / 25G	ISL, MetroCluster	
22/1-4			
23/1-4			
24/1-4			
25	MetroCluster 4, MetroCluster interface	e2a	e2b
26		e2a	e2b
27 - 28	Unused	disabled	
29	MetroCluster 4, Local Cluster interface	e4a	e4b
30		e4a	e4b
25 - 32	Unused	disabled	
33 - 36	Unused (Cisco 9336C-FX2 only)	disabled	

Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 7)

Review the platform port assignments to cable a FAS9000, AFF A700, AFF C800, ASA C800, AFF A800, ASA A800, FAS9500, AFF A900, or ASA A900 system to a Cisco 3232C or 9336C-FX2 switch:

Switch Port	Port use	FAS9000 AFF A700		AFF C800 ASA C800 AFF A800 ASA A800		FAS9500 AFF A900 ASA A900	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4e / e8a	e0a	e1a	e4a	e4b(e) / e8a Note 1
2							
3							
4	MetroCluster 2, Local Cluster interface	e4a	e4e / e8a	e0a	e1a	e4a	e4b(e) / e8a Note 1
5							
6	MetroCluster 3, Local Cluster interface	e4a	e4e / e8a	e0a	e1a	e4a	e4b(e) / e8a Note 1
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
8							
9	MetroCluster 1, MetroCluster interface	e5a	e5b	e0b	e1b	e5b	e7b
10							
11	MetroCluster 2, MetroCluster interface	e5a	e5b	e0b	e1b	e5b	e7b
12							
13	MetroCluster 3, MetroCluster interface	e5a	e5b	e0b	e1b	e5b	e7b
14							
15	ISL, MetroCluster native speed 40G / 100G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
16							
17							
18							
19							
20							
21/1-4	ISL, MetroCluster breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
22/1-4							
23/1-4							
24/1-4							
25	MetroCluster 4, MetroCluster interface	e5a	e5b	e0b	e1b	e5b	e7b
26							
27 - 28	Unused	disabled		disabled		disabled	
29	MetroCluster 4, Local Cluster interface	e4a	e4e / e8a	e0a	e1a	e4a	e4b(e) / e8a Note 1
30							
31 - 32	Unused	disabled		disabled		disabled	
33 - 34	Unused (Cisco 9336C-FX2 only)	disabled		disabled		disabled	

Note 1: Use either ports e4a and e4e or e4a and e8a if you are using an X91440A adapter (40Gbps). Use either ports e4a and e4b or e4a and e8a if you are using an X91153A adapter (100Gbps).



Using ports in the "MetroCluster 4" group requires ONTAP 9.13.1 or later.

Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 8)

Review the platform port assignments to cable an AFF A70, FAS70, AFF C80, FAS90, AFF A90, or AFF A1K system to a Cisco 3232C or 9336C-FX2 switch:

Switch Port	Port use	FAS70 AFF A70		AFF C80		FAS90 AFF A90		AFF A1K	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e1a	e7a	e1a	e7a	e1a	e7a	e1a	e7a
2									
3									
4	MetroCluster 2, Local Cluster interface	e1a	e7a	e1a	e7a	e1a	e7a	e1a	e7a
5									
6									
7	MetroCluster 3, Local Cluster interface	e1a	e7a	e1a	e7a	e1a	e7a	e1a	e7a
8									
9									
10	ISL, Local Cluster native speed / 100G	ISL, Local Cluster							
11									
12									
13	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b	e2b	e3b	e2b	e3b
14									
15									
16	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2b	e3b	e2b	e3b
17									
18									
19	MetroCluster 3, MetroCluster interface	e2a	e2b	e2a	e2b	e2b	e3b	e2b	e3b
20									
21									
22	ISL, MetroCluster native speed 40G / 100G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
23									
24									
25									
26	ISL, MetroCluster breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
27									
28									
29									
30	MetroCluster 4, MetroCluster interface	e2a	e2b	e2a	e2b	e2b	e3b	e2b	e3b
31									
32									
33	Unused	disabled		disabled		disabled		disabled	
34									
35									
36	MetroCluster 4, Local Cluster interface	e1a	e7a	e1a	e7a	e1a	e7a	e1a	e7a
37									
38									
39	Unused	disabled		disabled		disabled		disabled	
40									
41									
42	Unused (Cisco 9336C-FX2 only)	disabled		disabled		disabled		disabled	
43									
44									

Platform port assignments for 12-port Cisco 9336C-FX2 switches in a MetroCluster IP configuration

The port usage in a MetroCluster IP configuration depends on the switch model and platform type.

Review the following considerations before using the configuration tables:

- The tables in this section are for 12-port Cisco 9336C-FX2 switches.

If you have a 36-port Cisco 9336C-FX2 switch that isn't connecting NS224 shelves, use the tables in [Platform port assignments for Cisco 3232C or 36-port Cisco 9336C-FX2 switches](#).

If you have a 36-port Cisco 9336C-FX2 switch and at least one MetroCluster configuration or DR group is connecting NS224 shelves to the MetroCluster switch, use the tables in [Platform port assignments for a 36-port Cisco 9336C-FX2 switch connecting NS224 storage](#).



The 12-port Cisco 9336C-FX2 switch doesn't support connecting NS224 shelves to the MetroCluster switch.

- The following tables show the port usage for site A. The same cabling is used for site B.
- You cannot configure the switches with ports of different speeds (for example, a mix of 100 Gbps ports and 40 Gbps ports).
- If you are configuring a single MetroCluster with the switches, use the **MetroCluster 1** port group.

Keep track of the MetroCluster port group (MetroCluster 1, MetroCluster 2). You'll need it when using the RcfFileGenerator tool as described later in this configuration procedure.

- The RcfFileGenerator for MetroCluster IP also provides a per-port cabling overview for each switch.

Choose the correct cabling table for your configuration

Use the following table to determine which cabling table you should follow.

If your system is...	Use this cabling table...
AFF A150, ASA A150 FAS500f AFF C250, ASA C250 AFF A250, ASA A250	Cisco 9336C-FX2 12-port platform port assignments (group 1)
AFF A20	Cisco 9336C-FX2 12-port platform port assignments (group 2)
AFF A30, AFF C30 FAS50 AFF C60	The table you follow depends on whether you are using a 25G (group 3a) or 100G (group 3b) Ethernet card. <ul style="list-style-type: none"> • Cisco 9336C-FX2 12-port platform port assignments (group 3a - 25G) • Cisco 9336C-FX2 12-port platform port assignments (group 3b - 100G)
FAS8300, AFF C400, ASA C400, FAS8700 AFF A400, ASA A400	Cisco 9336C-FX2 12-port platform port assignments (group 4)
AFF A50	Cisco 9336C-FX2 12-port platform port assignments (group 5)
AFF C800, ASA C800, AFF A800, ASA A800 FAS9500, AFF A900, ASA A900	Cisco 9336C-FX2 12-port platform port assignments (group 6)
FAS70, AFF A70 AFF C80 FAS90, AFF A90 AFF A1K	Cisco 9336C-FX2 12-port platform port assignments (group 7)

Cisco 9336C-FX2 12-port platform port assignments (group 1)

Review the platform port assignments to cable an AFF A150, ASA A150, FAS500f, AFF C250, ASA C250, AFF A250, or ASA A250 system to a 12-port Cisco 9336C-FX2 switch:

Switch Port	Port use	AFF A150 ASA A150		FAS500f AFF C250 ASA C250 AFF A250 ASA A250	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1-4	Unused	disabled		disabled	
5-6	Ports disallowed to use	blocked		blocked	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster	
8					
9/1	MetroCluster 1, Shared Cluster and MetroCluster interface	e0a	e0b	e0c	e0d
9/2-4		disabled		disabled	
10/1		e0a	e0b	e0c	e0d
10/2-4		disabled		disabled	
11/1	MetroCluster 2, Shared Cluster and MetroCluster interface	e0a	e0b	e0c	e0d
11/2-4		disabled		disabled	
12/1		e0a	e0b	e0c	e0d
12/2-4		disabled		disabled	
13-18	Ports disallowed to use	blocked		blocked	
19	ISL, MetroCluster native speed 40G / 100G (note 1)	ISL, MetroCluster		ISL, MetroCluster	
20					
21/1-4	ISL, MetroCluster breakout mode 10G / 25G (note 1)	ISL, MetroCluster		ISL, MetroCluster	
22/1-4					
23-36	Ports disallowed to use	blocked		blocked	

Note 1: You can only configure ports 19 and 20 **or** ports 21 and 22. If you use ports 19 and 20 first, then ports 21 and 22 are blocked. If you use ports 21 and 22 first, then ports 19 and 20 are blocked.

Cisco 9336C-FX2 12-port platform port assignments (group 2)

Review the platform port assignments to cable an AFF A20 system to a 12-port Cisco 9336C-FX2 switch:

Switch Port	Port use	AFF A20	
		IP_Switch_x_1	IP_Switch_x_2
1/1	MetroCluster 1, Local Cluster interface	e2a	e4a
1/2-4		disabled	
2/1		e2a	e4a
2/2-4		disabled	
3/1	MetroCluster 2, Local Cluster interface	e2a	e4a
3/2-4		disabled	
4/1		e2a	e4a
4/2-4		disabled	
5-6	Ports disallowed to use	blocked	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
8			
9/1	MetroCluster 1, MetroCluster interface	e2b	e4b
9/2-4		disabled	
10/1		e2b	e4b
10/2-4		disabled	
11/1	MetroCluster 2, MetroCluster interface	e2b	e4b
11/2-4		disabled	
12/1		e2b	e4b
12/2-4		disabled	
13-18	Ports disallowed to use	blocked	
19	ISL, MetroCluster native speed 40G / 100G (note 1)	ISL, MetroCluster	
20			
21/1-4	ISL, MetroCluster breakout mode 10G / 25G (note 1)	ISL, MetroCluster	
22/1-4			
23-36	Ports disallowed to use	blocked	

Note 1: You can only configure ports 19 and 20 **or** ports 21 and 22. If you use ports 19 and 20 first, then ports 21 and 22 are blocked. If you use ports 21 and 22 first, then ports 19 and 20 are blocked.

Cisco 9336C-FX2 12-port platform port assignments (group 3a)

Review the platform port assignments to cable an AFF A30, AFF C30, AFF C60, or FAS50 system to a 12-port Cisco 9336C-FX2 switch using a four-port 25G Ethernet card.



This configuration requires a four-port 25G Ethernet card in slot 4 to connect the local cluster and HA interfaces.

Switch Port	Port use	AFF C30 (25G Cluster/HA) AFF A30 (25G Cluster/HA)		FAS50 (25G Cluster/HA)		AFF C60 (25G Cluster/HA)	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1/1	MetroCluster 1, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
1/2-4		disabled		disabled		disabled	
2/1		e4a	e4b	e4a	e4b	e4a	e4b
2/2-4		disabled		disabled		disabled	
3/1	MetroCluster 2, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
3/2-4		disabled		disabled		disabled	
4/1		e4a	e4b	e4a	e4b	e4a	e4b
4/2-4		disabled		disabled		disabled	
5-6	Ports disallowed to use	blocked		blocked		blocked	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
8		ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
9	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
10		e2a	e2b	e2a	e2b	e2a	e2b
11	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
12		e2a	e2b	e2a	e2b	e2a	e2b
13-18	Ports disallowed to use	blocked		blocked		blocked	
19	ISL, MetroCluster native speed 40G / 100G (note 1)	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
20		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
21/1-4	ISL, MetroCluster breakout mode 10G / 25G (note 1)	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
22/1-4		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
23-36	Ports disallowed to use	blocked		blocked		blocked	

Note 1: You can only configure ports 19 and 20 **or** ports 21 and 22. If you use ports 19 and 20 first, then ports 21 and 22 are blocked. If you use ports 21 and 22 first, then ports 19 and 20 are blocked.

Cisco 9336C-FX2 12-port platform port assignments (group 3b)

Review the platform port assignments to cable an AFF A30, AFF C30, AFF C60, or FAS50 system to a 12-port Cisco 9336C-FX2 switch using a two-port 100G Ethernet card.



This configuration requires a two-port 100G Ethernet card in slot 4 to connect the local cluster and HA interfaces.

Switch Port	Port use	AFF C30 (100G Cluster/HA) AFF A30 (100G Cluster/HA)		FAS50 (100G Cluster/HA)		AFF C60 (100G Cluster/HA)	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
2		e4a	e4b	e4a	e4b	e4a	e4b
3	MetroCluster 2, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
4		e4a	e4b	e4a	e4b	e4a	e4b
5-6	Ports disallowed to use	blocked		blocked		blocked	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
8		ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
9	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
10		e2a	e2b	e2a	e2b	e2a	e2b
11	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
12		e2a	e2b	e2a	e2b	e2a	e2b
13-18	Ports disallowed to use	blocked		blocked		blocked	
19	ISL, MetroCluster native speed 40G / 100G (note 1)	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
20		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
21/1-4	ISL, MetroCluster breakout mode 10G / 25G (note 1)	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
22/1-4		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
23-36	Ports disallowed to use	blocked		blocked		blocked	

Note 1: You can only configure ports 19 and 20 **or** ports 21 and 22. If you use ports 19 and 20 first, then ports 21 and 22 are blocked. If you use ports 21 and 22 first, then ports 19 and 20 are blocked.

Cisco 9336C-FX2 12-port platform port assignments (group 4)

Review the platform port assignments to cable an FAS8300, AFF C400, ASA C400, FAS8700, AFF A400, or ASA A400 system to a 12-port Cisco 9336C-FX2 switch:

Switch Port	Port use	FAS8300 AFF C400 ASA C400 FAS8700		AFF A400 ASA A400	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e0c	e0d	e3a	e3b
2					
3	MetroCluster 2, Local Cluster interface	e0c	e0d	e3a	e3b
4					
5-6	Ports disallowed to use	blocked		blocked	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster	
8					
9	MetroCluster 1, MetroCluster interface	e1a	e1b	e1a	e1b
10					
11	MetroCluster 2, MetroCluster interface	e1a	e1b	e1a	e1b
12					
13-18	Ports disallowed to use	blocked		blocked	
19	ISL, MetroCluster native speed 40G / 100G (note 1)	ISL, MetroCluster		ISL, MetroCluster	
20					
21/1-4	ISL, MetroCluster breakout mode 10G / 25G (note 1)	ISL, MetroCluster		ISL, MetroCluster	
22/1-4					
23-36	Ports disallowed to use	blocked		blocked	

Note 1: You can only configure ports 19 and 20 **or** ports 21 and 22. If you use ports 19 and 20 first, then ports 21 and 22 are blocked. If you use ports 21 and 22 first, then ports 19 and 20 are blocked.

Cisco 9336C-FX2 12-port platform port assignments (group 5)

Review the platform port assignments to cable an AFF A50 system to a 12-port Cisco 9336C-FX2 switch:

Switch Port	Port use	AFF A50	
		IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4b
2		e4a	e4b
3	MetroCluster 2, Local Cluster interface	e4a	e4b
4		e4a	e4b
5-6	Ports disallowed to use	blocked	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
8			
9	MetroCluster 1, MetroCluster interface	e2a	e2b
10		e2a	e2b
11	MetroCluster 2, MetroCluster interface	e2a	e2b
12		e2a	e2b
13-18	Ports disallowed to use	blocked	
19	ISL, MetroCluster native speed 40G / 100G (note 1)	ISL, MetroCluster	
20			
21/1-4	ISL, MetroCluster breakout mode 10G / 25G (note 1)	ISL, MetroCluster	
22/1-4			
23-36	Ports disallowed to use	blocked	

Note 1: You can only configure ports 19 and 20 **or** ports 21 and 22. If you use ports 19 and 20 first, then ports 21 and 22 are blocked. If you use ports 21 and 22 first, then ports 19 and 20 are blocked.

Cisco 9336C-FX2 12-port platform port assignments (group 6)

Review the platform port assignments to cable an AFF C800, ASA C800, AFF A800, ASA A800, FAS9500, AFF A900, or ASA A900 system to a 12-port Cisco 9336C-FX2 switch:

Switch Port	Port use	AFF C800 ASA C800 AFF A800 ASA A800		FAS9500 AFF A900 ASA A900	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e0a	e1a	e4a	e4b(e) / e8a (note 2)
2					
3	MetroCluster 2, Local Cluster interface	e0a	e1a	e4a	e4b(e) / e8a (note 2)
4					
5-6	Ports disallowed to use	blocked		blocked	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster	
8					
9	MetroCluster 1, MetroCluster interface	e0b	e1b	e5b	e7b
10					
11	MetroCluster 2, MetroCluster interface	e0b	e1b	e5b	e7b
12					
13-18	Ports disallowed to use	blocked		blocked	
19	ISL, MetroCluster native speed 40G / 100G (note 1)	ISL, MetroCluster		ISL, MetroCluster	
20					
21/1-4	ISL, MetroCluster breakout mode 10G / 25G (note 1)	ISL, MetroCluster		ISL, MetroCluster	
22/1-4					
23-36	Ports disallowed to use	blocked		blocked	

Note 1: You can only configure ports 19 and 20 **or** ports 21 and 22. If you use ports 19 and 20 first, then ports 21 and 22 are blocked. If you use ports 21 and 22 first, then ports 19 and 20 are blocked.

Note 2: Use either ports e4a and e4e or e4a and e8a if you are using an X91440A adapter (40Gbps). Use either ports e4a and e4b or e4a and e8a if you are using an X91153A adapter (100Gbps).

Cisco 9336C-FX2 12-port platform port assignments (group 7)

Review the platform port assignments to cable an AFF A70, FAS70, AFF C80, FAS90, AFF A90, or AFF A1K system to a 12-port Cisco 9336C-FX2 switch:

Switch Port	Port use	FAS70 AFF A70		AFF C80		FAS90 AFF A90		AFF A1K	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e1a	e7a	e1a	e7a	e1a	e7a	e1a	e7a
2									
3	MetroCluster 2, Local Cluster interface	e1a	e7a	e1a	e7a	e1a	e7a	e1a	e7a
4									
5-6	Ports disallowed to use	blocked		blocked		blocked		blocked	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster							
8									
9	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b	e2b	e3b	e2b	e3b
10									
11	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2b	e3b	e2b	e3b
12									
13-18	Ports disallowed to use	blocked		blocked		blocked		blocked	
19	ISL, MetroCluster native speed 40G / 100G (note 1)	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
20									
21/1-4	ISL, MetroCluster breakout mode 10G / 25G (note 1)	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
22/1-4									
23-36	Ports disallowed to use	blocked		blocked		blocked		blocked	

Note 1: You can only configure ports 19 and 20 **or** ports 21 and 22. If you use ports 19 and 20 first, then ports 21 and 22 are blocked. If you use ports 21 and 22 first, then ports 19 and 20 are blocked.

Platform port assignments for a 36-port Cisco 9336C-FX2 switch connecting NS224 storage in a MetroCluster IP configuration

The port usage in a MetroCluster IP configuration depends on the switch model and platform type.

Review the following considerations before using the configuration tables:

- The tables in this section are for 36-port Cisco 9336C-FX2 switches when at least one MetroCluster configuration or DR group is connecting NS224 shelves to the MetroCluster switch.

If you have a 36-port Cisco 9336C-FX2 switch and you don't plan to connect NS224 storage to the switch, use the tables in [Platform port assignments for Cisco 3232C or 36-port Cisco 9336C-FX2 switches](#).

If you have a 12-port Cisco 9336C-FX2 switch, use the tables in [Platform port assignments for 12-port Cisco 9336C-FX2 switches](#).



The 12-port Cisco 9336C-FX2 switch doesn't support connecting NS224 shelves to the MetroCluster switch.

- When you cable a Cisco 9336C-FX2 switch connecting NS224 storage, you can only have a maximum of two MetroCluster configurations or DR groups. At least one MetroCluster configuration or DR group must be connecting NS224 shelves to the MetroCluster switch.

If one of your MetroCluster configurations or DR groups doesn't connect NS224 shelves to the MetroCluster switch, follow the [cabling tables for controllers not connecting switch-attached NS224 shelves](#).

- The RcfFileGenerator only shows eligible platforms when the first platform is selected.
- If you are expanding a regular MetroCluster configuration that doesn't support switch-attached NS224 storage and you want to add a MetroCluster configuration that has switch-attached NS224 storage, review the following information:
 - In the [RcfFileGenerator](#) tool, you must select the platform that supports switch-attached NS224 storage first, then select the platform that doesn't support switch-attached NS224 storage.

This means that when you generate the RCF for your new configuration, you might need to select the platforms in the RcfFileGenerator tool in the following order:

1. Select the platform that supports switch-attached NS224 storage from the **second drop-down field**.
2. Select the platform that doesn't support switch-attached NS224 storage from the **first drop-down field**.

- Connecting one eight-node or two four-node MetroCluster configurations requires ONTAP 9.14.1 or later.

Choose the correct cabling table for your configuration

Review the correct port assignments table for your configuration. There are two sets of cabling tables in this section:

- [Cabling tables for controllers connecting switch-attached NS224 shelves](#)
- [Cabling tables for controllers not connecting switch-attached NS224 shelves](#)

Controllers connecting switch-attached NS224 shelves

Determine which port assignments table you should follow for controllers connecting switch-attached NS224 shelves.

Platform	Use this cabling table...
AFF C30, AFF A30 AFF C60	The table you follow depends on whether you are using a 25G (group 1a) or 100G (group 1b) Ethernet card. <ul style="list-style-type: none">• Cisco 9336C-FX2 switch connecting NS224 storage platform port assignments (group 1a - 25G)• Cisco 9336C-FX2 switch connecting NS224 storage platform port assignments (group 1b - 100G)
AFF A320 AFF C400, ASA C400 AFF A400, ASA A400	Cisco 9336C-FX2 switch connecting NS224 storage platform port assignments (group 2)
AFF A50	Cisco 9336C-FX2 switch connecting NS224 storage platform port assignments (group 3)
AFF A700 AFF C800, ASA C800, AFF A800 AFF A900, ASA A900	Cisco 9336C-FX2 switch connecting NS224 storage platform port assignments (group 4)
AFF A70 AFF C80 AFF A90 AFF A1K	Cisco 9336C-FX2 switch connecting NS224 storage platform port assignments (group 5)

Cisco 9336C-FX2 switch connecting NS224 storage platform port assignments (group 1a)

Review the platform port assignments to cable an AFF A30, AFF C30, or AFF C60 system that is connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch using a four-port 25G Ethernet card.



This configuration requires a four-port 25G Ethernet card in slot 4 to connect the local cluster and HA interfaces.

Controllers connecting switch-attached shelves					
Switch Port	Port Use	AFF C30 (25G Cluster/HA) AFF A30 (25G Cluster/HA)		AFF C60 (25G Cluster/HA)	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1/1	MetroCluster 1, Local Cluster interface	e4a	e4b	e4a	e4b
1/2-4		disabled		disabled	
2/1		e4a	e4b	e4a	e4b
2/2-4		disabled		disabled	
3/1	MetroCluster 2, Local Cluster interface	e4a	e4b	e4a	e4b
3/2-4		disabled		disabled	
4/1		e4a	e4b	e4a	e4b
4/2-4		disabled		disabled	
5	Storage shelf 1 (9)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
6		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster	
8					
9	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b
10		e2a	e2b	e2a	e2b
11	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b
12		e2a	e2b	e2a	e2b
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster	
14					
15					
16					
17	MetroCluster 1, Ethernet Storage Interface	e3a	e3b	e3a	e3b
18					
19	MetroCluster 2, Ethernet Storage Interface	e3a	e3b	e3a	e3b
20					
21	Storage shelf 2 (8)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
22		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
23	Storage shelf 3 (7)				
24					
25	Storage shelf 4 (6)				
26					
27	Storage shelf 5 (5)				
28					
29	Storage shelf 6 (4)				
30					
31	Storage shelf 7 (3)				
32					
33	Storage shelf 8 (2)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
34		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
35	Storage shelf 9 (1)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
36		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b

Cisco 9336C-FX2 switch connecting NS224 storage platform port assignments (group 1b)

Review the platform port assignments to cable an AFF A30, AFF C30, or AFF C60 system that is connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch using a two-port 100G Ethernet card.



This configuration requires a two-port 100G Ethernet card in slot 4 to connect the local cluster and HA interfaces.

Controllers connecting switch-attached shelves					
Switch Port	Port Use	AFF C30 (100G Cluster/HA) AFF A30 (100G Cluster/HA)		AFF C60 (100G Cluster/HA)	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4b	e4a	e4b
2		e4a	e4b	e4a	e4b
3	MetroCluster 2, Local Cluster interface	e4a	e4b	e4a	e4b
4		e4a	e4b	e4a	e4b
5	Storage shelf 1 (9)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
6		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster	
8					
9	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b
10		e2a	e2b	e2a	e2b
11	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b
12		e2a	e2b	e2a	e2b
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster	
14					
15					
16					
17	MetroCluster 1, Ethernet Storage Interface	e3a	e3b	e3a	e3b
18		e3a	e3b	e3a	e3b
19	MetroCluster 2, Ethernet Storage Interface	e3a	e3b	e3a	e3b
20		e3a	e3b	e3a	e3b
21	Storage shelf 2 (8)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
22		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
23	Storage shelf 3 (7)				
24					
25	Storage shelf 4 (6)				
26					
27	Storage shelf 5 (5)				
28					
29	Storage shelf 6 (4)				
30					
31	Storage shelf 7 (3)				
32					
33	Storage shelf 8 (2)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
34		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
35	Storage shelf 9 (1)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
36		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b

Cisco 9336C-FX2 switch connecting NS224 storage platform port assignments (group 2)

Review the platform port assignments to cable an AFF A320, AFF C400, ASA C400, AFF A400, or ASA A400 system that is connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch:

Controllers connecting switch-attached shelves							
Switch Port	Port Use	AFF A320		AFF C400 ASA C400		AFF A400 ASA A400	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e0a	e0d	e0c	e0d	e3a	e3b
2							
3	MetroCluster 2, Local Cluster interface	e0a	e0d	e0c	e0d	e3a	e3b
4							
5	Storage shelf 1 (9)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
6		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
8							
9	MetroCluster 1, MetroCluster interface	e0g	e0h	e1a	e1b	e1a	e1b
10							
11	MetroCluster 2, MetroCluster interface	e0g	e0h	e1a	e1b	e1a	e1b
12							
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
14							
15							
16							
17	MetroCluster 1, Ethernet Storage Interface	e0c	e0f	e4a	e4b / e5b	e0c	e0d / e5b
18							
19	MetroCluster 2, Ethernet Storage Interface	e0c	e0f	e4a	e4b / e5b	e0c	e0d / e5b
20							
21	Storage shelf 2 (8)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
22		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
23	Storage shelf 3 (7)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
24		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
25	Storage shelf 4 (6)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
26		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
27	Storage shelf 5 (5)	NSM-1, e0a	NSM-1, e0b				
28		NSM-2, e0a	NSM-2, e0b				
29	Storage shelf 6 (4)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
30		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
31	Storage shelf 7 (3)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
32		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
33	Storage shelf 8 (2)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
34		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
35	Storage shelf 9 (1)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
36		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b

Cisco 9336C-FX2 switch connecting NS224 storage platform port assignments (group 3)

Review the platform port assignments to cable an AFF A50 system that is connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch:

Controllers connecting switch-attached shelves			
Switch Port	Port Use	AFF A50	
		IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4b
2		e4a	e4b
3	MetroCluster 2, Local Cluster interface	e4a	e4b
4		e4a	e4b
5	Storage shelf 1 (9)	NSM-1, e0a	NSM-1, e0b
6		NSM-2, e0a	NSM-2, e0b
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
8			
9	MetroCluster 1, MetroCluster interface	e2a	e2b
10		e2a	e2b
11	MetroCluster 2, MetroCluster interface	e2a	e2b
12		e2a	e2b
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster	
14			
15			
16			
17	MetroCluster 1, Ethernet Storage Interface	e3a	e3b
18			
19	MetroCluster 2, Ethernet Storage Interface	e3a	e3b
20			
21	Storage shelf 2 (8)	NSM-1, e0a	NSM-1, e0b
22		NSM-2, e0a	NSM-2, e0b
23	Storage shelf 3 (7)	NSM-1, e0a	NSM-1, e0b
24		NSM-2, e0a	NSM-2, e0b
25	Storage shelf 4 (6)	NSM-1, e0a	NSM-1, e0b
26		NSM-2, e0a	NSM-2, e0b
27	Storage shelf 5 (5)		
28			
29	Storage shelf 6 (4)	NSM-1, e0a	NSM-1, e0b
30		NSM-2, e0a	NSM-2, e0b
31	Storage shelf 7 (3)	NSM-1, e0a	NSM-1, e0b
32		NSM-2, e0a	NSM-2, e0b
33	Storage shelf 8 (2)	NSM-1, e0a	NSM-1, e0b
34		NSM-2, e0a	NSM-2, e0b
35	Storage shelf 9 (1)	NSM-1, e0a	NSM-1, e0b
36		NSM-2, e0a	NSM-2, e0b

Cisco 9336C-FX2 switch connecting NS224 storage platform port assignments (group 4)

Review the platform port assignments to cable an AFF A700, AFF C800, ASA C800, AFF A800, AFF A900, or ASA A900 system that is connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch:

Controllers connecting switch-attached shelves							
Switch Port	Port Use	AFF A700		AFF C800 ASA C800 AFF A800		AFF A900 ASA A900	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4e / e8a	e0a	e1a	e4a	e4b(e) / e8a Note 1
2							
3	MetroCluster 2, Local Cluster interface	e4a	e4e / e8a	e0a	e1a	e4a	e4b(e) / e8a Note 1
4							
5	Storage shelf 1 (9)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
6		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
8							
9	MetroCluster 1, MetroCluster interface	e5a	e5b	e0b	e1b	e5b	e7b
10							
11	MetroCluster 2, MetroCluster interface	e5a	e5b	e0b	e1b	e5b	e7b
12							
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
14							
15							
16							
17	MetroCluster 1, Ethernet Storage Interface	e3a	e3b / e7b	e5a	e5b / e3b	e3a (option 1) e2a (option 2)	e3b (option 1) e10b (option 2)
18							
19	MetroCluster 2, Ethernet Storage Interface	e3a	e3b / e7b	e5a	e5b / e3b	e3a (option 1) e2a (option 2)	e3b (option 1) e10b (option 2)
20							
21	Storage shelf 2 (8)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
22		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
23	Storage shelf 3 (7)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
24		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
25	Storage shelf 4 (6)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
26		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
27	Storage shelf 5 (5)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
28		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
29	Storage shelf 6 (4)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
30		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
31	Storage shelf 7 (3)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
32		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
33	Storage shelf 8 (2)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
34		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
35	Storage shelf 9 (1)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
36		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b

Note 1: Use either ports e4a and e4e or e4a and e8a if you are using an X91440A adapter (40Gbps). Use either ports e4a and e4b or e4a and e8a if you are using an X91153A adapter (100Gbps).

Cisco 9336C-FX2 switch connecting NS224 storage platform port assignments (group 5)

Review the platform port assignments to cable an AFF A70, AFF C80, AFF A90, or AFF A1K system that is connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch:

Controllers connecting switch-attached shelves									
Switch Port	Port Use	AFF A70		AFF C80		AFF A90		AFF A1K	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e1a	e7a	e1a	e7a	e1a	e7a	e1a	e7a
2									
3	MetroCluster 2, Local Cluster interface	e1a	e7a	e1a	e7a	e1a	e7a	e1a	e7a
4									
5	Storage shelf 1 (9)	NSM-1, e0a	NSM-1, e0b						
6		NSM-2, e0a	NSM-2, e0b						
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster							
8									
9	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b	e2b	e3b	e2b	e3b
10									
11	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2b	e3b	e2b	e3b
12									
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
14									
15									
16									
17	MetroCluster 1, Ethernet Storage Interface	e8a (option 1)	e8b (option 1)						
18		e11a (option 2)	e11b (option 2)						
19	MetroCluster 2, Ethernet Storage Interface	e8a (option 1)	e8b (option 1)						
20		e11a (option 2)	e11b (option 2)						
21	Storage shelf 2 (8)	NSM-1, e0a	NSM-1, e0b						
22		NSM-2, e0a	NSM-2, e0b						
23	Storage shelf 3 (7)	NSM-1, e0a	NSM-1, e0b						
24		NSM-2, e0a	NSM-2, e0b						
25	Storage shelf 4 (6)	NSM-1, e0a	NSM-1, e0b						
26		NSM-2, e0a	NSM-2, e0b						
27	Storage shelf 5 (5)	NSM-1, e0a	NSM-1, e0b						
28		NSM-2, e0a	NSM-2, e0b						
29	Storage shelf 6 (4)	NSM-1, e0a	NSM-1, e0b						
30		NSM-2, e0a	NSM-2, e0b						
31	Storage shelf 7 (3)	NSM-1, e0a	NSM-1, e0b						
32		NSM-2, e0a	NSM-2, e0b						
33	Storage shelf 8 (2)	NSM-1, e0a	NSM-1, e0b						
34		NSM-2, e0a	NSM-2, e0b						
35	Storage shelf 9 (1)	NSM-1, e0a	NSM-1, e0b						
36		NSM-2, e0a	NSM-2, e0b						

Controllers not connecting switch-attached NS224 shelves

Determine which port assignments table you should follow for controllers that are not connecting switch-attached NS224 shelves.

Platform	Use this cabling table...
AFF A150, ASA A150 FAS2750, AFF A220	Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 6)
AFF A20	Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 7)
FAS500f AFF C250, ASA C250 AFF A250, ASA A250	Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 8)

Platform	Use this cabling table...
AFF C30, AFF A30 FAS50 AFF C60	The table you follow depends on whether you are using a 25G (group 9a) or 100G (group 9b) Ethernet card. <ul style="list-style-type: none"> • Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 9a) • Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 9b)
FAS8200, AFF A300	Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 10)
AFF A320 FAS8300, AFF C400, ASA C400, FAS8700 AFF A400, ASA A400	Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 11)
AFF A50	Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 12)
FAS9000, AFF A700 AFF C800, ASA C800, AFF A800, ASA A800 FAS9500, AFF A900, ASA A900	Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 13)
FAS70, AFF A70 AFF C80 FAS90, AFF A90 AFF A1K	Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 14)

Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 6)

Review the platform port assignments to cable an AFF A150, ASA A150, FAS2750, or AFF A220 system that isn't connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch:

Controllers not connecting switch-attached shelves			
Switch Port	Port Use	AFF A150 ASA A150 FAS2750 AFF A220	
		IP_Switch_x_1	IP_Switch_x_2
1 - 6	Unused	disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
8			
9/1	MetroCluster 1, Shared Cluster and MetroCluster interface	e0a	e0b
9/2-4		disabled	
10/1		e0a	e0b
10/2-4		disabled	
11/1	MetroCluster 2, Shared Cluster and MetroCluster interface	e0a	e0b
11/2-4		disabled	
12/1		e0a	e0b
12/2-4		disabled	
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster	
14			
15			
16			
17-36	Unused	disabled	

Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 7)

Review the platform port assignments to cable an AFF A20 system that isn't connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch:

Controllers not connecting switch-attached shelves			
Switch Port	Port Use	AFF A20	
		IP_Switch_x_1	IP_Switch_x_2
1/1	MetroCluster 1, Local Cluster interface	e2a	e4a
1/2-4		disabled	
2/1		e2a	e4a
2/2-4		disabled	
3/1	MetroCluster 2, Local Cluster interface	e2a	e4a
3/2-4		disabled	
4/1		e2a	e4a
4/2-4		disabled	
5-6	Unused	disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
8			
9/1	MetroCluster 1, MetroCluster interface	e2b	e4b
9/2-4		disabled	
10/1		e2b	e4b
10/2-4		disabled	
11/1	MetroCluster 2, MetroCluster interface	e2b	e4b
11/2-4		disabled	
12/1		e2b	e4b
12/2-4		disabled	
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster	
14			
15			
16			
17-36	Unused	disabled	

Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 8)

Review the platform port assignments to cable a FAS500f, AFF C250, ASA C250, AFF A250, or ASA A250 system that isn't connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch:

Controllers not connecting switch-attached shelves			
Switch Port	Port Use	FAS500f AFF C250 ASA C250 AFF A250 ASA A250	
		IP_Switch_x_1	IP_Switch_x_2
1 - 6	Unused	disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
8			
9/1	MetroCluster 1, Shared Cluster and MetroCluster interface	e0c	e0d
9/2-4		disabled	
10/1		e0c	e0d
10/2-4		disabled	
11/1	MetroCluster 2, Shared Cluster and MetroCluster interface	e0c	e0d
11/2-4		disabled	
12/1		e0c	e0d
12/2-4		disabled	
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster	
14			
15			
16			
17-36	Unused	disabled	

Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 9a)

Review the platform port assignments to cable an AFF A30, AFF C30, AFF C60, or FAS50 system that isn't connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch using a four-port 25G Ethernet card:



This configuration requires a four-port 25G Ethernet card in slot 4 to connect the local cluster and HA interfaces.

Controllers not connecting switch-attached shelves							
Switch Port	Port use	AFF C30 (25G Cluster/HA) AFF A30 (25G Cluster/HA)		FAS50 (25G Cluster/HA)		AFF C60 (25G Cluster/HA)	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1/1	MetroCluster 1, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
1/2-4		disabled		disabled		disabled	
2/1		e4a	e4b	e4a	e4b	e4a	e4b
2/2-4		disabled		disabled		disabled	
3/1	MetroCluster 2, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
3/2-4		disabled		disabled		disabled	
4/1		e4a	e4b	e4a	e4b	e4a	e4b
4/2-4		disabled		disabled		disabled	
5-6	Unused	disabled		disabled		disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
8		ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
9	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
10	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
11	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
12	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
14		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
15		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
16		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
17-36	Unused	disabled		disabled		disabled	

Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 9b)

Review the platform port assignments to cable an AFF A30, AFF C30, AFF C60, or FAS50 system that isn't connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch using a two-port 100G Ethernet card:



This configuration requires a two-port 100G Ethernet card in slot 4 to connect the local cluster and HA interfaces.

Controllers not connecting switch-attached shelves							
Switch Port	Port use	AFF C30 (100G Cluster/HA) AFF A30 (100G Cluster/HA)		FAS50 (100G Cluster/HA)		AFF C60 (100G Cluster/HA)	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
2		e4a	e4b	e4a	e4b	e4a	e4b
3	MetroCluster 2, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
4		e4a	e4b	e4a	e4b	e4a	e4b
5-6	Unused	disabled		disabled		disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
8		ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
9	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
10	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
11	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
12	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
14		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
15		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
16		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
17-36	Unused	disabled		disabled		disabled	

Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 10)

Review the platform port assignments to cable a FAS8200 or AFF A300 system that isn't connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch:

Controllers not connecting switch-attached shelves			
Switch Port	Port Use	FAS8200 AFF A300	
		IP_Switch_x_1	IP_Switch_x_2
1/1	MetroCluster 1, Local Cluster interface	e0a	e0b
1/2-4		disabled	
2/1		e0a	e0b
2/2-4		disabled	
3/1	MetroCluster 2, Local Cluster interface	e0a	e0b
3/2-4		disabled	
4/1		e0a	e0b
4/2-4		disabled	
5-6	Unused	disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
8			
9/1	MetroCluster 1, MetroCluster interface	e1a	e1b
9/2-4		disabled	
10/1		e1a	e1b
10/2-4		disabled	
11/1	MetroCluster 2, MetroCluster interface	e1a	e1b
11/2-4		disabled	
12/1		e1a	e1b
12/2-4		disabled	
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster	
14			
15			
16			
17-36	Unused	disabled	

Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 11)

Review the platform port assignments to cable an AFF A320, FAS8300, AFF C400, ASA C400, FAS8700, AFF A400, or ASA A400 system that isn't connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch:

Controllers not connecting switch-attached shelves							
Switch Port	Port Use	AFF A320		FAS8300 AFF C400 ASA C400 FAS8700		AFF A400 ASA A400	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e0a	e0d	e0c	e0d	e3a	e3b
2							
3	MetroCluster 2, Local Cluster interface	e0a	e0d	e0c	e0d	e3a	e3b
4							
5-6	Unused	disabled		disabled		disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
8							
9	MetroCluster 1, MetroCluster interface	e0g	e0h	e1a	e1b	e1a	e1b
10							
11	MetroCluster 2, MetroCluster interface	e0g	e0h	e1a	e1b	e1a	e1b
12							
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
14							
15							
16							
17-36	Unused	disabled		disabled		disabled	

Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 12)

Review the platform port assignments to cable an AFF A50 system that isn't connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch:

Controllers not connecting switch-attached shelves			
Switch Port	Port use	AFF A50	
		IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4b
2			
3	MetroCluster 2, Local Cluster interface	e4a	e4b
4			
5-6	Unused	disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
8			
9	MetroCluster 1, MetroCluster interface	e2a	e2b
10			
11	MetroCluster 2, MetroCluster interface	e2a	e2b
12			
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster	
14			
15			
16			
17-36	Unused	disabled	

Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 13)

Review the platform port assignments to cable a FAS9000, AFF A700, AFF C800, ASA C800, AFF A800, ASA A800, FAS9500, AFF A900, or ASA A900 system that isn't connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch:

Controllers not connecting switch-attached shelves							
Switch Port	Port Use	FAS9000 AFF A700		AFF C800 ASA C800 AFF A800 ASA A800		FAS9500 AFF A900 ASA A900	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4e / e8a	e0a	e1a	e4a	e4b(e) / e8a Note 1
2							
3	MetroCluster 2, Local Cluster interface	e4a	e4e / e8a	e0a	e1a	e4a	e4b(e) / e8a Note 1
4							
5-6	Unused	disabled		disabled		disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
8							
9	MetroCluster 1, MetroCluster interface	e5a	e5b	e0b	e1b	e5b	e7b
10							
11	MetroCluster 2, MetroCluster interface	e5a	e5b	e0b	e1b	e5b	e7b
12							
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
14							
15							
16							
17-36	Unused	disabled		disabled		disabled	

Note 1: Use either ports e4a and e4e or e4a and e8a if you are using an X91440A adapter (40Gbps). Use either ports e4a and e4b or e4a and e8a if you are using an X91153A adapter (100Gbps).

Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 14)

Review the platform port assignments to cable an AFF A70, FAS70, AFF C80, FAS90, AFF A90, or AFF A1K system that isn't connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch:

Controllers not connecting switch-attached shelves									
Switch Port	Port Use	FAS70 AFF A70		AFF C80		FAS90 AFF A90		AFF A1K	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e1a	e7a	e1a	e7a	e1a	e7a	e1a	e7a
2									
3	MetroCluster 2, Local Cluster interface	e1a	e7a	e1a	e7a	e1a	e7a	e1a	e7a
4									
5-6	Unused	disabled		disabled		disabled		disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster							
8									
9	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b	e2b	e3b	e2b	e3b
10									
11	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2b	e3b	e2b	e3b
12									
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
14									
15									
16									
17-36	Unused	disabled		disabled		disabled		disabled	

Platform port assignments for Broadcom supported BES-53248 IP switches in a MetroCluster IP configuration

The port usage in a MetroCluster IP configuration depends on the switch model and platform type.

Review the following considerations before using the configuration tables:

- You cannot use the switches with remote ISL ports of different speeds (for example, a 25 Gbps port connected to a 10 Gbps ISL port).
- If you configure the switch for MetroCluster FC to IP Transition, the following ports are used depending on

the target platform that you choose:

Target platform	Port
FAS500f, AFF C250, ASA C250, AFF A250, ASA A250, FAS8300, AFF C400, ASA C400, AFF A400, ASA A400, or FAS8700 platforms	ports 1 - 6, 10Gbps
FAS8200 or AFF A300 platforms	ports 3 - 4 and 9 - 12, 10Gbps

- AFF A320 systems configured with Broadcom BES-53248 switches might not support all features.

Any configuration or feature that requires that the local cluster connections are connected to a switch is not supported. For example, the following configurations and procedures are not supported:

- Eight-node MetroCluster configurations
- Transitioning from MetroCluster FC to MetroCluster IP configurations
- Refreshing a four-node MetroCluster IP configuration (ONTAP 9.8 and later)

Choose the correct cabling table for your configuration

Use the following table to determine which cabling table you should follow.

If your system is...	Use this cabling table...
AFF A150, ASA A150 FAS2750 AFF A220	Broadcom BES-53248 platform port assignments (group 1)
FAS500f AFF C250, ASA C250 AFF A250, ASA A250	Broadcom BES-53248 platform port assignments (group 2)
AFF A20	Broadcom BES-53248 platform port assignments (group 3)
AFF C30, AFF A30 FAS50 AFF C60	Broadcom BES-53248 platform port assignments (group 4)
FAS8200, AFF A300	Broadcom BES-53248 platform port assignments (group 5)
AFF A320	Broadcom BES-53248 platform port assignments (group 6)

If your system is...	Use this cabling table...
FAS8300	Broadcom BES-53248 platform port assignments (group 7)
AFF C400, ASA C400	
AFF A400, ASA A400	
FAS8700	

Broadcom BES-53248 platform port assignments (group 1)

Review the platform port assignments to cable an AFF A150, ASA A150, FAS2750, or AFF A220 system to a Broadcom BES-53248 switch:

Physical Port	Port use	AFF A150 ASA A150 FAS2750 AFF A220	
		IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Shared Cluster and MetroCluster interface	e0a	e0b
2			
3	MetroCluster 2, Shared Cluster and MetroCluster interface	e0a	e0b
4			
5-8	Unused	disabled	
9	MetroCluster 3, Shared Cluster and MetroCluster interface	e0a	e0b
10			
11	MetroCluster 4, Shared Cluster and MetroCluster interface	e0a	e0b
12			
13	ISL, MetroCluster native speed 10G / 25G	ISL, MetroCluster	
14			
15			
16			
..	Ports not licensed (17 - 54)		
53	ISL, MetroCluster, native speed 40G / 100G (Note 1)	ISL, MetroCluster	
54			
55	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
56			

- **Note 1:** Using these ports requires an additional license.
- If both MetroCluster configurations are using the same the platform, NetApp recommends selecting group "MetroCluster 3" for one configuration and group "MetroCluster 4" for the other configuration. If the platforms are different, then you must select "MetroCluster 3" or "MetroCluster 4" for the first configuration, and "MetroCluster 1" or "MetroCluster 2" for the second configuration.

Broadcom BES-53248 platform port assignments (group 2)

Review the platform port assignments to cable a FAS500f, AFF C250, ASA C250, AFF A250, or ASA A250 system to a Broadcom BES-53248 switch:

Physical Port	Port use	FAS500f AFF C250 ASA C250 AFF A250 ASA A250	
		IP_Switch_x_1	IP_Switch_x_2
1 - 4	Unused	disabled	
5	MetroCluster 1, Shared Cluster and MetroCluster interface	e0c	e0d
6			
7	MetroCluster 2, Shared Cluster and MetroCluster interface	e0c	e0d
8			
9	MetroCluster 3, Shared Cluster and MetroCluster interface	e0c	e0d
10			
11	MetroCluster 4, Shared Cluster and MetroCluster interface	e0c	e0d
12			
13	ISL, MetroCluster native speed 10G / 25G	ISL, MetroCluster	
14			
15			
16			
..	Ports not licensed (17 - 54)		
53	ISL, MetroCluster, native speed 40G / 100G (Note 1)	ISL, MetroCluster	
54			
55	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
56			

- **Note 1:** Using these ports requires an additional license.
- If both MetroCluster configurations are using the same the platform, NetApp recommends selecting group "MetroCluster 3" for one configuration and group "MetroCluster 4" for the other configuration. If the platforms are different, then you must select "MetroCluster 3" or "MetroCluster 4" for the first configuration, and "MetroCluster 1" or "MetroCluster 2" for the second configuration.

Broadcom BES-53248 platform port assignments (group 3)

Review the platform port assignments to cable an AFF A20 system to a Broadcom BES-53248 switch:

Physical Port	Port use	AFF A20	
		IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e2a	e4a
2			
3	MetroCluster 2, Local Cluster interface	e2a	e4a
4			
5	MetroCluster 1, MetroCluster interface	e2b	e4b
6			
7	MetroCluster 2, MetroCluster interface	e2b	e4b
8			
9 - 12	Unused	disabled	
13	ISL, MetroCluster native speed 10G / 25G	ISL, MetroCluster	
14			
15			
16			
17	MetroCluster 3, Local Cluster interface (note 1)	e2a	e4a
18			
19	MetroCluster 3, MetroCluster interface (note 1)	e2b	e4b
20			
21	MetroCluster 4, Local Cluster interface (note 1)	e2a	e4a
22			
23	MetroCluster 4, MetroCluster interface (note 1)	e2b	e4b
24			
..	Ports not licensed (25 - 54)		
53	ISL, MetroCluster, native speed 40G / 100G (note 1)	ISL, MetroCluster	
54			
55	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
56			

- **Note 1:** Using these ports requires an additional license.

Broadcom BES-53248 platform port assignments (group 4)

Review the platform port assignments to cable an AFF A30, AFF C30, AFF C60, or FAS50 system to a Broadcom BES-53248 switch using a four-port 25G Ethernet card.



- This configuration requires a four-port 25G Ethernet card in slot 4 to connect the local cluster and HA interfaces.
- This configuration requires a QSFP-to-SFP+ adapter in the card on the controller to support a 25Gbps network speed.

Physical Port	Port use	AFF C30 (25G Cluster/HA) AFF A30 (25G Cluster/HA)		FAS50 (25G Cluster/HA)		AFF C60 (25G Cluster/HA)	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
2							
3	MetroCluster 2, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
4							
5	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
6							
7	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
8							
9 - 12	Unused	disabled		disabled		disabled	
13	ISL, MetroCluster native speed 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
14							
15							
16							
17	MetroCluster 3, Local Cluster interface (note 1)	e4a	e4b	e4a	e4b	e4a	e4b
18							
19	MetroCluster 3, MetroCluster interface (note 1)	e2a	e2b	e2a	e2b	e2a	e2b
20							
21	MetroCluster 4, Local Cluster interface (note 1)	e4a	e4b	e4a	e4b	e4a	e4b
22							
23	MetroCluster 4, MetroCluster interface (note 1)	e2a	e2b	e2a	e2b	e2a	e2b
24							
..	Ports not licensed (25 - 54)						
53	ISL, MetroCluster, native speed 40G / 100G (note 1)	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
54							
55	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
56							

- **Note 1:** Using these ports requires an additional license.

Broadcom BES-53248 platform port assignments (group 5)

Review the platform port assignments to cable a FAS8200 or AFF A300 system to a Broadcom BES-53248 switch:

Physical Port	Port use	FAS8200 AFF A300	
		IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e0a	e0b
2			
3	MetroCluster 2, Local Cluster interface	e0a	e0b
4			
5	MetroCluster 1, MetroCluster interface	e1a	e1b
6			
7	MetroCluster 2, MetroCluster interface	e1a	e1b
8			
9 - 12	Unused	disabled	
13	ISL, MetroCluster native speed 10G / 25G	ISL, MetroCluster	
14			
15			
16			
..	Ports not licensed (17 - 54)		
53	ISL, MetroCluster, native speed 40G / 100G (note 1)	ISL, MetroCluster	
54			
55	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
56			

- **Note 1:** Using these ports requires an additional license.

Broadcom BES-53248 platform port assignments (group 6)

Review the platform port assignments to cable an AFF A320 system to a Broadcom BES-53248 switch:

Physical Port	Port use	AFF A320	
		IP_Switch_x_1	IP_Switch_x_2
1 - 12	Ports not used (Note 2)	disabled	
13	ISL, MetroCluster native speed 10G / 25G	ISL, MetroCluster	
14			
15			
16			
..	Ports not licensed (17 - 54)		
53	ISL, MetroCluster, native speed 40G / 100G (see Note 1)	ISL, MetroCluster	
54			
55	MetroCluster 1, MetroCluster interface (Note 2)	e0g	e0h
56			

- **Note 1:** Using these ports requires an additional license.
- **Note 2:** Only a single four-node MetroCluster using AFF A320 systems can be connected to the switch.

Features that require a switched cluster are not supported in this configuration. This includes the MetroCluster FC to IP transition and tech refresh procedures.

Broadcom BES-53248 platform port assignments (group 7)

Review the platform port assignments to cable a FAS8300, AFF C400, ASA C400, AFF A400, ASA A400, or FAS8700 system to a Broadcom BES-53248 switch:

Physical Port	Port use	FAS8300 AFF C400 ASA C400 FAS8700		AFF A400 ASA A400	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1 - 12	Ports not used (see Note 2)	disabled		disabled	
13	ISL, MetroCluster native speed 10G / 25G	ISL, MetroCluster		ISL, MetroCluster	
14					
15					
16					
..	Ports not licensed (17 - 48)				
49	MetroCluster 5, Local Cluster interface (Note 1)	e0c	e0d	e3a	e3b
50					
51	MetroCluster 5, MetroCluster interface (Note 1)	e1a	e1b	e1a	e1b
52					
53	ISL, MetroCluster, native speed 40G / 100G (Note 1)	ISL, MetroCluster		ISL, MetroCluster	
54					
55	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster	
56					

- **Note 1:** Using these ports requires an additional license.

- **Note 2:** Only a single four-node MetroCluster using AFF A320 systems can be connected to the switch.

Features that require a switched cluster are not supported in this configuration. This includes the MetroCluster FC to IP transition and tech refresh procedures.

Platform port assignments for NVIDIA supported SN2100 IP switches in a MetroCluster IP configuration

The port usage in a MetroCluster IP configuration depends on the switch model and platform type.

Review the following considerations before using the configuration tables:

- Connecting an eight-node or two four-node MetroCluster configurations requires ONTAP 9.14.1 or later and RCF file version 2.00 or later.



The RCF file version is different to the version of the RCFfilegenerator tool used to generate the file. For example, you can generate an RCF file version 2.00 using RCFfilegenerator v1.6c.

- If you cable multiple MetroCluster configurations then follow the respective table. For example:
 - If you cable two four-node MetroCluster configurations of type AFF A700, then connect the first MetroCluster shown as "MetroCluster 1", and the second MetroCluster shown as "MetroCluster 2" in the AFF A700 table.



Ports 13 and 14 can be used in native speed mode supporting 40 Gbps and 100 Gbps, or in breakout mode to support 4 × 25 Gbps or 4 × 10 Gbps. If they use native speed mode they are represented as ports 13 and 14. If they use breakout mode, either 4 × 25 Gbps or 4 × 10 Gbps, then they are represented as ports 13s0-3 and 14s0-3.

The following sections describe the physical cabling outline. You can also refer to the [RcfFileGenerator](#) for detailed cabling information.

Choose the correct cabling table for your configuration

Use the following table to determine which cabling table you should follow.

If your system is...	Use this cabling table...
AFF A150, ASA A150 FAS500f AFF C250, ASA C250 AFF A250, ASAA250	NVIDIA SN2100 platform port assignments (group 1)
AFF A20	NVIDIA SN2100 platform port assignments (group 2)

If your system is...	Use this cabling table...
AFF C30, AFF A30 FAS50 AFF C60	The table you follow depends on whether you are using a 25G (group 3a) or 100G (group 3b) Ethernet card. <ul style="list-style-type: none"> • NVIDIA SN2100 platform port assignments (group 3a -25G) • NVIDIA SN2100 platform port assignments (group 3b -100G)
FAS8300 AFF C400, ASA C400 AFF A400, ASA A400 FAS8700 FAS9000, AFF A700	NVIDIA SN2100 platform port assignments (group 4)
AFF A50	NVIDIA SN2100 platform port assignments (group 5)
AFF C800, ASA C800 AFF A800, ASA A800 FAS9500 AFF A900, ASA A900	NVIDIA SN2100 platform port assignments (group 6)
FAS70, AFF A70 AFF C80 FAS90, AFF A90 AFF A1K	NVIDIA SN2100 platform port assignments (group 7)

NVIDIA SN2100 platform port assignments (group 1)

Review the platform port assignments to cable an AFF A150, ASA A150, FAS500f, AFF C250, ASA C250, AFF A250, or ASA A250 system to a NVIDIA SN2100 switch:

Switch Port	Port use	AFF A150 ASA A150		FAS500F AFF C250 ASA C250 AFF A250 ASA A250	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1 - 6	Unused	disabled		disabled	
7s0	MetroCluster 1, Shared Cluster and MetroCluster interface	e0c	e0d	e0c	e0d
7s1-3		disabled		disabled	
8s0		e0c	e0d	e0c	e0d
8s1-3		disabled		disabled	
9s0	MetroCluster 2, Shared Cluster and MetroCluster interface	e0c	e0d	e0c	e0d
9s1-3		disabled		disabled	
10s0		e0c	e0d	e0c	e0d
10s1-3		disabled		disabled	
11s0	MetroCluster 3, Shared Cluster and MetroCluster interface	e0c	e0d	e0c	e0d
11s1-3		disabled		disabled	
12s0		e0c	e0d	e0c	e0d
12s1-3		disabled		disabled	
13 / 13s0-3	MetroCluster ISL 40/100G or 4x25G or 4x10G	ISL, MetroCluster		ISL, MetroCluster	
14 / 14s0-3		ISL, MetroCluster		ISL, MetroCluster	
15	ISL, Local Cluster	ISL, Local Cluster		ISL, Local Cluster	
16	100G	ISL, Local Cluster		ISL, Local Cluster	

NVIDIA SN2100 platform port assignments (group 2)

Review the platform port assignments to cable an AFF A20 system to a NVIDIA SN2100 switch:

Switch Port	Port use	AFF A20	
		IP_Switch_x_1	IP_Switch_x_2
1s0	MetroCluster 1, Local Cluster interface	e2a	e4a
s1s1-3		disabled	
2s0		e2a	e4a
2s1-3		disabled	
3s0	MetroCluster 2, Local Cluster interface	e2a	e4a
3s1-3		disabled	
4s0		e2a	e4a
4s1-3		disabled	
5s0	MetroCluster 3, Local Cluster interface	e2a	e4a
5s1-3		disabled	
6s0		e2a	e4a
6s1-3		disabled	
7	MetroCluster 1, MetroCluster interface	e2b	e4b
8			
9	MetroCluster 2, MetroCluster interface	e2b	e4b
10			
11	MetroCluster 3, MetroCluster interface	e2b	e4b
12			
13 / 13s0-3	MetroCluster ISL 40/100G or 4x25G or 4x10G	ISL, MetroCluster	
14 / 14s0-3			
15	ISL, Local Cluster 100G	ISL, Local Cluster	
16			

NVIDIA SN2100 platform port assignments (group 3a)

Review the platform port assignments to cable an AFF A30, AFF C30, AFF C60, or FAS50 system to a NVIDIA SN2100 switch using a four-port 25G Ethernet card:



This configuration requires a four-port 25G Ethernet card in slot 4 to connect the local cluster and HA interfaces.

Switch Port	Port use	AFF C30 (25G Cluster/HA) AFF A30 (25G Cluster/HA)		FAS50 (25G Cluster/HA)		AFF C60 (25G Cluster/HA)	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1s0	MetroCluster 1, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
3s1-3		disabled		disabled		disabled	
2s0		e4a	e4b	e4a	e4b	e4a	e4b
2s1-3		disabled		disabled		disabled	
3s0	MetroCluster 2, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
3s1-3		disabled		disabled		disabled	
4s0		e4a	e4b	e4a	e4b	e4a	e4b
4s1-3		disabled		disabled		disabled	
5s0	MetroCluster 3, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
5s1-3		disabled		disabled		disabled	
6s0		e4a	e4b	e4a	e4b	e4a	e4b
6s1-3		disabled		disabled		disabled	
7	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
8							
9	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
10	MetroCluster 3, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
11	MetroCluster 3, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
12	MetroCluster 3, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
13 / 13s0-3	MetroCluster ISL 40/100G or 4x25G or 4x10G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
14 / 14s0-3		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
15	ISL, Local Cluster 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
16	ISL, Local Cluster 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	

NVIDIA SN2100 platform port assignments (group 3b)

Review the platform port assignments to cable an AFF A30, AFF C30, AFF C60, or FAS50 system to a NVIDIA SN2100 switch using a two-port 100G Ethernet card:



This configuration requires a two-port 100G Ethernet card in slot 4 to connect the local cluster and HA interfaces.

Switch Port	Port use	AFF C30 (100G Cluster/HA) AFF A30 (100G Cluster/HA)		FAS50 (100G Cluster/HA)		AFF C60 (100G Cluster/HA)	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
2		e4a	e4b	e4a	e4b	e4a	e4b
3	MetroCluster 2, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
4		e4a	e4b	e4a	e4b	e4a	e4b
5	MetroCluster 3, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
6		e4a	e4b	e4a	e4b	e4a	e4b
7	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
8							
9	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
10	MetroCluster 3, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
11	MetroCluster 3, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
12	MetroCluster 3, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
13 / 13s0-3	MetroCluster ISL 40/100G or 4x25G or 4x10G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
14 / 14s0-3		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
15	ISL, Local Cluster 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
16	ISL, Local Cluster 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	

NVIDIA SN2100 platform port assignments (group 4)

Review the platform port assignments to cable a FAS8300, AFF C400, ASA C400, AFF A400, ASA A400, FAS8700, FAS9000, or AFF A700 system to a NVIDIA SN2100 switch:

Switch Port	Port use	FAS8300 AFF C400 ASA C400 FAS8700		AFF A400 ASA A400		FAS9000 AFF A700	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e0c	e0d	e3a	e3b	e4a	e4e / e8a Note 1
2							
3	MetroCluster 2, Local Cluster interface	e0c	e0d	e3a	e3b	e4a	e4e / e8a Note 1
4							
5	MetroCluster 3, Local Cluster interface	e0c	e0d	e3a	e3b	e4a	e4e / e8a Note 1
6							
7	MetroCluster 1, MetroCluster interface	e1a	e1b	e1a	e1b	e5a	e5b
8							
9	MetroCluster 2, MetroCluster interface	e1a	e1b	e1a	e1b	e5a	e5b
10							
11	MetroCluster 3, MetroCluster interface	e1a	e1b	e1a	e1b	e5a	e5b
12							
13 / 13s0-3	MetroCluster ISL 40/100G or 4x25G or 4x10G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
14 / 14s0-3							
15	ISL, Local Cluster 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
16							

Note 1: Use either ports e4a and e4e or e4a and e8a if you are using an X91440A adapter (40Gbps). Use either ports e4a and e4b or e4a and e8a if you are using an X91153A adapter (100Gbps).

NVIDIA SN2100 platform port assignments (group 5)

Review the platform port assignments to cable an AFF A50 system to a NVIDIA SN2100 switch:

Switch Port	Port use	AFF A50	
		IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4b
2			
3	MetroCluster 2, Local Cluster interface	e4a	e4b
4			
5	MetroCluster 3, Local Cluster interface	e4a	e4b
6			
7	MetroCluster 1, MetroCluster interface	e2a	e2b
8			
9	MetroCluster 2, MetroCluster interface	e2a	e2b
10			
11	MetroCluster 3, MetroCluster interface	e2a	e2b
12			
13 / 13s0-3	MetroCluster ISL 40/100G or 4x25G or 4x10G	ISL, MetroCluster	
14 / 14s0-3			
15	ISL, Local Cluster 100G	ISL, Local Cluster	
16			

NVIDIA SN2100 platform port assignments (group 6)

Review the platform port assignments to cable an AFF C800, ASA C800, AFF A800, ASA A800, FAS9500, AFF A900, or ASA A900 system to a NVIDIA SN2100 switch:

Switch Port	Port use	AFF C80 ASA C800 AFF A800 ASA A800		FAS9500 AFF A900 ASA A900	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e0a	e1a	e4a	e4b(e) / e8a Note 1
2					
3	MetroCluster 2, Local Cluster interface	e0a	e1a	e4a	e4b(e) / e8a Note 1
4					
5	MetroCluster 3, Local Cluster interface	e0a	e1a	e4a	e4b(e) / e8a Note 1
6					
7	MetroCluster 1, MetroCluster interface	e0b	e1b	e5b	e7b
8					
9	MetroCluster 2, MetroCluster interface	e0b	e1b	e5b	e7b
10					
11	MetroCluster 3, MetroCluster interface	e0b	e1b	e5b	e7b
12					
13 / 13s0-3	MetroCluster ISL 40/100G or 4x25G or 4x10G	ISL, MetroCluster		ISL, MetroCluster	
14 / 14s0-3					
15	ISL, Local Cluster 100G	ISL, Local Cluster		ISL, Local Cluster	
16					

Note 1: Use either ports e4a and e4e or e4a and e8a if you are using an X91440A adapter (40Gbps). Use either ports e4a and e4b or e4a and e8a if you are using an X91153A adapter (100Gbps).

NVIDIA SN2100 platform port assignments (group 7)

Review the platform port assignments to cable a FAS70, AFF A70, AFF C80, FAS90, AFF A90, or AFF A1K system to a NVIDIA SN2100 switch:

Switch Port	Port use	FAS70 AFF A70		AFF C80		FAS90 AFF A90		AFF A1K	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e1a	e7a	e1a	e7a	e1a	e7a	e1a	e7a
2									
3	MetroCluster 2, Local Cluster interface	e1a	e7a	e1a	e7a	e1a	e7a	e1a	e7a
4									
5	MetroCluster 3, Local Cluster interface	e1a	e7a	e1a	e7a	e1a	e7a	e1a	e7a
6									
7	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e3b	e2b	e3b	e2b	e3b
8									
9	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e3b	e2b	e3b	e2b	e3b
10									
11	MetroCluster 3, MetroCluster interface	e2a	e2b	e2a	e3b	e2b	e3b	e2b	e3b
12									
13 / 13s0-3	MetroCluster ISL 40/100G or 4x25G or 4x10G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
14 / 14s0-3									
15	ISL, Local Cluster 100G	ISL, Local Cluster							
16									

Cable the ONTAP controller module ports in a MetroCluster IP configuration

You must cable the controller module ports used for cluster peering, management and data connectivity.

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.

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.

[Cluster and SVM peering express configuration](#)

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

Use the installation instructions for your platform at the [ONTAP Hardware Systems Documentation](#).



MetroCluster IP systems do not have dedicated high-availability (HA) ports. Depending on your platform, HA traffic is served using the MetroCluster, local cluster, or shared cluster/MetroCluster interface. When using *ONTAP Hardware Systems Documentation* to install your platform, you should not follow the instructions to cable the cluster and HA ports.

Configure the MetroCluster IP switches

Choose the correct MetroCluster IP switch configuration procedure

You must configure the IP switches to provide backend MetroCluster IP connectivity. The procedure you follow depends on your switch vendor.

- [Configure Broadcom IP switches](#)
- [Configure Cisco IP switches](#)
- [Configure NVIDIA IP switches](#)

Configure Broadcom IP switches for cluster interconnect and backend MetroCluster IP connectivity

You must configure the Broadcom IP switches for use as the cluster interconnect and for backend MetroCluster IP connectivity.



Your configuration requires additional licenses (6 x 100-Gb port license) in the following scenarios:

- You use ports 53 and 54 as a 40-Gbps or 100-Gbps MetroCluster ISL.
- You use a platform that connects the local cluster and MetroCluster interfaces to ports 49 - 52.

Resetting the Broadcom IP switch to factory defaults

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

About this task

- You must repeat these steps on each of the IP switches in the MetroCluster IP configuration.
- You must be connected to the switch using the serial console.
- This task resets the configuration of the management network.

Steps

1. Change to the elevated command prompt (#): enable

```
(IP_switch_A_1)> enable
(IP_switch_A_1) #
```

2. Erase the startup configuration and remove the banner
 - a. Erase the startup configuration:

erase startup-config

```
(IP_switch_A_1) #erase startup-config

Are you sure you want to clear the configuration? (y/n) y

(IP_switch_A_1) #
```

This command does not erase the banner.

- b. Remove the banner:

no set clibanner

```
(IP_switch_A_1) #configure
(IP_switch_A_1) (Config) # no set clibanner
(IP_switch_A_1) (Config) #
```

3. Reboot the switch: **(IP_switch_A_1) #reload**

```
Are you sure you would like to reset the system? (y/n) y
```



If the system asks whether to save the unsaved or changed configuration before reloading the switch, select **No**.

4. Wait for the switch to reload, and then log in to the switch.

The default user is “admin”, and no password is set. A prompt similar to the following is displayed:

```
(Routing)>
```

5. Change to the elevated command prompt:

```
enable
```

```
Routing)> enable
(Routing) #
```

6. Set the service port protocol to none:

```
serviceport protocol none
```

```
(Routing) #serviceport protocol none
Changing protocol mode will reset ip configuration.
Are you sure you want to continue? (y/n) y

(Routing) #
```

7. Assign the IP address to the service port:

```
serviceport ip ip-address netmask gateway
```

The following example shows a service port assigned IP address "10.10.10.10" with subnet "255.255.255.0" and gateway "10.10.10.1":

```
(Routing) #serviceport ip 10.10.10.10 255.255.255.0 10.10.10.1
```

8. Verify that the service port is correctly configured:

```
show serviceport
```

The following example shows that the port is up and the correct addresses have been assigned:

```

(Routing) #show serviceport

Interface Status..... Up
IP Address..... 10.10.10.10
Subnet Mask..... 255.255.255.0
Default Gateway..... 10.10.10.1
IPv6 Administrative Mode..... Enabled
IPv6 Prefix is .....
fe80::dac4:97ff:fe56:87d7/64
IPv6 Default Router..... fe80::222:bdff:fef8:19ff
Configured IPv4 Protocol..... None
Configured IPv6 Protocol..... None
IPv6 AutoConfig Mode..... Disabled
Burned In MAC Address..... D8:C4:97:56:87:D7

(Routing) #

```

9. Configure the SSH server.



- The RCF file disables the Telnet protocol. If you do not configure the SSH server, you can only access the bridge using the serial port connection.
- You must configure the SSH server in order to use log collection and other external tools.

a. Generate RSA keys.

```

(Routing) #configure
(Routing) (Config)#crypto key generate rsa

```

b. Generate DSA keys (optional)

```

(Routing) #configure
(Routing) (Config)#crypto key generate dsa

```

c. If you are using the FIPS compliant version of EFOS, generate the ECDSA keys. The following example creates the keys with a length of 521. Valid values are 256, 384 or 521.

```

(Routing) #configure
(Routing) (Config)#crypto key generate ecdsa 521

```

d. Enable the SSH server.

If necessary, exit the configuration context.

```
(Routing) (Config)#end
(Routing) #ip ssh server enable
```



If keys already exist, then you might be asked to overwrite them.

10. If desired, configure the domain and name server:

```
configure
```

The following example shows the `ip domain` and `ip name server` commands:

```
(Routing) # configure
(Routing) (Config)#ip domain name lab.netapp.com
(Routing) (Config)#ip name server 10.99.99.1 10.99.99.2
(Routing) (Config)#exit
(Routing) (Config)#
```

11. If desired, configure the time zone and time synchronization (SNTP).

The following example shows the `sntp` commands, specifying the IP address of the SNTP server and the relative time zone.

```
(Routing) #
(Routing) (Config)#sntp client mode unicast
(Routing) (Config)#sntp server 10.99.99.5
(Routing) (Config)#clock timezone -7
(Routing) (Config)#exit
(Routing) (Config)#
```

For EFOS version 3.10.0.3 and later, use the `ntp` command, as shown in the following example:

```

> (Config)# ntp ?

authenticate          Enables NTP authentication.
authentication-key    Configure NTP authentication key.
broadcast             Enables NTP broadcast mode.
broadcastdelay        Configure NTP broadcast delay in microseconds.
server               Configure NTP server.
source-interface      Configure the NTP source-interface.
trusted-key           Configure NTP authentication key number for
trusted time source.
vrf                  Configure the NTP VRF.

>(Config)# ntp server ?

ip-address|ipv6-address|hostname  Enter a valid IPv4/IPv6 address or
hostname.

>(Config)# ntp server 10.99.99.5

```

12. Configure the switch name:

```
hostname IP_switch_A_1
```

The switch prompt will display the new name:

```

(Routing) # hostname IP_switch_A_1

(IP_switch_A_1) #

```

13. Save the configuration:

```
write memory
```

You receive prompts and output similar to the following example:

```
(IP_switch_A_1) #write memory
```

This operation may take a few minutes.
Management interfaces will not be available during this time.

```
Are you sure you want to save? (y/n) y
```

```
Config file 'startup-config' created successfully .
```

```
Configuration Saved!
```

```
(IP_switch_A_1) #
```

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

Downloading and installing the Broadcom switch EFOS software

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

About this task

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

Note the following:

- When upgrading from EFOS 3.4.x.x to EFOS 3.7.x.x or later, the switch must be running EFOS 3.4.4.6 (or later 3.4.x.x release). If you are running a release prior to that, then upgrade the switch to EFOS 3.4.4.6 (or later 3.4.x.x release) first, then upgrade the switch to EFOS 3.7.x.x or later.
- The configuration for EFOS 3.4.x.x and 3.7.x.x or later are different. Changing the EFOS version from 3.4.x.x to 3.7.x.x or later, or vice versa, requires the switch to be reset to factory defaults and the RCF files for the corresponding EFOS version to be (re)applied. This procedure requires access through the serial console port.
- Beginning with EFOS version 3.7.x.x or later, a non-FIPS compliant and a FIPS compliant version is available. Different steps apply when moving to from a non-FIPS compliant to a FIPS compliant version or vice versa. Changing EFOS from a non-FIPS compliant to a FIPS compliant version or vice versa will reset the switch to factory defaults. This procedure requires access through the serial console port.

Steps

1. Download the switch firmware from the [Broadcom support site](#).
2. Check if your version of EFOS is FIPS compliant or non-FIPS compliant by using the `show fips status` command. In the following examples, `IP_switch_A_1` is using FIPS compliant EFOS and `IP_switch_A_2` is using non-FIPS compliant EFOS.

Example 1

```
IP_switch_A_1 #show fips status

System running in FIPS mode

IP_switch_A_1 #
```

Example 2

```
IP_switch_A_2 #show fips status
                ^
% Invalid input detected at ``^` marker.

IP_switch_A_2 #
```

3. Use the following table to determine which method you must follow:

Procedure	Current EFOS version	New EFOS version	High level steps
Steps to upgrade EFOS between two (non) FIPS compliant versions	3.4.x.x	3.4.x.x	Install the new EFOS image using method 1) The configuration and license information is retained
	3.4.4.6 (or later 3.4.x.x)	3.7.x.x or later non-FIPS compliant	Upgrade EFOS using method 1. Reset the switch to factory defaults and apply the RCF file for EFOS 3.7.x.x or later
	3.7.x.x or later non-FIPS compliant	3.4.4.6 (or later 3.4.x.x)	Downgrade EFOS using method 1. Reset the switch to factory defaults and apply the RCF file for EFOS 3.4.x.x
		3.7.x.x or later non-FIPS compliant	Install the new EFOS image using method 1. The configuration and license information is retained
3.7.x.x or later FIPS compliant	3.7.x.x or later FIPS compliant	Install the new EFOS image using method 1. The configuration and license information is retained	

Steps to upgrade to/from a FIPS compliant EFOS version	Non-FIPS compliant	FIPS compliant	Installation of the EFOS image using method 2. The switch configuration and license information will be lost.
	FIPS compliant	Non-FIPS compliant	

- Method 1: [Steps to upgrade EFOS with downloading the software image to the backup boot partition](#)
- Method 2: [Steps to upgrade EFOS using the ONIE OS installation](#)

Steps to upgrade EFOS with downloading the software image to the backup boot partition

You can perform the following steps only if both EFOS versions are non-FIPS compliant or both EFOS versions are FIPS compliant.



Do not use these steps if one version is FIPS compliant and the other version is non-FIPS compliant.

Steps

1. Copy the switch software to the switch: `copy sftp://user@50.50.50.50/switchsoftware/efos-3.4.4.6.stk backup`

In this example, the efos-3.4.4.6.stk operating system file is copied from the SFTP server at 50.50.50.50 to the backup partition. You need to 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://user@50.50.50.50/switchsoftware/efos-3.4.4.6.stk backup
Remote Password:*****
```

```
Mode..... SFTP
Set Server IP..... 50.50.50.50
Path..... /switchsoftware/
Filename..... efos-3.4.4.6.stk
Data Type..... Code
Destination Filename..... backup
```

```
Management access will be blocked for the duration of the transfer
Are you sure you want to start? (y/n) y
```

```
File transfer in progress. Management access will be blocked for the
duration of the transfer. Please wait...
SFTP Code transfer starting...
```

```
File transfer operation completed successfully.
```

```
(IP_switch_A_1) #
```

2. Set the switch to boot from the backup partition on the next switch reboot:

```
boot system backup
```

```
(IP_switch_A_1) #boot system backup
Activating image backup ..
```

```
(IP_switch_A_1) #
```

3. Verify that the new boot image will be active on the next boot:

```
show bootvar
```

```
(IP_switch_A_1) #show bootvar
```

```
Image Descriptions
```

```
active :
```

```
backup :
```

```
Images currently available on Flash
```

unit	active	backup	current-active	next-active
1	3.4.4.2	3.4.4.6	3.4.4.2	3.4.4.6

```
(IP_switch_A_1) #
```

4. Save the configuration:

```
write memory
```

```
(IP_switch_A_1) #write memory
```

```
This operation may take a few minutes.
```

```
Management interfaces will not be available during this time.
```

```
Are you sure you want to save? (y/n) y
```

```
Configuration Saved!
```

```
(IP_switch_A_1) #
```

5. Reboot the switch:

```
reload
```

```
(IP_switch_A_1) #reload
```

```
Are you sure you would like to reset the system? (y/n) y
```

6. Wait for the switch to reboot.



In rare scenarios the switch may fail to boot. Follow the [Steps to upgrade EFOS using the ONIE OS installation](#) to install the new image.

7. If you change the switch from EFOS 3.4.x.x to EFOS 3.7.x.x or vice versa then follow the following two procedures to apply the correct configuration (RCF):
 - a. [Resetting the Broadcom IP switch to factory defaults](#)
 - b. [Downloading and installing the Broadcom RCF files](#)
8. Repeat these steps on the remaining three IP switches in the MetroCluster IP configuration.

Steps to upgrade EFOS using the ONIE OS installation

You can perform the following steps if one EFOS version is FIPS compliant and the other EFOS version is non-FIPS compliant. These steps can be used to install the non-FIPS or FIPS compliant EFOS 3.7.x.x image from ONIE if the switch fails to boot.

Steps

1. Boot the switch into ONIE installation mode.

During boot, select ONIE when the following screen appears:

```
+-----+
| EFOS  |
| *ONIE |
|       |
|       |
|       |
|       |
|       |
|       |
|       |
|       |
|       |
|       |
+-----+
```

After selecting "ONIE", the switch will then load and present you with the following choices:

```

+-----+
|*ONIE: Install OS
| ONIE: Rescue
| ONIE: Uninstall OS
| ONIE: Update ONIE
| ONIE: Embed ONIE
| DIAG: Diagnostic Mode
| DIAG: Burn-In Mode
|
|
|
|
|
+-----+

```

The switch now will boot into ONIE installation mode.

2. Stop the ONIE discovery and configure the ethernet interface

Once the following message appears press <enter> to invoke the ONIE console:

```

Please press Enter to activate this console. Info: eth0: Checking
link... up.
ONIE:/ #

```



The ONIE discovery will continue and messages will be printed to the console.

```

Stop the ONIE discovery
ONIE:/ # onie-discovery-stop
discover: installer mode detected.
Stopping: discover... done.
ONIE:/ #

```

3. Configure the ethernet interface and add the route using `ifconfig eth0 <ipAddress> netmask <netmask> up` and `route add default gw <gatewayAddress>`

```

ONIE:/ # ifconfig eth0 10.10.10.10 netmask 255.255.255.0 up
ONIE:/ # route add default gw 10.10.10.1

```

4. Verify that the server hosting the ONIE installation file is reachable:

```

ONIE:/ # ping 50.50.50.50
PING 50.50.50.50 (50.50.50.50): 56 data bytes
64 bytes from 50.50.50.50: seq=0 ttl=255 time=0.429 ms
64 bytes from 50.50.50.50: seq=1 ttl=255 time=0.595 ms
64 bytes from 50.50.50.50: seq=2 ttl=255 time=0.369 ms
^C
--- 50.50.50.50 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 0.369/0.464/0.595 ms
ONIE:/ #

```

5. Install the new switch software

```

ONIE:/ # onie-nos-install http:// 50.50.50.50/Software/onie-installer-
x86_64
discover: installer mode detected.
Stopping: discover... done.
Info: Fetching http:// 50.50.50.50/Software/onie-installer-3.7.0.4 ...
Connecting to 50.50.50.50 (50.50.50.50:80)
installer          100% |*****| 48841k
0:00:00 ETA
ONIE: Executing installer: http:// 50.50.50.50/Software/onie-installer-
3.7.0.4
Verifying image checksum ... OK.
Preparing image archive ... OK.

```

The software will install and then reboot the switch. Let the switch reboot normally into the new EFOS version.

6. Verify that the new switch software is installed

show bootvar

```

(Routing) #show bootvar
Image Descriptions
active :
backup :
Images currently available on Flash
-----
unit      active      backup      current-active  next-active
-----
1    3.7.0.4      3.7.0.4      3.7.0.4          3.7.0.4
(Routing) #

```

7. Complete the installation

The switch will reboot with no configuration applied and reset to factory defaults. Follow the two procedures to configure the switch basic settings and apply the RCF file as outlined in the following two documents:

- a. Configure the switch basic settings. Follow step 4 and later: [Resetting the Broadcom IP switch to factory defaults](#)
- b. Create and apply the RCF file as outlined in [Downloading and installing the Broadcom RCF files](#)

Downloading and installing the Broadcom RCF files

You must generate and install the switch RCF file to each switch in the MetroCluster IP 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 IP switches in the MetroCluster IP configuration.

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

Switch	RCF file
IP_switch_A_1	v1.32_Switch-A1.txt
IP_switch_A_2	v1.32_Switch-A2.txt
IP_switch_B_1	v1.32_Switch-B1.txt
IP_switch_B_2	v1.32_Switch-B2.txt



The RCF files for EFOS version 3.4.4.6 or later 3.4.x.x. release and EFOS version 3.7.0.4 are different. You need to make sure that you have created the correct RCF files for the EFOS version that the switch is running.

EFOS version	RCF file version
3.4.x.x	v1.3x, v1.4x
3.7.x.x	v2.x

Steps

1. Generate the Broadcom RCF files for MetroCluster IP.
 - a. Download the [RcfFileGenerator for MetroCluster IP](#)
 - b. Generate the RCF file for your configuration using the RcfFileGenerator for MetroCluster IP.



Modifications to the RCF files after download are not supported.

2. Copy the RCF files to the switches:

- a. Copy the RCF files to the first switch: `copy sftp://user@FTP-server-IP-address/RcfFiles/switch-specific-RCF/BES-53248_v1.32_Switch-A1.txt nvram:script BES-53248_v1.32_Switch-A1.scr`

In this example, the "BES-53248_v1.32_Switch-A1.txt" RCF file is copied from the SFTP server at "50.50.50.50" to the local bootflash. You need to 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://user@50.50.50.50/RcfFiles/BES-53248_v1.32_Switch-A1.txt nvram:script BES-53248_v1.32_Switch-A1.scr
```

```
Remote Password:*****
```

```
Mode..... SFTP
Set Server IP..... 50.50.50.50
Path..... /RcfFiles/
Filename..... BES-53248_v1.32_Switch-A1.txt
Data Type..... Config Script
Destination Filename..... BES-53248_v1.32_Switch-A1.scr
```

```
Management access will be blocked for the duration of the transfer
Are you sure you want to start? (y/n) y
```

```
File transfer in progress. Management access will be blocked for the
duration of the transfer. Please wait...
File transfer operation completed successfully.
```

```
Validating configuration script...
```

```
config
```

```
set clibanner
```

```
*****
*****
```

```
* NetApp Reference Configuration File (RCF)
```

```
*
```

```
* Switch : BES-53248
```

```
...
```

```
The downloaded RCF is validated. Some output is being logged here.
```

```
...
```

```
Configuration script validated.
```

```
File transfer operation completed successfully.
```

```
(IP_switch_A_1) #
```

b. Verify that the RCF file is saved as a script:

```
script list
```

```
(IP_switch_A_1) #script list

Configuration Script Name          Size(Bytes)  Date of Modification
-----
BES-53248_v1.32_Switch-A1.scr      852         2019 01 29 18:41:25

1 configuration script(s) found.
2046 Kbytes free.
(IP_switch_A_1) #
```

c. Apply the RCF script:

```
script apply BES-53248_v1.32_Switch-A1.scr
```

```
(IP_switch_A_1) #script apply BES-53248_v1.32_Switch-A1.scr

Are you sure you want to apply the configuration script? (y/n) y

config

set clibanner
"*****
*****

* NetApp Reference Configuration File (RCF)

*

* Switch      : BES-53248

...
The downloaded RCF is validated. Some output is being logged here.
...

Configuration script 'BES-53248_v1.32_Switch-A1.scr' applied.

(IP_switch_A_1) #
```

d. Save the configuration:

```
write memory
```

```
(IP_switch_A_1) #write memory
```

This operation may take a few minutes.
Management interfaces will not be available during this time.

```
Are you sure you want to save? (y/n) y
```

```
Configuration Saved!
```

```
(IP_switch_A_1) #
```

e. Reboot the switch:

```
reload
```

```
(IP_switch_A_1) #reload
```

```
Are you sure you would like to reset the system? (y/n) y
```

f. Repeat the previous steps for each of the other three switches, being sure to copy the matching RCF file to the corresponding switch.

3. Reload the switch:

```
reload
```

```
IP_switch_A_1# reload
```

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

Disable unused ISL ports and port channels

NetApp recommends disabling unused ISL ports and port channels to avoid unnecessary health alerts.

1. Identify the unused ISL ports and port channels using the RCF file banner:



If the port is in breakout mode, the port name you specify in the command might be different than the name stated in the RCF banner. You can also use the RCF cabling files to find the port name.

For ISL port details

Run the command `show port all`.

For port channel details

Run the command `show port-channel all`.

2. Disable the unused ISL ports and port channels.

You must run the following commands for each identified unused port or port channel.

```
(SwtichA_1)> enable
(SwtichA_1)# configure
(SwtichA_1) (Config)# <port_name>
(SwtichA_1) (Interface 0/15)# shutdown
(SwtichA_1) (Interface 0/15)# end
(SwtichA_1)# write memory
```

Configure Cisco IP switches

Configure Cisco IP switches for cluster interconnect and backend MetroCluster IP connectivity

You must configure the Cisco IP switches for use as the cluster interconnect and for backend MetroCluster IP connectivity.

About this task

Several of the procedures in this section are independent procedures and you only need to execute those you are directed to or are relevant to your task.

Resetting the Cisco IP switch to factory defaults

Before installing any RCF file, you must erase the Cisco switch configuration and perform basic configuration. This procedure is required when you want to reinstall the same RCF file after a previous installation failed, or if you want to install a new version of an RCF file.

About this task

- You must repeat these steps on each of the IP switches in the MetroCluster IP configuration.
- You must be connected to the switch using the serial console.
- This task resets the configuration of the management network.

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 (RSA)

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 configuring 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.

You enter basic information in the next set of prompts, including the switch name, management address, and gateway, and select SSH with RSA.



This example shows the minimum information required to configure the RCF, additional options can be configured after the RCF has been applied. For example, you can configure SNMPv3, NTP, or SCP/SFTP after you have applied the RCF.

The following configuration will be applied:

```
password strength-check
switchname IP_switch_A_1
vrf context management
ip route 0.0.0.0/0 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.
```

```
[#####] 100%
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 IP configuration.

Downloading and installing the Cisco switch NX-OS software

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

About this task

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

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

You must use the supported switch software version.

[NetApp Hardware Universe](#)

Steps

1. Download the supported NX-OS software file.

[Cisco Software Download](#)

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.I4.6.bin file and EPLD image 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)...
Copy complete.

IP_switch_A_1# copy sftp://root@10.10.99.99/tftpboot/n9000-
epld.9.3.5.img bootflash: vrf management
root@10.10.99.99's password: password
sftp> progress
Progress meter enabled
sftp> get /tftpboot/n9000-epld.9.3.5.img /bootflash/n9000-
epld.9.3.5.img
Fetching /tftpboot/n9000-epld.9.3.5.img to /bootflash/n9000-
epld.9.3.5.img
/tftpboot/n9000-epld.9.3.5.img          161MB 9.5MB/s 00:16
sftp> exit
Copy complete, now saving to disk (please wait)...
Copy complete.

```

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:

```

IP_switch_A_1# dir bootflash:
      .
      .
      .
698629632   Jun 13 21:37:44 2017  nxos.7.0.3.I4.6.bin
      .
      .
      .

Usage for bootflash://sup-local
 1779363840 bytes used
13238841344 bytes free
15018205184 bytes total
IP_switch_A_1#

```

4. Install the switch software:

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

The switch will reload (reboot) automatically after the switch software has been installed.

The following example shows the software installation on IP_switch_A_1:

```

IP_switch_A_1# install all nxos bootflash:nxos.7.0.3.I4.6.bin
Installer will perform compatibility check first. Please wait.
Installer is forced disruptive

Verifying image bootflash:/nxos.7.0.3.I4.6.bin for boot variable "nxos".
[#####] 100% -- SUCCESS

Verifying image type.
[#####] 100% -- SUCCESS

Preparing "nxos" version info using image
bootflash:/nxos.7.0.3.I4.6.bin.
[#####] 100% -- SUCCESS

Preparing "bios" version info using image
bootflash:/nxos.7.0.3.I4.6.bin.
[#####] 100% -- SUCCESS          [#####] 100%
-- SUCCESS

Performing module support checks.          [#####] 100%
-- SUCCESS

Notifying services about system upgrade.   [#####] 100%

```

```
-- SUCCESS
```

```
Compatibility check is done:
```

Module	bootable	Impact	Install-type	Reason
1	yes	disruptive	reset	default upgrade is not hitless

```
Images will be upgraded according to following table:
```

Module Required	Image	Running-Version(pri:alt)	New-Version	Upg-
1	nxos	7.0(3)I4(1)	7.0(3)I4(6)	yes
1	bios	v04.24(04/21/2016)	v04.24(04/21/2016)	no

```
Switch will be reloaded for disruptive upgrade.
```

```
Do you want to continue with the installation (y/n)? [n] y
```

```
Install is in progress, please wait.
```

```
Performing runtime checks. [#####] 100% --  
SUCCESS
```

```
Setting boot variables.  
[#####] 100% -- SUCCESS
```

```
Performing configuration copy.  
[#####] 100% -- SUCCESS
```

```
Module 1: Refreshing compact flash and upgrading bios/loader/bootrom.  
Warning: please do not remove or power off the module at this time.  
[#####] 100% -- SUCCESS
```

```
Finishing the upgrade, switch will reboot in 10 seconds.  
IP_switch_A_1#
```

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)I4(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. Upgrade the EPLD image and reboot the switch.

```

IP_switch_A_1# install epld bootflash:n9000-epld.9.3.5.img module 1
Compatibility check:
Module          Type          Upgradable    Impact        Reason
-----
1              SUP              Yes           disruptive    Module Upgradable

Retrieving EPLD versions.... Please wait.
Images will be upgraded according to following table:
Module  Type  EPLD          Running-Version  New-Version  Upg-
Required
-----
1  SUP  MI FPGA          0x07            0x07        No
1  SUP  IO FPGA          0x17            0x19        Yes
1  SUP  MI FPGA2        0x02            0x02        No

The above modules require upgrade.
The switch will be reloaded at the end of the upgrade
Do you want to continue (y/n) ? [n] y

Proceeding to upgrade Modules.

Starting Module 1 EPLD Upgrade

Module 1 : IO FPGA [Programming] : 100.00% (      64 of      64 sectors)
Module 1 EPLD upgrade is successful.
Module  Type  Upgrade-Result
-----
1  SUP  Success

EPLDs upgraded.

Module 1 EPLD upgrade is successful.

```

- After the switch reboot, log in again and verify that the new version of EPLD loaded successfully.

```
show version module 1 epld
```

- Repeat these steps on the remaining three IP switches in the MetroCluster IP configuration.

Downloading and installing the Cisco IP RCF files

You must generate and install the RCF file to each switch in the MetroCluster IP configuration.

About this task

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

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

You must use the supported switch software version.

NetApp Hardware Universe

If you are using a QSFP-to-SFP+ adapter, you might need to configure the ISL port in native speed mode instead of breakout speed mode. Refer to your switch vendor documentation to determine the ISL port speed mode.

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

Switch	RCF file
IP_switch_A_1	NX3232_v1.80_Switch-A1.txt
IP_switch_A_2	NX3232_v1.80_Switch-A2.txt
IP_switch_B_1	NX3232_v1.80_Switch-B1.txt
IP_switch_B_2	NX3232_v1.80_Switch-B2.txt

Steps

1. Generate the Cisco RCF files for MetroCluster IP.
 - a. Download the [RcfFileGenerator for MetroCluster IP](#)
 - b. Generate the RCF file for your configuration using the RcfFileGenerator for MetroCluster IP.



Modifications to the RCF files after download are not supported.

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_v1.80_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:

```

IP_switch_A_1# dir bootflash:
.
.
.
5514   Jun 13 22:09:05 2017  NX3232_v1.80_Switch-A1.txt
.
.
.

Usage for bootflash://sup-local
1779363840 bytes used
13238841344 bytes free
15018205184 bytes total
IP_switch_A_1#

```

4. Configure the TCAM regions on Cisco 3132Q-V and Cisco 3232C switches.



Skip this step if you do not have Cisco 3132Q-V or Cisco 3232C switches.

a. On Cisco 3132Q-V switch, set the following TCAM regions:

```
conf t
hardware access-list tcam region span 0
hardware access-list tcam region racl 256
hardware access-list tcam region e-racl 256
hardware access-list tcam region qos 256
```

- b. On Cisco 3232C switch, set the following TCAM regions:

```
conf t
hardware access-list tcam region span 0
hardware access-list tcam region racl-lite 0
hardware access-list tcam region racl 256
hardware access-list tcam region e-racl 256
hardware access-list tcam region qos 256
```

- c. After setting the TCAM regions, save the configuration and reload the switch:

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

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

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

6. 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
```

7. Reload the switch:

```
reload
```

```
IP_switch_A_1# reload
```

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

Setting Forward Error Correction for systems using 25-Gbps connectivity

If your system is configured using 25-Gbps connectivity, you need to set the Forward Error Correction (fec) parameter manually to off after applying the RCF file. The RCF file does not apply this setting.

About this task

The 25-Gbps ports must be cabled prior to performing this procedure.

[Platform port assignments for Cisco 3232C or Cisco 9336C switches](#)

This task only applies to platforms using 25-Gbps connectivity:

- AFF A300
- FAS 8200
- FAS 500f
- AFF A250

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

Steps

1. Set the fec parameter to off on each 25-Gbps port that is connected to a controller module, and then copy the running configuration to the startup configuration:
 - a. Enter configuration mode: `conf t`
 - b. Specify the 25-Gbps interface to configure: `interface interface-ID`
 - c. Set fec to off: `fec off`
 - d. Repeat the previous steps for each 25-Gbps port on the switch.
 - e. Exit configuration mode: `exit`

The following example shows the commands for interface Ethernet1/25/1 on switch IP_switch_A_1:

```
IP_switch_A_1# conf t
IP_switch_A_1(config)# interface Ethernet1/25/1
IP_switch_A_1(config-if)# fec off
IP_switch_A_1(config-if)# exit
IP_switch_A_1(config-if)# end
IP_switch_A_1# copy running-config startup-config
```

2. Repeat the previous step on the other three switches in the MetroCluster IP configuration.

Disable unused ISL ports and port channels

NetApp recommends disabling unused ISL ports and port channels to avoid unnecessary health alerts.

1. Identify the unused ISL ports and port channels:

```
show interface brief
```

2. Disable the unused ISL ports and port channels.

You must run the following commands for each identified unused port or port channel.

```
SwitchA_1# config t
Enter configuration commands, one per line. End with CNTL/Z.
SwitchA_1(config)# int Eth1/14
SwitchA_1(config-if)# shutdown
SwitchA_12(config-if)# exit
SwitchA_1(config-if)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete.
```

Configure MACsec encryption on Cisco 9336C switches in a MetroCluster IP site

You can configure MACsec encryption on the WAN ISL ports that run between the sites. You must configure MACsec after applying the correct RCF file.



MACsec encryption can only be applied to the WAN ISL ports.

Configure MACsec encryption on Cisco 9336C switches

You must only configure MACsec encryption on the WAN ISL ports that run between the sites. You must configure MACsec after applying the correct RCF file.

Licensing requirements for MACsec

MACsec requires a security license. For a complete explanation of the Cisco NX-OS licensing scheme and how to obtain and apply for licenses, see the [Cisco NX-OS Licensing Guide](#)

Enable Cisco MACsec Encryption WAN ISLs in MetroCluster IP configurations

You can enable MACsec encryption for Cisco 9336C switches on the WAN ISLs in a MetroCluster IP configuration.

Steps

1. Enter global configuration mode:

```
configure terminal
```

```
IP_switch_A_1# configure terminal
IP_switch_A_1(config)#
```

2. Enable MACsec and MKA on the device:

```
feature macsec
```

```
IP_switch_A_1(config)# feature macsec
```

3. Copy the running configuration to the startup configuration:

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

```
IP_switch_A_1(config)# copy running-config startup-config
```

Configure a MACsec key chain and keys

You can create a MACsec key chain or keys on your configuration.

Key Lifetime and Hitless Key Rollover

A MACsec keychain can have multiple pre-shared keys (PSKs), each configured with a key ID and an optional lifetime. A key lifetime specifies at which time the key activates and expires. In the absence of a lifetime configuration, the default lifetime is unlimited. When a lifetime is configured, MKA rolls over to the next configured pre-shared key in the keychain after the lifetime is expired. The time zone of the key can be local or UTC. The default time zone is UTC. A key can roll over to a second key within the same keychain if you configure the second key (in the keychain) and configure a lifetime for the first key. When the lifetime of the first key expires, it automatically rolls over to the next key in the list. If the same key is configured on both sides of the link at the same time, then the key rollover is hitless (that is, the key rolls over without traffic interruption).

Steps

1. Enter the global configuration mode:

```
configure terminal
```

```
IP_switch_A_1# configure terminal  
IP_switch_A_1(config)#
```

2. To hide the encrypted key octet string, replace the string with a wildcard character in the output of the `show running-config` and `show startup-config` commands:

```
IP_switch_A_1(config)# key-chain macsec-psk no-show
```



The octet string is also hidden when you save the configuration to a file.

By default, PSK keys are displayed in encrypted format and can easily be decrypted. This command applies only to MACsec key chains.

3. Create a MACsec key chain to hold a set of MACsec keys and enter MACsec key chain configuration mode:

```
key chain name macsec
```

```
IP_switch_A_1(config)# key chain 1 macsec
IP_switch_A_1(config-macseckeychain)#
```

4. Create a MACsec key and enter MACsec key configuration mode:

```
key key-id
```

The range is from 1 to 32 hex digit key-string, and the maximum size is 64 characters.

```
IP_switch_A_1 switch(config-macseckeychain)# key 1000
IP_switch_A_1 (config-macseckeychain-macseckey)#
```

5. Configure the octet string for the key:

```
key-octet-string octet-string cryptographic-algorithm AES_128_CMAC |
AES_256_CMAC
```

```
IP_switch_A_1(config-macseckeychain-macseckey)# key-octet-string
abcdef0123456789abcdef0123456789abcdef0123456789abcdef0123456789
cryptographic-algorithm AES_256_CMAC
```



The octet-string argument can contain up to 64 hexadecimal characters. The octet key is encoded internally, so the key in clear text does not appear in the output of the `show running-config macsec` command.

6. Configure a send lifetime for the key (in seconds):

```
send-lifetime start-time duration duration
```

```
IP_switch_A_1(config-macseckeychain-macseckey)# send-lifetime 00:00:00
Oct 04 2020 duration 100000
```

By default, the device treats the start time as UTC. The start-time argument is the time of day and date that the key becomes active. The duration argument is the length of the lifetime in seconds. The maximum length is 2147483646 seconds (approximately 68 years).

7. Copy the running configuration to the startup configuration:

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

```
IP_switch_A_1(config)# copy running-config startup-config
```

8. Displays the keychain configuration:

```
show key chain name
```

```
IP_switch_A_1(config-macseckeychain-macseckey)# show key chain 1
```

Configure a MACsec policy

Steps

1. Enter global configuration mode:

```
configure terminal
```

```
IP_switch_A_1# configure terminal  
IP_switch_A_1(config)#
```

2. Create a MACsec policy:

```
macsec policy name
```

```
IP_switch_A_1(config)# macsec policy abc  
IP_switch_A_1(config-macsec-policy)#
```

3. Configure one of the following ciphers, GCM-AES-128, GCM-AES-256, GCM-AES-XPN-128, or GCM-AES-XPN-256:

```
cipher-suite name
```

```
IP_switch_A_1(config-macsec-policy)# cipher-suite GCM-AES-256
```

4. Configure the key server priority to break the tie between peers during a key exchange:

```
key-server-priority number
```

```
switch(config-macsec-policy)# key-server-priority 0
```

5. Configure the security policy to define the handling of data and control packets:

```
security-policy security policy
```

Choose a security policy from the following options:

- must-secure — packets not carrying MACsec headers are dropped
- should-secure — packets not carrying MACsec headers are permitted (this is the default value)

```
IP_switch_A_1(config-macsec-policy)# security-policy should-secure
```

6. Configure the replay protection window so the secured interface does not accept a packet that is less than the configured window size: `window-size number`



The replay protection window size represents the maximum out-of-sequence frames that MACsec accepts and are not discarded. The range is from 0 to 596000000.

```
IP_switch_A_1(config-macsec-policy)# window-size 512
```

7. Configure the time in seconds to force an SAK rekey:

```
sak-expiry-time time
```

You can use this command to change the session key to a predictable time interval. The default is 0.

```
IP_switch_A_1(config-macsec-policy)# sak-expiry-time 100
```

8. Configure one of the following confidentiality offsets in the layer 2 frame where encryption begins:

```
conf-offsetconfidentiality offset
```

Choose from the following options:

- CONF-OFFSET-0.
- CONF-OFFSET-30.
- CONF-OFFSET-50.

```
IP_switch_A_1(config-macsec-policy)# conf-offset CONF-OFFSET-0
```



This command might be necessary for intermediate switches to use packet headers (dmac, smac, etype) like MPLS tags.

9. Copy the running configuration to the startup configuration:

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

```
IP_switch_A_1(config)# copy running-config startup-config
```

10. Display the MACsec policy configuration:

```
show macsec policy
```

```
IP_switch_A_1(config-macsec-policy)# show macsec policy
```

Enable Cisco MACsec encryption on the interfaces

1. Enter global configuration mode:

```
configure terminal
```

```
IP_switch_A_1# configure terminal
IP_switch_A_1(config)#
```

2. Select the interface that you configured with MACsec encryption.

You can specify the interface type and identity. For an Ethernet port, use ethernet slot/port.

```
IP_switch_A_1(config)# interface ethernet 1/15
switch(config-if)#
```

3. Add the keychain and policy to be configured on the interface to add the MACsec configuration:

```
macsec keychain keychain-name policy policy-name
```

```
IP_switch_A_1(config-if)# macsec keychain 1 policy abc
```

4. Repeat steps 1 and 2 on all interfaces where MACsec encryption is to be configured.
5. Copy the running configuration to the startup configuration:

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

```
IP_switch_A_1(config)# copy running-config startup-config
```

Disable Cisco MACsec Encryption WAN ISLs in MetroCluster IP configurations

You might need to disable MACsec encryption for Cisco 9336C switches on the WAN ISLs in a MetroCluster IP configuration.

Steps

1. Enter global configuration mode:

```
configure terminal
```

```
IP_switch_A_1# configure terminal
IP_switch_A_1(config)#
```

2. Disable the MACsec configuration on the device:

```
macsec shutdown
```

```
IP_switch_A_1(config)# macsec shutdown
```



Selecting the “no” option restores the MACsec feature.

3. Select the interface that you already configured with MACsec.

You can specify the interface type and identity. For an Ethernet port, use ethernet slot/port.

```
IP_switch_A_1(config)# interface ethernet 1/15
switch(config-if)#
```

4. Remove the keychain and policy configured on the interface to remove the MACsec configuration:

```
no macsec keychain keychain-name policy policy-name
```

```
IP_switch_A_1(config-if)# no macsec keychain 1 policy abc
```

5. Repeat steps 3 and 4 on all interfaces where MACsec is configured.
6. Copy the running configuration to the startup configuration:

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

```
IP_switch_A_1(config)# copy running-config startup-config
```

Verifying the MACsec configuration

Steps

1. Repeat **all** of the previous procedures on the second switch within the configuration to establish a MACsec session.
2. Run the following commands to verify that both switches are successfully encrypted:
 - a. Run: `show macsec mka summary`
 - b. Run: `show macsec mka session`
 - c. Run: `show macsec mka statistics`

You can verify the MACsec configuration using the following commands:

Command	Displays information about...
<code>show macsec mka session interface typeslot/port number</code>	The MACsec MKA session for a specific interface or for all interfaces
<code>show key chain name</code>	The key chain configuration
<code>show macsec mka summary</code>	The MACsec MKA configuration
<code>show macsec policy policy-name</code>	The configuration for a specific MACsec policy or for all MACsec policies

Configure NVIDIA IP switches

Configure NVIDIA IP SN2100 switch for cluster interconnect and backend MetroCluster IP connectivity

You must configure the NVIDIA SN2100 IP switches for use as the cluster interconnect and for backend MetroCluster IP connectivity.

Reset the NVIDIA IP SN2100 switch to factory defaults

You can choose from the following methods to reset a switch to factory default settings.

- [Reset the switch using the RCF file option](#)
- [Download and install the Cumulus software](#)

Reset the switch using the RCF file option

Before installing a new RCF configuration you must revert the NVIDIA switch settings.

About this task

To restore the switch to default settings, run the RCF file with the `restoreDefaults` option. This option copies the original backed up files to their original location and then reboots the switch. After reboot, the switch comes online with the original configuration that existed when you first ran the RCF file to configure the switch.

The following configuration details are not reset:

- User and credential configuration
- Configuration of the management network port, eth0



All other configuration changes that occur during application of the RCF file are reverted to the original configuration.

Before you begin

- You must configure the switch according to [Download and install the NVIDIA RCF file](#). If you have not configured in this manner, or you have configured additional features before running the RCF file, you cannot use this procedure.

- You must repeat these steps on each of the IP switches in the MetroCluster IP configuration.
- You must be connected to the switch with a serial console connection.
- This task resets the configuration of the management network.

Steps

1. Verify that the RCF configuration was successfully applied with the same or a compatible RCF file version and that the backup files exist.



The output can show backup files, preserved files, or both. If backup files or preserved files do not appear in the output, you cannot use this procedure.

```

cumulus@IP_switch_A_1:mgmt:~$ sudo python3
SN2100_v2.0.0_IP_switch_A_1.py
[sudo] password for cumulus:
>>> Opened RcfApplyLog
A RCF configuration has been successfully applied.
Backup files exist.
Preserved files exist.
Listing completion of the steps:
    Success: Step: 1: Performing Backup and Restore
    Success: Step: 2: updating MOTD file
    Success: Step: 3: Disabling apt-get
    Success: Step: 4: Disabling cdp
    Success: Step: 5: Adding lldp config
    Success: Step: 6: Creating interfaces
    Success: Step: 7: Configuring switch basic settings: Hostname,
SNMP
    Success: Step: 8: Configuring switch basic settings: bandwidth
allocation
    Success: Step: 9: Configuring switch basic settings: ecn
    Success: Step: 10: Configuring switch basic settings: cos and
dscp remark
    Success: Step: 11: Configuring switch basic settings: generic
egress cos mappings
    Success: Step: 12: Configuring switch basic settings: traffic
classification
    Success: Step: 13: Configuring LAG load balancing policies
    Success: Step: 14: Configuring the VLAN bridge
    Success: Step: 15: Configuring local cluster ISL ports
    Success: Step: 16: Configuring MetroCluster ISL ports
    Success: Step: 17: Configuring ports for MetroCluster-1, local
cluster and MetroCluster interfaces
    Success: Step: 18: Configuring ports for MetroCluster-2, local
cluster and MetroCluster interfaces
    Success: Step: 19: Configuring ports for MetroCluster-3, local
cluster and MetroCluster interfaces
    Success: Step: 20: Configuring L2FC for MetroCluster interfaces
    Success: Step: 21: Configuring the interface to UP
    Success: Step: 22: Final commit
    Success: Step: 23: Final reboot of the switch
Exiting ...
<<< Closing RcfApplyLog
cumulus@IP_switch_A_1:mgmt:~$

```

2. Run the RCF file with the option to restore defaults: `restoreDefaults`

```

cumulus@IP_switch_A_1:mgmt:~$ sudo python3
SN2100_v2.0.0_IP_switch_A_2.py restoreDefaults
[sudo] password for cumulus:
>>> Opened RcfApplyLog
Can restore from backup directory. Continuing.
This will reboot the switch !!!
Enter yes or no: yes

```

3. Respond 'yes' to the prompt. The switch reverts to the original configuration and reboots.
4. Wait for the switch to reboot.

The switch is reset and retains the initial configuration such as management network configuration and current credentials as they existed before applying the RCF file. After reboot, you can apply a new configuration by using the same or a different version of the RCF file.

Download and install the Cumulus software

About this task

Use these steps if you want to reset the switch completely by applying the Cumulus image.

Before you begin

- You must be connected to the switch with a serial console connection.
- The Cumulus switch software image is accessible through HTTP.



For more information about installing Cumulus Linux, see [Overview of installation and configuration for NVIDIA SN2100 switches](#)

- You must have the root password for `sudo` access to the commands.

Steps

1. From the Cumulus console download and queue the switch software installation with the command `onie-install -a -i` followed by the file path to the switch software:

In this example, the firmware file `cumulus-linux-4.4.3-mlx-amd64.bin` is copied from the HTTP server '50.50.50.50' to the local switch.

```

cumulus@IP_switch_A_1:mgmt:~$ sudo onie-install -a -i
http://50.50.50.50/switchsoftware/cumulus-linux-4.4.3-mlx-amd64.bin
Fetching installer: http://50.50.50.50/switchsoftware/cumulus-linux-
4.4.3-mlx-amd64.bin
Downloading URL: http://50.50.50.50/switchsoftware/cumulus-linux-4.4.3-
mlx-amd64.bin
#####
# 100.0%
Success: HTTP download complete.
tar: ./sysroot.tar: time stamp 2021-01-30 17:00:58 is 53895092.604407122

```

```
s in the future
tar: ./kernel: time stamp 2021-01-30 17:00:58 is 53895092.582826352 s in
the future
tar: ./initrd: time stamp 2021-01-30 17:00:58 is 53895092.509682557 s in
the future
tar: ./embedded-installer/bootloader/grub: time stamp 2020-12-10
15:25:16 is 49482950.509433937 s in the future
tar: ./embedded-installer/bootloader/init: time stamp 2020-12-10
15:25:16 is 49482950.509336507 s in the future
tar: ./embedded-installer/bootloader/uboot: time stamp 2020-12-10
15:25:16 is 49482950.509213637 s in the future
tar: ./embedded-installer/bootloader: time stamp 2020-12-10 15:25:16 is
49482950.509153787 s in the future
tar: ./embedded-installer/lib/init: time stamp 2020-12-10 15:25:16 is
49482950.509064547 s in the future
tar: ./embedded-installer/lib/logging: time stamp 2020-12-10 15:25:16 is
49482950.508997777 s in the future
tar: ./embedded-installer/lib/platform: time stamp 2020-12-10 15:25:16
is 49482950.508913317 s in the future
tar: ./embedded-installer/lib/utility: time stamp 2020-12-10 15:25:16 is
49482950.508847367 s in the future
tar: ./embedded-installer/lib/check-onie: time stamp 2020-12-10 15:25:16
is 49482950.508761477 s in the future
tar: ./embedded-installer/lib: time stamp 2020-12-10 15:25:47 is
49482981.508710647 s in the future
tar: ./embedded-installer/storage/blk: time stamp 2020-12-10 15:25:16 is
49482950.508631277 s in the future
tar: ./embedded-installer/storage/gpt: time stamp 2020-12-10 15:25:16 is
49482950.508523097 s in the future
tar: ./embedded-installer/storage/init: time stamp 2020-12-10 15:25:16
is 49482950.508437507 s in the future
tar: ./embedded-installer/storage/mbr: time stamp 2020-12-10 15:25:16 is
49482950.508371177 s in the future
tar: ./embedded-installer/storage/mtd: time stamp 2020-12-10 15:25:16 is
49482950.508293856 s in the future
tar: ./embedded-installer/storage: time stamp 2020-12-10 15:25:16 is
49482950.508243666 s in the future
tar: ./embedded-installer/platforms.db: time stamp 2020-12-10 15:25:16
is 49482950.508179456 s in the future
tar: ./embedded-installer/install: time stamp 2020-12-10 15:25:47 is
49482981.508094606 s in the future
tar: ./embedded-installer: time stamp 2020-12-10 15:25:47 is
49482981.508044066 s in the future
tar: ./control: time stamp 2021-01-30 17:00:58 is 53895092.507984316 s
in the future
tar: .: time stamp 2021-01-30 17:00:58 is 53895092.507920196 s in the
```

```
future
Staging installer image...done.
WARNING:
WARNING: Activating staged installer requested.
WARNING: This action will wipe out all system data.
WARNING: Make sure to back up your data.
WARNING:
Are you sure (y/N)? y
Activating staged installer...done.
Reboot required to take effect.
cumulus@IP_switch_A_1:mgmt:~$
```

2. Respond `y` to the prompt to confirm the installation when the image is downloaded and verified.
3. Reboot the switch to install the new software: `sudo reboot`

```
cumulus@IP_switch_A_1:mgmt:~$ sudo reboot
```



The switch reboots and enters the switch software installation which takes some time. When the installation is complete, the switch reboots and remains at the 'log-in' prompt.

4. Configure the basic switch settings
 - a. When the switch is booted and at the log-in prompt, log in and change the password.



The username is 'cumulus' and the default password is 'cumulus'.

```
Debian GNU/Linux 10 cumulus ttyS0

cumulus login: cumulus
Password:
You are required to change your password immediately (administrator
enforced)
Changing password for cumulus.
Current password:
New password:
Retype new password:
Linux cumulus 4.19.0-cl-1-amd64 #1 SMP Cumulus 4.19.206-1+cl4.4.3u1
(2021-12-18) x86_64

Welcome to NVIDIA Cumulus (R) Linux (R)

For support and online technical documentation, visit
http://www.cumulusnetworks.com/support

The registered trademark Linux (R) is used pursuant to a sublicense from
LMI,
the exclusive licensee of Linus Torvalds, owner of the mark on a world-
wide
basis.

cumulus@cumulus:~$
```

5. Configure the management network interface.

The commands you use depend on the switch firmware version you are running.



The following example commands configure the hostname as `IP_switch_A_1`, the IP address as `10.10.10.10`, the netmask as `255.255.255.0 (24)`, and the gateway address as `10.10.10.1`.

Cumulus 4.4.x

The following example commands configure the hostname, IP address, netmask, and gateway on a switch running Cumulus 4.4.x.

```
cumulus@cumulus:mgmt:~$ net add hostname IP_switch_A_1
cumulus@cumulus:mgmt:~$ net add interface eth0 ip address
10.0.10.10/24
cumulus@cumulus:mgmt:~$ net add interface eth0 ip gateway 10.10.10.1
cumulus@cumulus:mgmt:~$ net pending
```

```
.
.
.
```

```
cumulus@cumulus:mgmt:~$ net commit
```

```
.
.
.
```

net add/del commands since the last "net commit"

User Timestamp Command

```
cumulus 2021-05-17 22:21:57.437099 net add hostname Switch-A-1
cumulus 2021-05-17 22:21:57.538639 net add interface eth0 ip address
10.10.10.10/24
cumulus 2021-05-17 22:21:57.635729 net add interface eth0 ip gateway
10.10.10.1
```

```
cumulus@cumulus:mgmt:~$
```

Cumulus 5.4.x and later

The following example commands configure the hostname, IP address, netmask, and gateway on a switch running Cumulus 5.4.x. or later.

```
cumulus@cumulus:mgmt:~$ nv set system hostname IP_switch_A_1

cumulus@cumulus:mgmt:~$ nv set interface eth0 ip address
10.0.10.10/24

cumulus@cumulus:mgmt:~$ nv set interface eth0 ip gateway 10.10.10.1

cumulus@cumulus:mgmt:~$ nv config apply

cumulus@cumulus:mgmt:~$ nv config save
```

6. Reboot the switch using the `sudo reboot` command.

```
cumulus@cumulus:~$ sudo reboot
```

When the switch reboots, you can apply a new configuration using the steps in [Download and install the NVIDIA RCF file](#).

Download and install the NVIDIA RCF files

You must generate and install the switch RCF file to each switch in the MetroCluster IP configuration.

Before you begin

- You must have the root password for `sudo` access to the commands.
- The switch software is installed and the management network is configured.
- You followed the steps to initially install the switch by using either method 1 or method 2.
- You did not apply any additional configuration after the initial installation.



If you perform further configuration after resetting the switch and before applying the RCF file, you cannot use this procedure.

About this task

You must repeat these steps on each of the IP switches in the MetroCluster IP configuration (new installation) or on the replacement switch (switch replacement).

If you are using a QSFP-to-SFP+ adapter, you might need to configure the ISL port in native speed mode instead of breakout speed mode. Refer to your switch vendor documentation to determine the ISL port speed mode.

Steps

1. Generate the NVIDIA RCF files for MetroCluster IP.
 - a. Download the [RcfFileGenerator for MetroCluster IP](#).
 - b. Generate the RCF file for your configuration by using the RcfFileGenerator for MetroCluster IP.

c. Navigate to your home directory. If you are logged as 'cumulus', the file path is /home/cumulus.

```
cumulus@IP_switch_A_1:mgmt:~$ cd ~
cumulus@IP_switch_A_1:mgmt:~$ pwd
/home/cumulus
cumulus@IP_switch_A_1:mgmt:~$
```

d. Download the RCF file to this directory. The following example shows that you use SCP to download the file SN2100_v2.0.0_IP_switch_A_1.txt from server '50.50.50.50' to your home directory and save it as SN2100_v2.0.0_IP_switch_A_1.py:

```
cumulus@Switch-A-1:mgmt:~$ scp
username@50.50.50.50:/RcfFiles/SN2100_v2.0.0_IP_switch_A_1.txt
./SN2100_v2.0.0_IP_switch-A1.py
The authenticity of host '50.50.50.50 (50.50.50.50)' can't be
established.
RSA key fingerprint is
SHA256:B5gBtOmNZvdKiY+dPhh8=ZK9DaKG7g6sv+2gFlGVF8E.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '50.50.50.50' (RSA) to the list of known
hosts.
*****
**
Banner of the SCP server
*****
**
username@50.50.50.50's password:
SN2100_v2.0.0_IP_switch_A1.txt 100% 55KB 1.4MB/s 00:00
cumulus@IP_switch_A_1:mgmt:~$
```

2. Execute the RCF file. The RCF file requires an option to apply one or more steps. Unless instructed by technical support, run the RCF file without the command line option. To verify the completion status of the various steps of the RCF file, use the option '-1' or 'all' to apply all (pending) steps.

```

cumulus@IP_switch_A_1:mgmt:~$ sudo python3
SN2100_v2.0.0_IP_switch_A_1.py
all
[sudo] password for cumulus:
The switch will be rebooted after the step(s) have been run.
Enter yes or no: yes

... the steps will apply - this is generating a lot of output ...

Running Step 24: Final reboot of the switch

... The switch will reboot if all steps applied successfully ...

```

3. If your configuration uses DAC cables, enable the DAC option on the switch ports:

```

cumulus@IP_switch_A_1:mgmt:~$ sudo python3 SN2100_v2.0.0-X10_Switch-
A1.py runCmd <switchport> DacOption [enable | disable]

```

The following example enables the DAC option for port swp7:

```

cumulus@IP_switch_A_1:mgmt:~$ sudo python3 SN2100_v2.00_Switch-A1.py
runCmd swp7 DacOption enable
Running cumulus version : 5.4.0
Running RCF file version : v2.00
Running command: Enabling the DacOption for port swp7
runCmd: 'nv set interface swp7 link fast-linkup on', ret: 0
runCmd: committed, ret: 0
Completion: SUCCESS
cumulus@IP_switch_A_1:mgmt:~$

```

4. Reboot the switch after enabling the DAC option on the switch ports:

```
sudo reboot
```



When you set the DAC option for multiple switch ports, you only need to reboot the switch once.

Set Forward Error Correction for systems using 25-Gbps connectivity

If your system is configured using 25-Gbps connectivity, set the Forward Error Correction (fec) parameter manually to off after applying the RCF. The RCF does not apply this setting.

About this task

- This task only applies to platforms using 25-Gbps connectivity. Refer to [Platform port assignments for NVIDIA supported SN2100 IP switches](#).
- This task must be performed on all four switches in the MetroCluster IP configuration.
- You must update each switch port individually, you cannot specify multiple ports or port ranges in the command.

Steps

1. Set the `fec` parameter to off for the first switch port that uses 25-Gbps connectivity:

```
sudo python3 SN2100_v2.0_Switch-A1.py runCmd <switchport> fec off
```

2. Repeat the step for each 25-Gbps switch port that is connected to a controller module.

Set the switch port speed for the MetroCluster IP interfaces

About this task

- Use this procedure to set the switch port speed to 100G for the following systems:
 - AFF A70, AFF A90, AFF A1K, AFF C80
 - AFF A30, AFF C30, AFF A50, AFF C60
 - FAS50, FAS70, FAS90
- You must update each switch port individually, you cannot specify multiple ports or port ranges in the command.

Step

1. Use the RCF file with the `runCmd` option to set the speed. This applies the setting and saves the configuration.

The following commands set the speed for the MetroCluster interfaces `swp7` and `swp8`:

```
sudo python3 SN2100_v2.20 _Switch-A1.py runCmd swp7 speed 100
```

```
sudo python3 SN2100_v2.20 _Switch-A1.py runCmd swp8 speed 100
```

Example

```
cumulus@Switch-A-1:mgmt:~$ sudo python3 SN2100_v2.20_Switch-A1.py runCmd
swp7 speed 100
[sudo] password for cumulus: <password>
Running cumulus version : 5.4.0
Running RCF file version : v2.20
Running command: Setting switchport swp7 to 100G speed
runCmd: 'nv set interface swp7 link auto-negotiate off', ret: 0
runCmd: 'nv set interface swp7 link speed 100G', ret: 0
runCmd: committed, ret: 0
Completion: SUCCESS
cumulus@Switch-A-1:mgmt:~$
```

Disable unused ISL ports and port channels

NetApp recommends disabling unused ISL ports and port channels to avoid unnecessary health alerts. You must disable each port or port channel individually, you cannot specify multiple ports or port ranges in the command.

Steps

1. Identify the unused ISL ports and port channels using the RCF file banner:



If the port is in breakout mode, the port name you specify in the command might be different than the name stated in the RCF banner. You can also use the RCF cabling files to find the port name.

```
net show interface
```

2. Disable the unused ISL ports and port channels using the RCF file.

```

cumulus@mcc1-integrity-a1:mgmt:~$ sudo python3 SN2100_v2.0_IP_Switch-
A1.py runCmd
[sudo] password for cumulus:
    Running cumulus version   : 5.4.0
    Running RCF file version  : v2.0
Help for runCmd:
    To run a command execute the RCF script as follows:
    sudo python3 <script> runCmd <option-1> <option-2> <option-x>
    Depending on the command more or less options are required. Example
to 'up' port 'swp1'
    sudo python3 SN2100_v2.0_IP_Switch-A1.py runCmd swp1 up
Available commands:
    UP / DOWN the switchport
        sudo python3 SN2100_v2.0_IP_Switch-A1.py runCmd <switchport>
state <up | down>
    Set the switch port speed
        sudo python3 SN2100_v2.0_Switch-A1.py runCmd <switchport>
speed <10 | 25 | 40 | 100 | AN>
    Set the fec mode on the switch port
        sudo python3 SN2100_v2.0_Switch-A1.py runCmd <switchport>
fec <default | auto | rs | baser | off>
    Set the [localISL | remoteISL] to 'UP' or 'DOWN' state
        sudo python3 SN2100_v2.0_Switch-A1.py runCmd [localISL |
remoteISL] state [up | down]
    Set the option on the port to support DAC cables. This option
does not support port ranges.
    You must reload the switch after changing this option for
the required ports. This will disrupt traffic.
    This setting requires Cumulus 5.4 or a later 5.x release.
        sudo python3 SN2100_v2.0_Switch-A1.py runCmd <switchport>
DacOption [enable | disable]
cumulus@mcc1-integrity-a1:mgmt:~$

```

The following example command disables port "swp14":

```
sudo python3 SN2100_v2.0_Switch-A1.py runCmd swp14 state down
```

Repeat this step for each identified unused port or port channel.

Install the Ethernet Switch Health Monitor configuration file for a NVIDIA SN2100 MetroCluster IP switch

To configure Ethernet switch health monitoring on NVIDIA Ethernet switches, follow this procedure.

These instructions apply if NVIDIA X190006-PE and X190006-PI switches are not detected properly, which can

be confirmed by running `system switch ethernet show` and checking if **OTHER** is shown for your model. To identify your NVIDIA switch model, find the part-number with the command `nv show platform hardware` for NVIDIA CL 5.8 and earlier or `nv show platform` for later versions.



These steps are also recommended if you want health monitoring and log collection to work as intended when using NVIDIA CL 5.11.x with the following ONTAP releases. While health monitoring and log collection might still function without these steps, following them ensures everything operates correctly.

- 9.10.1P20, 9.11.1P18, 9.12.1P16, 9.13.1P8, 9.14.1, 9.15.1 and later patch releases

Before you begin

- Make sure that the ONTAP cluster is up and running.
- Enable SSH on the switch to use all of the features available in CSHM.
- Clear the `/mroot/etc/cshm_nod/nod_sign/` directory on all nodes:

- a. Enter the nodeshell:

```
system node run -node <name>
```

- b. Change to advanced privilege:

```
priv set advanced
```

- c. List the configuration files in the `/etc/cshm_nod/nod_sign` directory. If the directory exists and contains configuration files, it lists the file names.

```
ls /etc/cshm_nod/nod_sign
```

- d. Delete all configuration files corresponding to your connected switch models.

If you are unsure, remove all configuration files for the supported models listed above, then download and install the latest configuration files for those same models.

```
rm /etc/cshm_nod/nod_sign/<filename>
```

- e. Confirm that the deleted configuration files are no longer in the directory:

```
ls /etc/cshm_nod/nod_sign
```

Steps

1. Download the Ethernet switch health monitor configuration zip file based on the corresponding ONTAP release version. This file is available from the [NVIDIA Ethernet switches](#) page.
 - a. On the NVIDIA SN2100 Software download page, select **Nvidia CSHM File**.
 - b. On the Caution/Must read page, select the check box to agree.
 - c. On the End User License Agreement page, select the check box to agree and click **Accept & Continue**.
 - d. On the Nvidia CSHM File - Download page, select the applicable configuration file. The following files are available:

ONTAP 9.15.1 and later

- MSN2100-CB2FC-v1.4.zip
- MSN2100-CB2RC-v1.4.zip
- X190006-PE-v1.4.zip
- X190006-PI-v1.4.zip

ONTAP 9.11.1 through 9.14.1

- MSN2100-CB2FC_PRIOR_R9.15.1-v1.4.zip
- MSN2100-CB2RC_PRIOR_R9.15.1-v1.4.zip
- X190006-PE_PRIOR_9.15.1-v1.4.zip
- X190006-PI_PRIOR_9.15.1-v1.4.zip

2. Upload the applicable zip file to your internal web server.
3. Access the advanced mode setting from one of the ONTAP systems in the cluster.

```
set -privilege advanced
```

4. Run the switch health monitor configure command.

```
cluster1::> system switch ethernet configure-health-monitor
```

5. Verify that the command output ends with the following text for your ONTAP version:

ONTAP 9.15.1 and later

Ethernet switch health monitoring installed the configuration file.

ONTAP 9.11.1 through 9.14.1

SHM installed the configuration file.

ONTAP 9.10.1

CSHM downloaded package processed successfully.

If an error occurs, contact NetApp support.

6. Wait up to twice the Ethernet switch health monitor polling interval, found by running `system switch ethernet polling-interval show`, before completing the next step.
7. Run the command `system switch ethernet configure-health-monitor show` on the ONTAP system and make sure that the cluster switches are discovered with the monitored field set to **True** and the serial number field not showing **Unknown**.

```
cluster1::> system switch ethernet configure-health-monitor show
```



If your model is still showing **OTHER** after applying the configuration file, contact NetApp support.

See the [system switch ethernet configure-health-monitor](#) command for further details.

What's next?

[Configure switch health monitoring.](#)

Monitor MetroCluster IP switch health

Learn about switch health monitoring in a MetroCluster IP configuration

The Ethernet switch health monitor (CSHM) is responsible for ensuring the operational health of Cluster and Storage network switches and collecting switch logs for debugging purposes.

Important notes for configuring CSHM in a MetroCluster IP configuration

This section contains the generic steps for configuring SNMPv3 and log collection on Cisco, Broadcom, and NVIDIA SN2100 switches. You must follow the steps for a switch firmware version that is supported in a MetroCluster IP configuration. Refer to the [Hardware Universe](#) to verify the supported firmware versions.

In a MetroCluster configuration, you configure health monitoring on the local cluster switches only.

For log collection with Broadcom and Cisco switches, a new user should be created on the switch for each cluster with log collection enabled. In a MetroCluster configuration, this means that MetroCluster 1, MetroCluster 2, MetroCluster 3, and MetroCluster 4 all require a separate user to be configured on the switches. These switches do not support multiple SSH keys for the same user. Any additional log collection setup performed overwrites any pre-existing SSH keys for the user.

Before you configure the CSHM, you should disable unused ISLs to avoid any unnecessary ISL alerts.

Configure SNMPv3 to monitor the health of MetroCluster IP switches

In MetroCluster IP configurations, you can configure SNMPv3 to monitor the health of IP switches.

This procedure shows the generic steps for configuring SNMPv3 on a switch. Some of the switch firmware versions listed might not be supported in a MetroCluster IP configuration.

You must follow the steps for a switch firmware version that is supported in a MetroCluster IP configuration. Refer to the [Hardware Universe](#) to verify the supported firmware versions.

- SNMPv3 is only supported on ONTAP 9.12.1 and later.
- ONTAP 9.13.1P12, 9.14.1P9, 9.15.1P5, 9.16.1 and later versions fix these two issues:
 - [For ONTAP health monitoring of Cisco switches, SNMPv2 traffic might still be seen after switching to SNMPv3 for monitoring](#)
 - [False-positive switch fan and power alerts when SNMP failures occur](#)



About this task

The following commands are used to configure an SNMPv3 username on **Broadcom**, **Cisco**, and **NVIDIA** switches:

Broadcom switches

Configure an SNMPv3 username NETWORK-OPERATOR on Broadcom BES-53248 switches.

- For **no authentication**:

```
snmp-server user SNMPv3UserNoAuth NETWORK-OPERATOR noauth
```

- For **MD5/SHA authentication**:

```
snmp-server user SNMPv3UserAuth NETWORK-OPERATOR [auth-md5|auth-sha]
```

- For **MD5/SHA authentication with AES/DES encryption**:

```
snmp-server user SNMPv3UserAuthEncrypt NETWORK-OPERATOR [auth-  
md5|auth-sha] [priv-aes128|priv-des]
```

The following command configures an SNMPv3 username on the ONTAP side:

```
security login create -user-or-group-name SNMPv3_USER -application snmp  
-authentication-method usm -remote-switch-ipaddress ADDRESS
```

The following command establishes the SNMPv3 username with CSHM:

```
cluster1::*> system switch ethernet modify -device DEVICE -snmp-version  
SNMPv3 -community-or-username SNMPv3_USER
```

Steps

1. Set up the SNMPv3 user on the switch to use authentication and encryption:

```
show snmp status
```

```
(sw1) (Config)# snmp-server user <username> network-admin auth-md5
<password> priv-aes128 <password>
```

```
(cs1) (Config)# show snmp user snmp
```

Name	Group Name	Auth Meth	Priv Meth	Remote Engine ID
<username>	network-admin	MD5	AES128	8000113d03d8c497710bee

2. Set up the SNMPv3 user on the ONTAP side:

```
security login create -user-or-group-name <username> -application
snmp -authentication-method usm -remote-switch-ipaddress
10.231.80.212
```

```
cluster1::*> security login create -user-or-group-name <username>
-application snmp -authentication-method usm -remote-switch
-ipaddress 10.231.80.212
```

Enter the authoritative entity's EngineID [remote EngineID]:

Which authentication protocol do you want to choose (none, md5, sha, sha2-256)

[none]: **md5**

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]: **aes128**

Enter privacy protocol password (minimum 8 characters long):

Enter privacy protocol password again:

3. Configure CSHM to monitor with the new SNMPv3 user:

```
system switch ethernet show-all -device "sw1" -instance
```

```

cluster1::*> system switch ethernet show-all -device "sw1
(b8:59:9f:09:7c:22)" -instance

Device Name: sw1
IP Address: 10.228.136.24
SNMP Version: SNMPv2c
Is Discovered: true
DEPRECATED-Community String or SNMPv3 Username: -
Community String or SNMPv3 Username: cshml!
Model Number: BES-53248
Switch Network: cluster-network
Software Version: 3.9.0.2
Reason For Not Monitoring: None <---- should
display this if SNMP settings are valid
Source Of Switch Version: CDP/ISDP
Is Monitored ?: true
Serial Number of the Device: QTFCU3826001C
RCF Version: v1.8X2 for

Cluster/HA/RDMA

cluster1::*>
cluster1::*> system switch ethernet modify -device "sw1" -snmp
-version SNMPv3 -community-or-username <username>

```

4. After waiting the CSHM polling period, verify that the serial number is populated for the Ethernet switch.

```
system switch ethernet polling-interval show
```

```

cluster1::*> system switch ethernet polling-interval show
Polling Interval (in minutes): 5

cluster1::*> system switch ethernet show-all -device "sw1" -instance
Device Name: sw1
IP Address: 10.228.136.24
SNMP Version: SNMPv3
Is Discovered: true
DEPRECATED-Community String or SNMPv3 Username: -
Community String or SNMPv3 Username: <username>
Model Number: BES-53248
Switch Network: cluster-network
Software Version: 3.9.0.2
Reason For Not Monitoring: None <---- should
display this if SNMP settings are valid
Source Of Switch Version: CDP/ISDP
Is Monitored?: true
Serial Number of the Device: QTFCU3826001C
RCF Version: v1.8X2 for

Cluster/HA/RDMA

```

Cisco switches

Configure an SNMPv3 username SNMPv3_USER on Cisco 9336C-FX2 switches:

- For no authentication:

```
snmp-server user SNMPv3_USER NoAuth
```

- For MD5/SHA authentication:

```
snmp-server user SNMPv3_USER auth [md5|sha] AUTH-PASSWORD
```

- For MD5/SHA authentication with AES/DES encryption:

```
snmp-server user SNMPv3_USER AuthEncrypt auth [md5|sha] AUTH-
PASSWORD priv aes-128 PRIV-PASSWORD
```

The following command configures an SNMPv3 username on the ONTAP side:

```
security login create -user-or-group-name SNMPv3_USER -application snmp
-authentication-method usm -remote-switch-ipaddress ADDRESS
```

The following command establishes the SNMPv3 username with CSHM:

```
system switch ethernet modify -device DEVICE -snmp-version SNMPv3
-community-or-username SNMPv3_USER
```

Steps

1. Set up the SNMPv3 user on the switch to use authentication and encryption:

```
show snmp user
```

```
(sw1) (Config)# snmp-server user SNMPv3User auth md5 <auth_password>
priv aes-128 <priv_password>
```

```
(sw1) (Config)# show snmp user
```

```
-----
-----
```

SNMP USERS

```
-----
-----
```

User	Auth	Priv(enforce)	Groups
acl_filter			
admin	md5	des(no)	network-admin
SNMPv3User	md5	aes-128(no)	network-operator

```
-----
-----
```

NOTIFICATION TARGET USERS (configured for sending V3 Inform)

```
-----
-----
```

User	Auth	Priv
------	------	------

```
-----
-----
```

```
(sw1) (Config)#
```

2. Set up the SNMPv3 user on the ONTAP side:

```
security login create -user-or-group-name <username> -application
snmp -authentication-method usm -remote-switch-ipaddress
10.231.80.212
```

```
cluster1::*> system switch ethernet modify -device "sw1
(b8:59:9f:09:7c:22)" -is-monitoring-enabled-admin true
```

```
cluster1::*> security login create -user-or-group-name <username>
-application snmp -authentication-method usm -remote-switch
-ipaddress 10.231.80.212
```

Enter the authoritative entity's EngineID [remote EngineID]:

Which authentication protocol do you want to choose (none, md5, sha, sha2-256)

[none]: **md5**

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]: **aes128**

Enter privacy protocol password (minimum 8 characters long):

Enter privacy protocol password again:

3. Configure CSHM to monitor with the new SNMPv3 user:

```
system switch ethernet show-all -device "sw1" -instance
```

```

cluster1::*> system switch ethernet show-all -device "sw1" -instance

                Device Name: sw1
                IP Address: 10.231.80.212
                SNMP Version: SNMPv2c
                Is Discovered: true
                SNMPv2c Community String or SNMPv3 Username: cshml!
                Model Number: N9K-C9336C-FX2
                Switch Network: cluster-network
                Software Version: Cisco Nexus
Operating System (NX-OS) Software, Version 9.3(7)
                Reason For Not Monitoring: None <---- displays
when SNMP settings are valid
                Source Of Switch Version: CDP/ISDP
                Is Monitored?: true
                Serial Number of the Device: QTFCU3826001C
                RCF Version: v1.8X2 for

Cluster/HA/RDMA

cluster1::*>
cluster1::*> system switch ethernet modify -device "sw1" -snmp
-version SNMPv3 -community-or-username <username>
cluster1::*>

```

4. Verify that the serial number to be queried with the newly created SNMPv3 user is the same as detailed in the previous step after the CSHM polling period has completed.

```

system switch ethernet polling-interval show

```

```

cluster1::*> system switch ethernet polling-interval show
                Polling Interval (in minutes): 5

cluster1::*> system switch ethernet show-all -device "sw1" -instance

                Device Name: sw1
                IP Address: 10.231.80.212
                SNMP Version: SNMPv3
                Is Discovered: true
                SNMPv2c Community String or SNMPv3 Username: SNMPv3User
                Model Number: N9K-C9336C-FX2
                Switch Network: cluster-network
                Software Version: Cisco Nexus
                Operating System (NX-OS) Software, Version 9.3(7)
                Reason For Not Monitoring: None <---- displays
when SNMP settings are valid
                Source Of Switch Version: CDP/ISDP
                Is Monitored ?: true
                Serial Number of the Device: QTFCU3826001C
                RCF Version: v1.8X2 for
                Cluster/HA/RDMA

cluster1::*>

```

NVIDIA - CL 5.4.0

Configure an SNMPv3 username SNMPv3_USER on NVIDIA SN2100 switches running CLI 5.4.0:

- For **no authentication**:

```
nv set service snmp-server username SNMPv3_USER auth-none
```

- For **MD5/SHA authentication**:

```
nv set service snmp-server username SNMPv3_USER [auth-md5|auth-sha]
AUTH-PASSWORD
```

- For **MD5/SHA authentication with AES/DES encryption**:

```
nv set service snmp-server username SNMPv3_USER [auth-md5|auth-sha]
AUTH-PASSWORD [encrypt-aes|encrypt-des] PRIV-PASSWORD
```

The following command configures an SNMPv3 username on the ONTAP side:

```
security login create -user-or-group-name SNMPv3_USER -application snmp
-authentication-method usm -remote-switch-ipaddress ADDRESS
```

The following command establishes the SNMPv3 username with CSHM:

```
system switch ethernet modify -device DEVICE -snmp-version SNMPv3
-community-or-username SNMPv3_USER
```

Steps

1. Set up the SNMPv3 user on the switch to use authentication and encryption:

```
net show snmp status
```

```
cumulus@sw1:~$ net show snmp status
Simple Network Management Protocol (SNMP) Daemon.
-----
Current Status          active (running)
Reload Status           enabled
Listening IP Addresses  all vrf mgmt
Main snmpd PID          4318
Version 1 and 2c Community String  Configured
Version 3 Usernames     Not Configured
-----

cumulus@sw1:~$
cumulus@sw1:~$ net add snmp-server username SNMPv3User auth-md5
<password> encrypt-aes <password>
cumulus@sw1:~$ net commit
--- /etc/snmp/snmpd.conf      2020-08-02 21:09:34.686949282 +0000
+++ /run/nclu/snmp/snmpd.conf 2020-08-11 00:13:51.826126655 +0000
@@ -1,26 +1,28 @@
# Auto-generated config file: do not edit. #
agentaddress udp:@mgmt:161
agentxperms 777 777 snmp snmp
agentxsocket /var/agentx/master
createuser _snmptrapusernameX
+createuser SNMPv3User MD5 <password> AES <password>
ifmib_max_num_ifaces 500
iquerysecname _snmptrapusernameX
master agentx
monitor -r 60 -o laNames -o laErrorMessage "laTable" laErrorFlag != 0
pass -p 10 1.3.6.1.2.1.1.1 /usr/share/snmp/sysDescr_pass.py
pass_persist 1.2.840.10006.300.43
```

```

/usr/share/snmp/ieee8023_lag_pp.py
  pass_persist 1.3.6.1.2.1.17 /usr/share/snmp/bridge_pp.py
  pass_persist 1.3.6.1.2.1.31.1.1.1.18
/usr/share/snmp/snmpifAlias_pp.py
  pass_persist 1.3.6.1.2.1.47 /usr/share/snmp/entity_pp.py
  pass_persist 1.3.6.1.2.1.99 /usr/share/snmp/entity_sensor_pp.py
  pass_persist 1.3.6.1.4.1.40310.1 /usr/share/snmp/resq_pp.py
  pass_persist 1.3.6.1.4.1.40310.2
/usr/share/snmp/cl_drop_cntrs_pp.py
  pass_persist 1.3.6.1.4.1.40310.3 /usr/share/snmp/cl_poe_pp.py
  pass_persist 1.3.6.1.4.1.40310.4 /usr/share/snmp/bgpun_pp.py
  pass_persist 1.3.6.1.4.1.40310.5 /usr/share/snmp/cumulus-status.py
  pass_persist 1.3.6.1.4.1.40310.6 /usr/share/snmp/cumulus-sensor.py
  pass_persist 1.3.6.1.4.1.40310.7 /usr/share/snmp/vrf_bgpun_pp.py
+rocommunity cshml! default
  rouser _snmptrapusernameX
+rouser SNMPv3User priv
  sysobjectid 1.3.6.1.4.1.40310
  sysservices 72
-rocommunity cshml! default

```

net add/del commands since the last "net commit"

```

User          Timestamp          Command
-----
-----
SNMPv3User   2020-08-11 00:13:51.826987 net add snmp-server username
SNMPv3User auth-md5 <password> encrypt-aes <password>

```

```

cumulus@sw1:~$
cumulus@sw1:~$ net show snmp status
Simple Network Management Protocol (SNMP) Daemon.
-----
Current Status          active (running)
Reload Status           enabled
Listening IP Addresses  all vrf mgmt
Main snmpd PID          24253
Version 1 and 2c Community String Configured
Version 3 Usernames     Configured      <---- Configured
here
-----
cumulus@sw1:~$

```

2. Set up the SNMPv3 user on the ONTAP side:

```
security login create -user-or-group-name SNMPv3User -application
snmp -authentication-method usm -remote-switch-ipaddress
10.231.80.212
```

```
cluster1::*> security login create -user-or-group-name SNMPv3User
-application snmp -authentication-method usm -remote-switch
-ipaddress 10.231.80.212
```

Enter the authoritative entity's EngineID [remote EngineID]:

Which authentication protocol do you want to choose (none, md5, sha, sha2-256)

[none]: **md5**

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]: **aes128**

Enter privacy protocol password (minimum 8 characters long):

Enter privacy protocol password again:

3. Configure CSHM to monitor with the new SNMPv3 user:

```
system switch ethernet show-all -device "sw1 (b8:59:9f:09:7c:22)"
-instance
```

```

cluster1::*> system switch ethernet show-all -device "sw1
(b8:59:9f:09:7c:22)" -instance
                                     Device Name: sw1
(b8:59:9f:09:7c:22)
                                     IP Address: 10.231.80.212
                                     SNMP Version: SNMPv2c
                                     Is Discovered: true
DEPRECATED-Community String or SNMPv3 Username: -
      Community String or SNMPv3 Username: cshml!
      Model Number: MSN2100-CB2FC
      Switch Network: cluster-network
      Software Version: Cumulus Linux
version 5.4.0 running on Mellanox Technologies Ltd. MSN2100
      Reason For Not Monitoring: None
      Source Of Switch Version: LLDP
      Is Monitored ?: true
      Serial Number of the Device: MT2110X06399 <----
serial number to check
      RCF Version: MSN2100-RCF-v1.9X6-
Cluster-LLDP Aug-18-2022

cluster1::*>
cluster1::*> system switch ethernet modify -device "sw1
(b8:59:9f:09:7c:22)" -snmp-version SNMPv3 -community-or-username
SNMPv3User

```

4. Verify that the serial number to be queried with the newly created SNMPv3 user is the same as detailed in the previous step after the CSHM polling period has completed.

```
system switch ethernet polling-interval show
```

```

cluster1::*> system switch ethernet polling-interval show
Polling Interval (in minutes): 5

cluster1::*> system switch ethernet show-all -device "sw1
(b8:59:9f:09:7c:22)" -instance
Device Name: sw1
IP Address: 10.231.80.212
SNMP Version: SNMPv3
Is Discovered: true
DEPRECATED-Community String or SNMPv3 Username: -
Community String or SNMPv3 Username: SNMPv3User
Model Number: MSN2100-CB2FC
Switch Network: cluster-network
Software Version: Cumulus Linux
version 5.4.0 running on Mellanox Technologies Ltd. MSN2100
Reason For Not Monitoring: None
Source Of Switch Version: LLDP
Is Monitored?: true
Serial Number of the Device: MT2110X06399 <----
serial number to check
RCF Version: MSN2100-RCF-v1.9X6-
Cluster-LLDP Aug-18-2022

```

NVIDIA - CL 5.11.0

Configure an SNMPv3 username SNMPv3_USER on NVIDIA SN2100 switches running CLI 5.11.0:

- For no authentication:

```
nv set system snmp-server username SNMPv3_USER auth-none
```

- For MD5/SHA authentication:

```
nv set system snmp-server username SNMPv3_USER [auth-md5|auth-sha]
AUTH-PASSWORD
```

- For MD5/SHA authentication with AES/DES encryption:

```
nv set system snmp-server username SNMPv3_USER [auth-md5|auth-sha]
AUTH-PASSWORD [encrypt-aes|encrypt-des] PRIV-PASSWORD
```

The following command configures an SNMPv3 username on the ONTAP side:

```
security login create -user-or-group-name SNMPv3_USER -application snmp
-authentication-method usm -remote-switch-ipaddress ADDRESS
```

The following command establishes the SNMPv3 username with CSHM:

```
system switch ethernet modify -device DEVICE -snmp-version SNMPv3
-community-or-username SNMPv3_USER
```

Steps

1. Set up the SNMPv3 user on the switch to use authentication and encryption:

```
nv show system snmp-server
```

```
cumulus@sw1:~$ nv show system snmp-server
                                applied
-----
[username]                       SNMPv3_USER
[username]                       limiteduser1
[username]                       testuserauth
[username]                       testuserauthaes
[username]                       testusernoauth
trap-link-up
  check-frequency                 60
trap-link-down
  check-frequency                 60
[listening-address]              all
[readonly-community]              $nvsec$94d69b56e921aec1790844eb53e772bf
state                             enabled
cumulus@sw1:~$
```

2. Set up the SNMPv3 user on the ONTAP side:

```
security login create -user-or-group-name SNMPv3User -application
snmp -authentication-method usm -remote-switch-ipaddress
10.231.80.212
```

```
cluster1::*> security login create -user-or-group-name SNMPv3User
-application snmp -authentication-method usm -remote-switch
-ipaddress 10.231.80.212
```

Enter the authoritative entity's EngineID [remote EngineID]:

Which authentication protocol do you want to choose (none, md5, sha, sha2-256)

[none]: **md5**

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]: **aes128**

Enter privacy protocol password (minimum 8 characters long):

Enter privacy protocol password again:

3. Configure CSHM to monitor with the new SNMPv3 user:

```
system switch ethernet show-all -device "sw1 (b8:59:9f:09:7c:22)"
-instance
```

```

cluster1::*> system switch ethernet show-all -device "sw1
(b8:59:9f:09:7c:22)" -instance
                                     Device Name: sw1
(b8:59:9f:09:7c:22)
                                     IP Address: 10.231.80.212
                                     SNMP Version: SNMPv2c
                                     Is Discovered: true
DEPRECATED-Community String or SNMPv3 Username: -
      Community String or SNMPv3 Username: cshml!
      Model Number: MSN2100-CB2FC
      Switch Network: cluster-network
      Software Version: Cumulus Linux
version 5.11.0 running on Mellanox Technologies Ltd. MSN2100
      Reason For Not Monitoring: None
      Source Of Switch Version: LLDP
      Is Monitored ?: true
      Serial Number of the Device: MT2110X06399 <----
serial number to check
      RCF Version: MSN2100-RCF-v1.9X6-
Cluster-LLDP Aug-18-2022

cluster1::*>
cluster1::*> system switch ethernet modify -device "sw1
(b8:59:9f:09:7c:22)" -snmp-version SNMPv3 -community-or-username
SNMPv3User

```

4. Verify that the serial number to be queried with the newly created SNMPv3 user is the same as detailed in the previous step after the CSHM polling period has completed.

```
system switch ethernet polling-interval show
```

```

cluster1::*> system switch ethernet polling-interval show
Polling Interval (in minutes): 5

cluster1::*> system switch ethernet show-all -device "sw1
(b8:59:9f:09:7c:22)" -instance
Device Name: sw1
IP Address: 10.231.80.212
SNMP Version: SNMPv3
Is Discovered: true
DEPRECATED-Community String or SNMPv3 Username: -
Community String or SNMPv3 Username: SNMPv3User
Model Number: MSN2100-CB2FC
Switch Network: cluster-network
Software Version: Cumulus Linux
version 5.11.0 running on Mellanox Technologies Ltd. MSN2100
Reason For Not Monitoring: None
Source Of Switch Version: LLDP
Is Monitored?: true
Serial Number of the Device: MT2110X06399 <----
serial number to check
RCF Version: MSN2100-RCF-v1.9X6-
Cluster-LLDP Aug-18-2022

```

Configure log collection on a MetroCluster IP switch

In a MetroCluster IP configuration, you can configure log collection to collect switch logs for debugging purposes.



On Broadcom and Cisco switches, a new user is required for each cluster with log collection. For example, MetroCluster 1, MetroCluster 2, MetroCluster 3, and MetroCluster 4 all require a separate user to be configured on the switches. Multiple SSH keys for the same user is not supported.

About this task

The Ethernet switch health monitor (CSHM) is responsible for ensuring the operational health of Cluster and Storage network switches and collecting switch logs for debugging purposes. This procedure guides you through the process of setting up collection, requesting detailed **Support** logs, and enabling an hourly collection of **Periodic** data that is collected by AutoSupport.

NOTE: If you enable FIPS mode, you must complete the following:



1. Regenerate SSH keys on the switch using the vendor instructions.
2. Regenerate SSH keys in ONTAP using `debug system regenerate-systemshell-key-pair`
3. Re-run log collection setup routine using the `system switch ethernet log setup-password` command

Before you begin

- The user must have access to the switch `show` commands. If these are not available, create a new user and grant the necessary permissions to the user.
- Switch health monitoring must be enabled for the switch. Verify this by ensuring the `Is Monitored:` field is set to **true** in the output of the `system switch ethernet show` command.
- For log collection with Broadcom and Cisco switches:
 - The local user must have network admin privileges.
 - A new user should be created on the switch for each cluster setup with log collection enabled. These switches do not support multiple SSH keys for the same user. Any additional log collection setup performed overwrites any pre-existing SSH keys for the user.
- For support log collection with NVIDIA switches, the *user* for log collection must be permitted to run the `cl-support` command without having to provide a password. To allow this usage, run the command:

```
echo '<user> ALL = NOPASSWD: /usr/cumulus/bin/cl-support' | sudo EDITOR='tee -a' visudo -f /etc/sudoers.d/cumulus
```

Steps

ONTAP 9.15.1 and later

1. To set up log collection, run the following command for each switch. You are prompted to enter the switch name, username, and password for log collection.

NOTE: If answering **y** to the user specification prompt, make sure that the user has the necessary permissions as outlined in [Before you begin](#).

```
system switch ethernet log setup-password
```

```
cluster1::*> system switch ethernet log setup-password
```

```
Enter the switch name: <return>
```

```
The switch name entered is not recognized.
```

```
Choose from the following list:
```

```
cs1
```

```
cs2
```

```
cluster1::*> system switch ethernet log setup-password
```

```
Enter the switch name: cs1
```

```
Would you like to specify a user other than admin for log  
collection? {y|n}: n
```

```
Enter the password: <enter switch password>
```

```
Enter the password again: <enter switch password>
```

```
cluster1::*> system switch ethernet log setup-password
```

```
Enter the switch name: cs2
```

```
Would you like to specify a user other than admin for log  
collection? {y|n}: n
```

```
Enter the password: <enter switch password>
```

```
Enter the password again: <enter switch password>
```



For CL 5.11.1, create the user **cumulus** and respond **y** to the following prompt: Would you like to specify a user other than admin for log collection? {y|n}: **y**

2. Enable periodic log collection:

```
system switch ethernet log modify -device <switch-name> -periodic  
-enabled true
```

```
cluster1::*> system switch ethernet log modify -device cs1 -periodic
-enabled true
```

Do you want to modify the cluster switch log collection configuration? {y|n}: [n] **y**

cs1: Periodic log collection has been scheduled to run every hour.

```
cluster1::*> system switch ethernet log modify -device cs2 -periodic
-enabled true
```

Do you want to modify the cluster switch log collection configuration? {y|n}: [n] **y**

cs2: Periodic log collection has been scheduled to run every hour.

```
cluster1::*> system switch ethernet log show
```

	Periodic	Periodic
Support		
Switch	Log Enabled	Log State
Log State		
cs1	true	scheduled
never-run		
cs2	true	scheduled
never-run		

2 entries were displayed.

3. Request support log collection:

```
system switch ethernet log collect-support-log -device <switch-name>
```

```
cluster1::*> system switch ethernet log collect-support-log -device cs1
```

```
cs1: Waiting for the next Ethernet switch polling cycle to begin support collection.
```

```
cluster1::*> system switch ethernet log collect-support-log -device cs2
```

```
cs2: Waiting for the next Ethernet switch polling cycle to begin support collection.
```

```
cluster1::*> *system switch ethernet log show
```

	Periodic	Periodic
Support		
Switch	Log Enabled	Log State
Log State		
cs1	false	halted
initiated		
cs2	true	scheduled
initiated		

2 entries were displayed.

4. To view all details of log collection, including the enablement, status message, previous timestamp and filename of periodic collection, the request status, status message, and previous timestamp and filename of support collection, use the following:

```
system switch ethernet log show -instance
```

```

cluster1::*> system switch ethernet log show -instance

                Switch Name: cs1
    Periodic Log Enabled: true
        Periodic Log Status: Periodic log collection has been
scheduled to run every hour.
    Last Periodic Log Timestamp: 3/11/2024 11:02:59
        Periodic Log Filename: cluster1:/mroot/etc/log/shm-
cluster-info.tgz
    Support Log Requested: false
        Support Log Status: Successfully gathered support logs
- see filename for their location.
    Last Support Log Timestamp: 3/11/2024 11:14:20
        Support Log Filename: cluster1:/mroot/etc/log/shm-
cluster-log.tgz

                Switch Name: cs2
    Periodic Log Enabled: false
        Periodic Log Status: Periodic collection has been
halted.
    Last Periodic Log Timestamp: 3/11/2024 11:05:18
        Periodic Log Filename: cluster1:/mroot/etc/log/shm-
cluster-info.tgz
    Support Log Requested: false
        Support Log Status: Successfully gathered support logs
- see filename for their location.
    Last Support Log Timestamp: 3/11/2024 11:18:54
        Support Log Filename: cluster1:/mroot/etc/log/shm-
cluster-log.tgz
2 entries were displayed.

```

ONTAP 9.14.1 and earlier

1. To set up log collection, run the following command for each switch. You are prompted to enter the switch name, username, and password for log collection.

NOTE: If answering `y` to the user specification prompt, make sure that the user has the necessary permissions as outlined in [Before you begin](#).

```
system switch ethernet log setup-password
```

```
cluster1::*> system switch ethernet log setup-password
```

```
Enter the switch name: <return>
```

```
The switch name entered is not recognized.
```

```
Choose from the following list:
```

```
cs1
```

```
cs2
```

```
cluster1::*> system switch ethernet log setup-password
```

```
Enter the switch name: cs1
```

```
Would you like to specify a user other than admin for log  
collection? {y|n}: n
```

```
Enter the password: <enter switch password>
```

```
Enter the password again: <enter switch password>
```

```
cluster1::*> system switch ethernet log setup-password
```

```
Enter the switch name: cs2
```

```
Would you like to specify a user other than admin for log  
collection? {y|n}: n
```

```
Enter the password: <enter switch password>
```

```
Enter the password again: <enter switch password>
```



For CL 5.11.1, create the user **cumulus** and respond **y** to the following prompt: Would you like to specify a user other than admin for log collection? {y|n}: **y**

2. To request support log collection and enable periodic collection, run the following command. This starts both types of log collection: the detailed Support logs and an hourly collection of Periodic data.

```
system switch ethernet log modify -device <switch-name> -log-request  
true
```

```
cluster1::*> system switch ethernet log modify -device cs1 -log
-request true
```

```
Do you want to modify the cluster switch log collection
configuration? {y|n}: [n] y
```

```
Enabling cluster switch log collection.
```

```
cluster1::*> system switch ethernet log modify -device cs2 -log
-request true
```

```
Do you want to modify the cluster switch log collection
configuration? {y|n}: [n] y
```

```
Enabling cluster switch log collection.
```

Wait for 10 minutes and then check that the log collection completes:

```
system switch ethernet log show
```



If any error statuses are reported by the log collection feature (visible in the output of `system switch ethernet log show`), see [Troubleshoot log collection](#) for further details.

Manage the monitoring of Ethernet switches in a MetroCluster IP configuration

In most cases, Ethernet switches are automatically discovered by ONTAP and monitored by CSHM. The Reference Configuration File (RCF) applied to the switch, among other things, enables the Cisco Discovery Protocol (CDP) and/or the Link Layer Discovery Protocol (LLDP). However, you might need to manually add a switch that is not discovered or remove a switch that is no longer in use. You can also stop active monitoring while retaining the switch in the configuration, such as during maintenance.

Create a switch entry so that ONTAP can monitor it

About this task

Use the `system switch ethernet create` command to manually configure and enable monitoring for a specified Ethernet switch. This is helpful if ONTAP does not add the switch automatically, or if you previously removed the switch and want to re-add it.

```
system switch ethernet create -device DeviceName -address 1.2.3.4 -snmp
-version SNMPv2c -community-or-username cshml! -model NX3132V -type
cluster-network
```

A typical example is adding a switch named [DeviceName], with IP address 1.2.3.4, and SNMPv2c credentials set to **cs hm1!**. Use `-type storage-network` instead of `-type cluster-network` if you are configuring a storage switch.

Disable monitoring without deleting the switch

If you want to pause or stop monitoring for a certain switch, but still retain it for future monitoring, modify its `is-monitoring-enabled-admin` parameter instead of deleting it.

For example:

```
system switch ethernet modify -device DeviceName -is-monitoring-enabled
-admin false
```

This lets you preserve switch details and configuration without generating new alerts or re-discoveries.

Remove a switch you no longer need

Use `system switch ethernet delete` to delete a switch that has been disconnected or is no longer required:

```
system switch ethernet delete -device DeviceName
```

By default, this command succeeds only if ONTAP does not currently detect the switch through CDP or LLDP. To remove a discovered switch, use the `-force` parameter:

```
system switch ethernet delete -device DeviceName -force
```

When `-force` is used, the switch might be re-added automatically if ONTAP detects it again.

Verify Ethernet switch monitoring in a MetroCluster IP configuration

The Ethernet switch health monitor (CSHM) automatically attempts to monitor the switches that it discovers; however, monitoring might not happen automatically if the switches are not configured correctly. You should verify that the health monitor is properly configured to monitor your switches.

Confirm monitoring of the connected Ethernet switches

About this task

To confirm that the connected Ethernet switches are being monitored, run:

```
system switch ethernet show
```

If the `Model` column displays **OTHER** or the `IS Monitored` field displays **false**, then ONTAP cannot monitor

the switch. A value of **OTHER** typically indicates that ONTAP does not support that switch for health monitoring.

The `IS Monitored` field is set to **false** for the reason specified in the `Reason` field.



If a switch is not listed in the command output, ONTAP likely has not discovered it. Confirm that the switch is cabled correctly. If necessary, you can add the switch manually. See [Manage the monitoring of Ethernet Switches](#) for further details.

Confirm firmware and RCF versions are up to date

Make sure that the switch is running the latest supported firmware and that a compatible reference configuration file (RCF) has been applied. More information is available on the [NetApp Support Downloads page](#).

By default, the health monitor uses SNMPv2c with the community string **csbm1!** for monitoring, but SNMPv3 can also be configured.

If you need to change the default SNMPv2c community string, make sure that the desired SNMPv2c community string has been configured on the switch.

```
system switch ethernet modify -device SwitchA -snmp-version SNMPv2c
-community-or-username newCommunity!
```



See [Optional: Configure SNMPv3](#) for details on configuring SNMPv3 for use.

Confirm management network connection

Verify that the switch's management port is connected to the management network.

A correct management port connection is required for ONTAP to perform SNMP queries and log collection.

Related information

- [Troubleshoot alerts](#)

Copyright information

Copyright © 2026 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.