



Maintenance procedures for MetroCluster IP configurations

ONTAP MetroCluster

NetApp
October 01, 2024

Table of Contents

- Maintenance procedures for MetroCluster IP configurations 1
 - Modify the properties of a MetroCluster IP interface 1
 - IP switch maintenance and replacement 6
 - Identifying storage in a MetroCluster IP configuration 31
 - Adding shelves to a MetroCluster IP using shared Storage MetroCluster switches 35
 - Configure end-to-end encryption in a MetroCluster IP configuration 51
 - Power off and power on a single site in a MetroCluster IP configuration 55
 - Powering off an entire MetroCluster IP configuration 61

Maintenance procedures for MetroCluster IP configurations

Modify the properties of a MetroCluster IP interface

Beginning with ONTAP 9.10.1, you can change the following properties of a MetroCluster IP interface: IP address and mask, and gateway. You can use any combination of parameters to update.

You might need to update these properties, for example, if a duplicate IP address is detected or if a gateway needs to change in the case of a layer 3 network due to router configuration changes.

Modify the IP address, netmask, and gateway

The procedure you follow depends on whether you are using ONTAP System Manager or the CLI.

System Manager

Use System Manager to modify the IP address, netmask, and gateway properties.

Step

Update the IP address, netmask, and gateway for each node and interface.

CLI

Use the CLI to modify the IP address, netmask, and gateway properties.

About this task

- You can only change one interface at a time. There will be traffic disruption on that interface until the other interfaces are updated and connections are reestablished.
- Use the `metrocluster configuration-settings interface modify` command to change any MetroCluster IP interface property.



These commands change the configuration on a particular node for a particular port. To restore complete network connectivity, similar commands are needed on other ports. Similarly, network switches also need to update their configuration. For example, if the gateway is updated, ideally it is changed on both nodes of an HA pair, since they are same. The switch connected to those nodes also needs to update its gateway.

- Use the `metrocluster configuration-settings interface show`, `metrocluster connection check`, and `metrocluster connection show` commands to verify that all connectivity is working in all interfaces.

Steps

1. Update the IP address, netmask, and gateway for a single node and interface: `metrocluster configuration-settings interface modify`

The following command shows how to update the IP address, netmask and gateway:

```

cluster_A::* metrocluster configuration-settings interface modify
-cluster-name cluster_A -home-node node_A_1 -home-port e0a-10
-address 192.168.12.101 -gateway 192.168.12.1 -netmask 255.255.254.0
(metrocluster configuration-settings interface modify)
Warning: This operation will disconnect and reconnect iSCSI and RDMA
connections used for DR protection through port "e0a-10". Partner
nodes may need modifications for port "e0a-10" in order to
completely establish network connectivity.
Do you want to continue?" yes
[Job 28] Setting up iSCSI target configuration. (pass2:iscsil3:0:-
1:0): xpt_action_default: CCB type 0xe XPT_DEV_ADVINFO not supported
[Job 28] Establishing iSCSI initiator connections.
(pass6:iscsil4:0:-1:0): xpt_action_default: CCB type 0xe
XPT_DEV_ADVINFO not supported
(pass8:iscsil5:0:-1:0): xpt_action_default: CCB type 0xe
XPT_DEV_ADVINFO not supported
(pass9:iscsil6:0:-1:0): xpt_action_default: CCB type 0xe
XPT_DEV_ADVINFO not supported
[Job 28] Job succeeded: Interface Modify is successful.
cluster_A::*> metrocluster configuration-settings interface modify
-cluster-name cluster_A -home-node node_A_2 -home-port e0a-10
-address 192.168.12.201 -gateway 192.168.12.1 -netmask 255.255.254.0
(metrocluster configuration-settings interface modify)
Warning: This operation will disconnect and reconnect iSCSI and RDMA
connections used for DR protection through port "e0a-10". Partner
nodes may need modifications for port "e0a-10" in order to
completely establish network connectivity.
Do you want to continue?" yes
[Job 28] Job succeeded: Interface Modify is successful

```

2. Verify that all connectivity is working for all interfaces: metrocluster configuration-settings interface show

The following command shows how to verify that all connectivity is working for all interfaces:

```

cluster_A::*> metrocluster configuration-settings interface show
(metrocluster configuration-settings interface show)
DR          Config
Group Cluster Node      Network Address Netmask      Gateway
State
-----
-----
1      cluster_A node_A_2
          Home Port: e0a-10
          192.168.12.201  255.255.254.0  192.168.12.1
completed
          Home Port: e0b-20
          192.168.20.200  255.255.255.0  192.168.20.1
completed
          node_A_1
          Home Port: e0a-10
          192.168.12.101  255.255.254.0  192.168.12.1
completed
          Home Port: e0b-20
          192.168.20.101  255.255.255.0  192.168.20.1
completed
      cluster_B node_B_1
          Home Port: e0a-10
          192.168.11.151  255.255.255.0  192.168.11.1
completed
          Home Port: e0b-20
          192.168.21.150  255.255.255.0  192.168.21.1
completed
          node_B_2
          Home Port: e0a-10
          192.168.11.250  255.255.255.0  192.168.11.1
completed
          Home Port: e0b-20
          192.168.21.250  255.255.255.0  192.168.21.1
completed
8 entries were displayed.

```

3. Verify that all connections are working:

```
metrocluster configuration-settings connection show
```

The following command shows how to verify that all connections are working:

```

cluster_A::*> metrocluster configuration-settings connection show
(metrocluster configuration-settings connection show)
DR
Group Cluster Node      Source          Destination
Config State           Network Address Network Address Partner Type
-----
1      cluster_A node_A_2
      Home Port: e0a-10
      192.168.10.200 192.168.10.101 HA Partner
completed
      Home Port: e0a-10
      192.168.10.200 192.168.11.250 DR Partner
completed
      Home Port: e0a-10
      192.168.10.200 192.168.11.151 DR Auxiliary
completed
      Home Port: e0b-20
      192.168.20.200 192.168.20.100 HA Partner
completed
      Home Port: e0b-20
      192.168.20.200 192.168.21.250 DR Partner
completed
      Home Port: e0b-20
      192.168.20.200 192.168.21.150 DR Auxiliary
completed
      node_A_1
      Home Port: e0a-10
      192.168.10.101 192.168.10.200 HA Partner
completed
      Home Port: e0a-10
      192.168.10.101 192.168.11.151 DR Partner
completed
      Home Port: e0a-10
      192.168.10.101 192.168.11.250 DR Auxiliary
completed
      Home Port: e0b-20
      192.168.20.100 192.168.20.200 HA Partner
completed
      Home Port: e0b-20
      192.168.20.100 192.168.21.150 DR Partner
completed
      Home Port: e0b-20
      192.168.20.100 192.168.21.250 DR Auxiliary
completed

```

IP switch maintenance and replacement

Replace an IP switch or change the use of existing MetroCluster IP switches

You might need to replace a failed switch, upgrade or downgrade a switch, or change the use of existing MetroCluster IP switches.

About this task

This procedure applies when you are using NetApp-validated switches. If you are using MetroCluster-compliant switches, refer to the switch vendor.

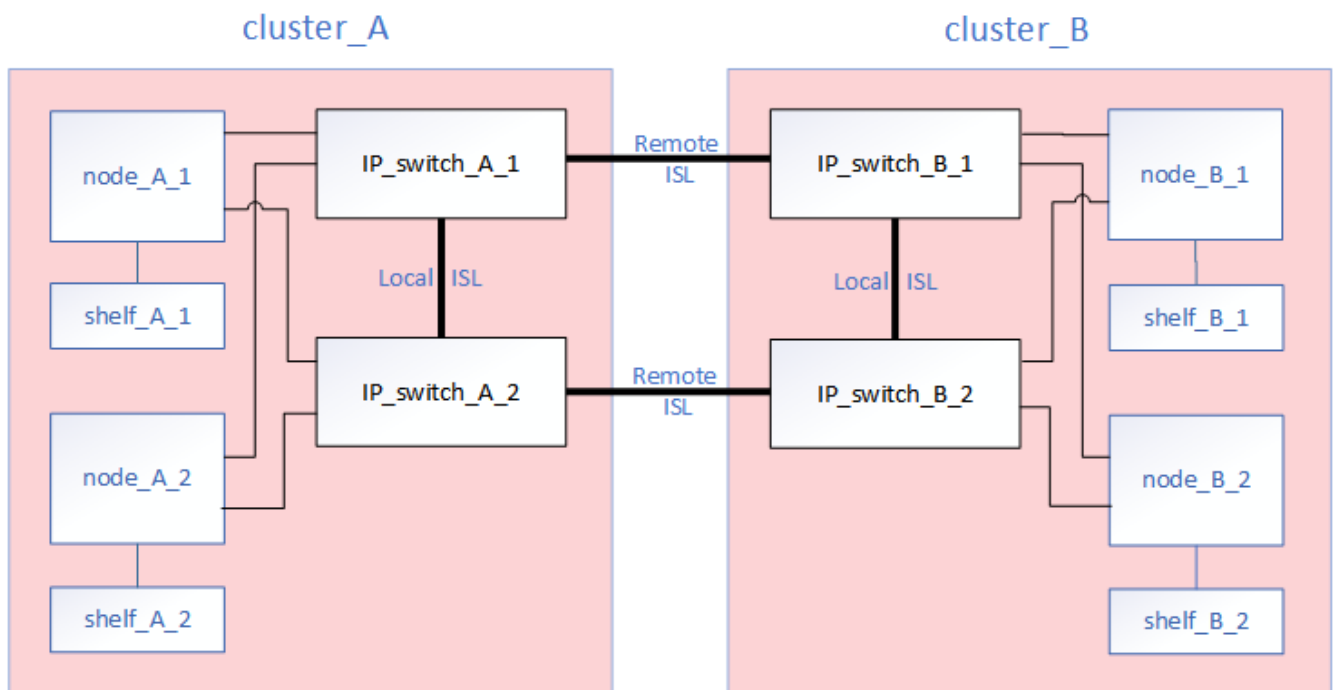
[Enable console logging](#) before performing this task.

This procedure supports the following conversions:

- Changing the switch vendor, type, or both. The new switch can be the same as the old switch when a switch has failed, or you can change the switch type (upgrade or downgrade the switch).

For example, to expand a MetroCluster IP configuration from a single four-node configuration using AFF A400 controllers and BES-53248 switches to an eight-node configuration using AFF A400 controllers, you must change the switches to a supported type for the configuration because BES-53248 switches are not supported in the new configuration.

If you want to replace a failed switch with the same type of switch, you only replace the failed switch. If you want to upgrade or downgrade a switch, you must adjust two switches that are in the same network. Two switches are in the same network when they are connected with an inter-switch link (ISL) and are not located at the same site. For example, Network 1 includes IP_switch_A_1 and IP_switch_B_1, and Network 2 includes IP_switch_A_2 and IP_switch_B_2, as shown in the diagram below:



If you replace a switch or upgrade to different switches, then you can pre-configure the switches by installing the switch firmware and RCF file.

- Convert a MetroCluster IP configuration to a MetroCluster IP configuration using shared storage MetroCluster switches.

For example, if you have a regular MetroCluster IP configuration using AFF A700 controllers and you want to reconfigure the MetroCluster to connect NS224 shelves to the same switches.



- If you are adding or removing shelves in a MetroCluster IP configuration using shared storage MetroCluster IP switches, follow the steps in [Adding shelves to a MetroCluster IP using shared storage MetroCluster switches](#)
- Your MetroCluster IP configuration might already directly connect to NS224 shelves or to dedicated storage switches.

Port usage worksheet

The following is an example worksheet for converting a MetroCluster IP configuration to a shared storage configuration connecting two NS224 shelves using the existing switches.

Worksheet definitions:

- Existing configuration: The cabling of the existing MetroCluster configuration.
- New configuration with NS224 shelves: The target configuration where the switches are shared between storage and the MetroCluster.

The highlighted fields in this worksheet indicate the following:

- Green: You do not need to change the cabling.
- Yellow: You must move ports with the same or a different configuration.
- Blue: Ports that are new connections.

PORT USAGE OVERVIEW

Example of expanding an existing 4Node MetroCluster with 2x NS224 shelves and changing the ISL's from 10G to 40/100G

Switch port	Existing configuration			New configuration with NS224 shelves		
	Port use	IP_switch_x_1	IP_switch_x_2	Port use	IP_switch_x_1	IP_switch_x_2
1	MetroCluster 1, Local Cluster Interface	Cluster Port 'A'	Cluster Port 'B'	MetroCluster 1, Local Cluster Interface	Cluster Port 'A'	Cluster Port 'B'
2		Cluster Port 'A'	Cluster Port 'B'		Cluster Port 'A'	Cluster Port 'B'
3						
4						
5				Storage shelf 1 (9)	NSM-A, e0a	NSM-A, e0b
6					NSM-B, e0a	NSM-B, e0b
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
8						
9	MetroCluster 1, MetroCluster interface	Port 'A'	Port 'B'	MetroCluster 1, MetroCluster interface	Port 'A'	Port 'B'
10		Port 'A'	Port 'B'		Port 'A'	Port 'B'
11						
12						
13				ISL, MetroCluster, native speed 40G / 100G breakout mode 10G	Remote ISL, 2x 40/100G	Remote ISL, 2x 40/100G
14						
15						
16						
17				MetroCluster 1, Storage Interface	Storage Port 'A'	Storage Port 'B'
18					Storage Port 'A'	Storage Port 'B'
19						
20						
21	ISL, MetroCluster breakout mode 10G	Remote ISL, 10G	Remote ISL, 10G	Storage shelf 2 (8)	NSM-A, e0a	NSM-A, e0b
22					NSM-B, e0a	NSM-B, e0b
23						
24						
25						
26						
27						
28						
29						
30						
31						
32						
33						
34						
35						
36						

Steps

1. Check the health of the configuration.
 - a. Check that the MetroCluster is configured and in normal mode on each cluster: **metrocluster show**

```
cluster_A::> metrocluster show
Cluster                               Entry Name                               State
-----                               -
Local: cluster_A                       Configuration state configured
Mode                                    normal
AUSO Failure Domain auso-on-cluster-
disaster
Remote: cluster_B                       Configuration state configured
Mode                                    normal
AUSO Failure Domain auso-on-cluster-
disaster
```

- b. Check that mirroring is enabled on each node: **metrocluster node show**

```
cluster_A::> metrocluster node show
DR                                     Configuration DR
Group Cluster Node                    State          Mirroring Mode
-----
1   cluster_A
    node_A_1    configured    enabled    normal
    cluster_B
    node_B_1    configured    enabled    normal
2 entries were displayed.
```

- c. Check that the MetroCluster components are healthy: **metrocluster check run**

```
cluster_A::> metrocluster check run
```

```
Last Checked On: 10/1/2014 16:03:37
```

Component	Result
nodes	ok
lifs	ok
config-replication	ok
aggregates	ok

4 entries were displayed.

Command completed. Use the "metrocluster check show -instance" command or sub-commands in "metrocluster check" directory for detailed results.

To check if the nodes are ready to do a switchover or switchback operation, run "metrocluster switchover -simulate" or "metrocluster switchback -simulate", respectively.

d. Check that there are no health alerts: **system health alert show**

2. Configure the new switch before installation.

If you are reusing existing switches, go to [Step 4](#).



If you are upgrading or downgrading the switches, you must configure all the switches in the network.

Follow the steps in the section *Configuring the IP switches* in the [MetroCluster IP installation and configuration](#).

Make sure that you apply the correct RCF file for switch `_A_1`, `_A_2`, `_B_1` or `_B_2`. If the new switch is the same as the old switch, you need to apply the same RCF file.

If you upgrade or downgrade a switch, apply the latest supported RCF file for the new switch.

3. Run the port show command to view information about the network ports:

network port show

a. Modify all cluster LIFs to disable auto-revert:

```
network interface modify -vserver <vserver_name> -lif <lif_name>  
-auto-revert false
```

4. Disconnect the connections from the old switch.



You only disconnect connections that are not using the same port in the old and new configurations. If you are using new switches, you must disconnect all connections.

Remove the connections in the following order:

- a. Disconnect the local cluster interfaces
- b. Disconnect the local cluster ISLs
- c. Disconnect the MetroCluster IP interfaces
- d. Disconnect the MetroCluster ISLs

In the example [Port usage worksheet](#), the switches do not change. The MetroCluster ISLs are relocated and must be disconnected. You do not need to disconnect the connections marked in green on the worksheet.

5. If you are using new switches, power off the old switch, remove the cables, and physically remove the old switch.

If you are reusing existing switches, go to [Step 6](#).



Do **not** cable the new switches except for the management interface (if used).

6. Configure the existing switches.

If you have pre-configured the switches already, you can skip this step.

To configure the existing switches, follow the steps to install and upgrade the firmware and RCF files:

- [Upgrading firmware on MetroCluster IP switches](#)
- [Upgrade RCF files on MetroCluster IP switches](#)

7. Cable the switches.

You can follow the steps in the *Cabling the IP switches* section in [MetroCluster IP installation and configuration](#).

Cable the switches in the following order (if required):

- a. Cable the ISLs to the remote site.
- b. Cable the MetroCluster IP interfaces.
- c. Cable the local cluster interfaces.



- The used ports might be different from those on the old switch if the switch type is different. If you are upgrading or downgrading the switches, do **NOT** cable the local ISLs. Only cable the local ISLs if you are upgrading or downgrading the switches in the second network and both switches at one site are the same type and cabling.
- If you are upgrading Switch-A1 and Switch-B1, you must perform steps 1 to 6 for switches Switch-A2 and Switch-B2.

8. Finalize the local cluster cabling.

- a. If the local cluster interfaces are connected to a switch:

- i. Cable the local cluster ISLs.
 - b. If the local cluster interfaces are **not** connected to a switch:
 - i. Use the [Migrate to a switched NetApp cluster environment](#) procedure to convert a switchless cluster to a switched cluster. Use the ports indicated in [MetroCluster IP installation and configuration](#) or the RCF cabling files to connect the local cluster interface.
9. Power up the switch or switches.

If the new switch is the same, power up the new switch. If you are upgrading or downgrading the switches, then power up both switches. The configuration can operate with two different switches at each site until the second network is updated.

10. Verify that the MetroCluster configuration is healthy by repeating [Step 1](#).

If you are upgrading or downgrading the switches in the first network, you might see some alerts related to local clustering.



If you upgrade or downgrade the networks, then repeat all of the steps for the second network.

11. Modify all cluster LIFs to re-enable auto-revert:

```
network interface modify -vserver <vserver_name> -lif <lif_name> -auto
-revert true
```

12. Optionally, move the NS224 shelves.

If you are reconfiguring a MetroCluster IP configuration that does not connect NS224 shelves to the MetroCluster IP switches, use the appropriate procedure to add or move the NS224 shelves:

- [Adding shelves to a MetroCluster IP using shared storage MetroCluster switches](#)
- [Migrate from a switchless cluster with direct-attached storage](#)
- [Migrate from a switchless configuration with switch-attached storage by reusing the storage switches](#)

Online or offline MetroCluster IP interface ports

When you perform maintenance tasks, you might need to bring a MetroCluster IP interface port offline or online.

About this task

[Enable console logging](#) before performing this task.

Steps

You can use the following steps to bring a MetroCluster IP interface port online or take it offline.

1. Set the privilege level to advanced.

```
set -privilege advanced
```

Example output

```
Cluster_A_1::> set -privilege advanced
Warning: These advanced commands are potentially dangerous; use them
only when
        directed to do so by NetApp personnel.
Do you want to continue? {y|n}: y
```

2. Take the MetroCluster IP interface port offline.

```
system ha interconnect link off -node <node_name> -link <link_num, 0 or
1>
```

Example output

```
Cluster_A1::*> system ha interconnect link off -node node-a1 -link 0
```

a. Verify the MetroCluster IP interface is offline.

```
Cluster_A1::*> system ha interconnect port show
```

Example output

```

Cluster_A1::*> system ha interconnect port show

```

Physical Node Down	Link Active Link	Monitor	Port	Physical Layer State	Link Layer State	Physical Link Up	Link
node-a1	off			0 disabled	down	4	
3 false				1 linkup	active	4	
2 true				0 linkup	active	4	
node-a2	off			1 linkup	active	4	
2 true							

2 entries were displayed.

3. Bring the MetroCluster IP interface port online.

```

system ha interconnect link on -node <node_name> -link <link_num, 0 or 1>

```

Example output

```

Cluster_A1::*> system ha interconnect link on -node node-a1 -link 0

```

a. Verify the MetroCluster IP interface port is online.

```

Cluster_A1::*> system ha interconnect port show

```

Example output


```

Cluster_A1::*> system ha interconnect port show
                Physical  Link
                Layer    Layer    Physical
Physical  Active
Node      Link Monitor  Port  State  State  Link Up  Link
Down  Link
-----  -----  ----  -----  -----  -----
node-a1      off
                0  linkup  active  5
3  true
                1  linkup  active  4
2  true
node-a2      off
                0  linkup  active  4
2  true
                1  linkup  active  4
2  true
2 entries were displayed.

```

Upgrading firmware on MetroCluster IP switches

You might need to upgrade the firmware on a MetroCluster IP switch.

About this task

You must repeat this task on each of the switches in succession.

[Enable console logging](#) before performing this task.

Steps

1. Check the health of the configuration.
 - a. Check that the MetroCluster is configured and in normal mode on each cluster:

```
metrocluster show
```

```

cluster_A::> metrocluster show
Cluster                               Entry Name                               State
-----                               -
Local: cluster_A                       Configuration state configured
Mode                                    normal
AUSO Failure Domain auso-on-cluster-
disaster
Remote: cluster_B                       Configuration state configured
Mode                                    normal
AUSO Failure Domain auso-on-cluster-
disaster

```

b. Check that mirroring is enabled on each node:

```
metrocluster node show
```

```

cluster_A::> metrocluster node show
DR                                     Configuration DR
Group Cluster Node                     State           Mirroring Mode
-----
-----
1      cluster_A
      node_A_1      configured     enabled   normal
      cluster_B
      node_B_1      configured     enabled   normal
2 entries were displayed.

```

c. Check that the MetroCluster components are healthy:

```
metrocluster check run
```

```
cluster_A::> metrocluster check run
```

```
Last Checked On: 10/1/2014 16:03:37
```

Component	Result
nodes	ok
lifs	ok
config-replication	ok
aggregates	ok

4 entries were displayed.

Command completed. Use the "metrocluster check show -instance" command or sub-commands in "metrocluster check" directory for detailed results.

To check if the nodes are ready to do a switchover or switchback operation, run "metrocluster switchover -simulate" or "metrocluster switchback -simulate", respectively.

d. Check that there are no health alerts:

```
system health alert show
```

2. Install the software on the first switch.



You must install the switch software on the switches in the following order: switch_A_1, switch_B_1, switch_A_2, switch_B_2.

Follow the steps for installing switch software in the relevant topic depending on whether the switch type is Broadcom, Cisco, or NVIDIA:

- [Download and install the Broadcom switch EFOS software](#)
- [Download and install the Cisco switch NX-OS software](#)
- [Download and install the NVIDIA SN2100 switch Cumulus software](#)

3. Repeat the previous step for each of the switches.

4. Repeat [Step 1](#) to check the health of the configuration.

Upgrade RCF files on MetroCluster IP switches

You might need to upgrade an RCF file on a MetroCluster IP switch. For example, if the RCF file version that you are running on the switches is not supported by the ONTAP version, the switch firmware version, or both.

Verify that the RCF file is supported

If you are changing the ONTAP version or the switch firmware version, you should verify that you have an RCF

file that is supported for that version. If you use the RCF generator, the correct RCF file will be generated for you.

Steps

1. Use the following commands from the switches to verify the version of the RCF file:

From this switch...	Issue this command...
Broadcom switch	(IP_switch_A_1) # show clibanner
Cisco switch	IP_switch_A_1# show banner motd

For either switch, find the line in the output that indicates the version of the RCF file. For example, the following output is from a Cisco switch, which indicates the RCF file version is “v1.80”.

```
Filename : NX3232_v1.80_Switch-A2.txt
```

2. To check which files are supported for a specific ONTAP version, switch, and platform, use the RcfFileGenerator. If you can generate the RCF file for the configuration that you have or that you want to upgrade to, then it is supported.
3. To verify that the switch firmware is supported, refer to the following:
 - [Hardware Universe](#)
 - [NetApp Interoperability Matrix](#)

Upgrade RCF files

If you are installing new switch firmware, you must install the switch firmware before upgrading the RCF file.

About this task

- This procedure disrupts traffic on the switch where the RCF file is upgraded. Traffic will resume once the new RCF file is applied.
- Perform the steps on one switch at a time, in the following order: Switch_A_1, Switch_B_1, Switch_A_2, Switch_B_2.
- [Enable console logging](#) before performing this task.

Steps

1. Verify the health of the configuration.
 - a. Verify that the MetroCluster components are healthy:

```
metrocluster check run
```

```
cluster_A::*> metrocluster check run
```

The operation runs in the background.

- a. After the `metrocluster check run` operation completes, run `metrocluster check show` to view the results.

After approximately five minutes, the following results are displayed:

```
-----  
::*> metrocluster check show  
  
Component          Result  
-----  
nodes              ok  
lifs               ok  
config-replication ok  
aggregates        ok  
clusters          ok  
connections        not-applicable  
volumes           ok  
7 entries were displayed.
```

- b. Check the status of the running MetroCluster check operation:

```
metrocluster operation history show -job-id 38
```

- c. Verify that there are no health alerts:

```
system health alert show
```

2. Prepare the IP switches for the application of the new RCF files.

Follow the steps for your switch vendor:

- [Reset the Broadcom IP switch to factory defaults](#)
- [Reset the Cisco IP switch to factory defaults](#)
- [Reset the NVIDIA IP SN2100 switch to factory defaults](#)

3. Download and install the IP RCF file, depending on your switch vendor.

- [Download and install the Broadcom IP RCF files](#)
- [Download and install the Cisco IP RCF files](#)
- [Download and install the NVIDIA IP RCF files](#)




If you have an L2 shared or L3 network configuration, you might need to adjust the ISL ports on the intermediate/customer switches. The switchport mode might change from 'access' to 'trunk' mode. Only proceed to upgrade the second switch pair (A_2, B_2) if the network connectivity between switches A_1 and B_1 is fully operational and the network is healthy.

Upgrade RCF files on Cisco IP switches using CleanUpFiles

You might need to upgrade an RCF file on a Cisco IP switch. For example, an ONTAP upgrade or a switch firmware upgrade both require a new RCF file.

About this task

- Beginning with RcfFileGenerator version 1.4a, there is a new option to change (upgrade, downgrade, or replace) the switch configuration on Cisco IP switches without the need to perform a 'write erase'.
- [Enable console logging](#) before performing this task.
- The Cisco 9336C-FX2 switch has two different switch storage types that are named differently in the RCF. Use the following table to determine the correct Cisco 9336C-FX2 storage type for your configuration:

If you are connecting the following storage...	Choose the Cisco 9336C-FX2 storage type...	Sample RCF file banner/MOTD
<ul style="list-style-type: none">• Directly connected SAS shelves• Directly connected NVMe shelves• NVMe shelves connected to dedicated storage switches	9336C-FX2 – Direct Storage only	* Switch : NX9336C (direct storage, L2 Networks, direct ISL)
<ul style="list-style-type: none">• Directly connected SAS shelves• NVMe shelves connected to the MetroCluster IP switches <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-top: 10px;"> At least one Ethernet connected NVMe shelf is required</div>	9336C-FX2 – SAS and Ethernet storage	* Switch : NX9336C (SAS and Ethernet storage, L2 Networks, direct ISL)

Before you begin

You can use this method if your configuration meets the following requirements:

- The standard RCF configuration is applied.
- The [RcfFileGenerator](#) must be able to create the same RCF file that is applied, with the same version and configuration (platforms, VLANs).
- The RCF file that is applied was not provided by NetApp for a special configuration.
- The RCF file was not altered before it was applied.
- The steps to reset the switch to factory defaults were followed before applying the current RCF file.
- No changes were made to the switch(port) configuration after the RCF was applied.

If you do not meet these requirements, then you cannot use the CleanUpFiles that are created when generating the RCF files. However, you can leverage the function to create generic CleanUpFiles — the cleanup using this method is derived from the output of `show running-config` and is best practice.



You must update the switches in the following order: Switch_A_1, Switch_B_1, Switch_A_2, Switch_B_2. Or, you can update the switches Switch_A_1 and Switch_B_1 at the same time followed by switches Switch_A_2 and Switch_B_2.

Steps

1. Determine the current RCF file version, and which ports and VLANs are used: `IP_switch_A_1# show banner motd`



You need to get this information from all four switches and complete the following information table.

```
* NetApp Reference Configuration File (RCF)
*
* Switch : NX9336C (SAS storage, L2 Networks, direct ISL)
* Filename : NX9336_v1.81_Switch-A1.txt
* Date : Generator version: v1.3c_2022-02-24_001, file creation time:
2021-05-11, 18:20:50
*
* Platforms : MetroCluster 1 : FAS8300, AFF-A400, FAS8700
*             MetroCluster 2 : AFF-A320, FAS9000, AFF-A700, AFF-A800
* Port Usage:
* Ports 1- 2: Intra-Cluster Node Ports, Cluster: MetroCluster 1, VLAN
111
* Ports 3- 4: Intra-Cluster Node Ports, Cluster: MetroCluster 2, VLAN
151
* Ports 5- 6: Ports not used
* Ports 7- 8: Intra-Cluster ISL Ports, local cluster, VLAN 111, 151
* Ports 9-10: MetroCluster 1, Node Ports, VLAN 119
* Ports 11-12: MetroCluster 2, Node Ports, VLAN 159
* Ports 13-14: Ports not used
* Ports 15-20: MetroCluster-IP ISL Ports, VLAN 119, 159, Port Channel 10
* Ports 21-24: MetroCluster-IP ISL Ports, VLAN 119, 159, Port Channel
11, breakout mode 10gx4
* Ports 25-30: Ports not used
* Ports 31-36: Ports not used
*
#
IP_switch_A_1#
```

From this output, you must collect the information shown in the following two tables.

Generic information	MetroCluster	Data
RCF file version		1.81

Switch type		NX9336
Network typology		L2 Networks, direct ISL
Storage type		SAS storage
Platforms	1	AFF A400
	2	FAS9000

VLAN information	Network	MetroCluster configuration	Switchports	Site A	Site B
VLAN local cluster	Network 1	1	1, 2	111	222
		2	3, 4	151	251
	Network 2	1	1, 2	111	222
		2	3, 4	151	251
VLAN MetroCluster	Network 1	1	9, 10	119	119
		2	11, 12	159	159
	Network 2	1	9, 10	219	219
		2	11, 12	259	259

2. Create the RCF files and CleanUpFiles, or create generic CleanUpFiles for the current configuration.

If your configuration meets the requirements outlined in the prerequisites, select **Option 1**. If your configuration does **not** meet the requirements outlined in the prerequisites, select **Option 2**.

Option 1: Create the RCF files and CleanUpFiles

Use this procedure if the configuration meets the requirements.

Steps

- a. Use the RcfFileGenerator 1.4a (or later) to create the RCF files with the information that you retrieved in Step 1. The new version of the RcfFileGenerator creates an additional set of CleanUpFiles that you can use to revert some configuration and prepare the switch to apply a new RCF configuration.
- b. Compare the banner motd with the RCF files that are currently applied. The platform types, switch type, port and VLAN usage must be the same.



You must use the CleanUpFiles from the same version as the RCF file and for the exact same configuration. Using any CleanUpFile will not work and might require a full reset of the switch.



The ONTAP version the RCF file is created for is not relevant. Only the RCF file version is important.



The RCF file (even it is the same version) might list fewer or more platforms. Make sure that your platform is listed.

Option 2: Create generic CleanUpFiles

Use this procedure if the configuration does **not** meet all the requirements.

Steps

- a. Retrieve the output of `show running-config` from each switch.
- b. Open the RcfFileGenerator tool and click 'Create generic CleanUpFiles' at the bottom of the window
- c. Copy the output that you retrieved in Step 1 from 'one' switch into the upper window. You can remove or leave the default output.
- d. Click 'Create CUF files'.
- e. Copy the output from the lower window into a text file (this file is the CleanUpFile).
- f. Repeat Steps c, d, and e for all switches in the configuration.

At the end of this procedure, you should have four text files, one for each switch. You can use these files in the same way as the CleanUpFiles that you can create by using Option 1.

3. Create the 'new' RCF files for the new configuration. Create these files in the same way that you created the files in the previous step, except choose the respective ONTAP and RCF file version.

After completing this step you should have two sets of RCF files, each set consisting of twelve files.

4. Download the files to the bootflash.
 - a. Download the CleanUpFiles that you created in [Create the RCF files and CleanUpFiles, or create generic CleanUpFiles for the current configuration](#)



This CleanUpFile is for the current RCF file that is applied and **NOT** for the new RCF that you want to upgrade to.

Example CleanUpFile for Switch-A1: Cleanup_NX9336_v1.81_Switch-A1.txt

- b. Download the 'new' RCF files that you created in [Create the 'new' RCF files for the new configuration](#).

Example RCF file for Switch-A1: NX9336_v1.90_Switch-A1.txt

- c. Download the CleanUpFiles that you created in [Create the 'new' RCF files for the new configuration](#). This step is optional — you can use the file in future to update the switch configuration. It matches the currently applied configuration.

Example CleanUpFile for Switch-A1: Cleanup_NX9336_v1.90_Switch-A1.txt



You must use the CleanUpFile for the correct (matching) RCF version. If you use a CleanUpFile for a different RCF version, or a different configuration then the cleanup of the configuration might not work correctly.

The following example copies the three files to the bootflash:

```
IP_switch_A_1# copy sftp://user@50.50.50.50/RcfFiles/NX9336-direct-
SAS_v1.81_MetroCluster-
IP_L2Direct_A400FAS8700_xxx_xxx_xxx_xxx/Cleanup_NX9336_v1.81_Switch-
A1.txt bootflash:
IP_switch_A_1# copy sftp://user@50.50.50.50/RcfFiles/NX9336-direct-
SAS_v1.90_MetroCluster-
IP_L2Direct_A400FAS8700A900FAS9500_xxx_xxx_xxx_xxxNX9336_v1.90//NX933
6_v1.90_Switch-A1.txt bootflash:
IP_switch_A_1# copy sftp://user@50.50.50.50/RcfFiles/NX9336-direct-
SAS_v1.90_MetroCluster-
IP_L2Direct_A400FAS8700A900FAS9500_xxx_xxx_xxx_xxxNX9336_v1.90//Clean
up_NX9336_v1.90_Switch-A1.txt bootflash:
```



You are prompted to specify Virtual Routing and Forwarding (VRF).

5. Apply the CleanUpFile or generic CleanUpFile.

Some of the configuration is reverted and switchports go 'offline'.

- a. Confirm that there are no pending changes to the startup configuration: `show running-config diff`

```
IP_switch_A_1# show running-config diff
IP_switch_A_1#
```

6. If you see system output, save the running configuration to the startup configuration: `copy running-`

```
config startup-config
```



System output indicates that the startup configuration and running configuration are different and pending changes. If you do not save the pending changes, you are unable to roll back using a reload of the switch.

a. Apply the CleanUpFile:

```
IP_switch_A_1# copy bootflash:Cleanup_NX9336_v1.81_Switch-A1.txt
running-config

IP_switch_A_1#
```



The script might take a while to return to the switch prompt. No output is expected.

7. View the running configuration to verify that the configuration is cleared: `show running-config`

The current configuration should show:

- No class maps and IP access lists are configured
- No policy maps are configured
- No service policies are configured
- No port-profiles are configured
- All Ethernet interfaces (except mgmt0 which should not show any configuration, and only VLAN 1 should be configured).

If you find that any of the above items are configured, you might not be able to apply a new RCF file configuration. However, you can revert to the previous configuration by reloading the switch **without** saving the running configuration to the startup configuration. The switch will come up with the previous configuration.

8. Apply the RCF file and verify that the ports are online.

a. Apply the RCF files.

```
IP_switch_A_1# copy bootflash:NX9336_v1.90-X2_Switch-A1.txt running-
config
```



Some warning messages appear while applying the configuration. Error messages are generally not expected. However, if you are logged in using SSH, you might receive the following error: `Error: Can't disable/re-enable ssh:Current user is logged in through ssh`

b. After the configuration is applied, verify that the cluster and MetroCluster ports are coming online with one of the following commands, `show interface brief`, `show cdp neighbors`, or `show lldp neighbors`



If you changed the VLAN for the local cluster and you upgraded the first switch at the site, then cluster health monitoring might not report the state as 'healthy' because the VLANs from the old and new configurations do not match. After the second switch is updated, the state should return to healthy.

If the configuration is not applied correctly, or you do not want to keep the configuration, you can revert to the previous configuration by reloading the switch **without** saving the running configuration to startup configuration. The switch will come up with the previous configuration.

9. Save the configuration and reload the switch.

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

```
IP_switch_A_1# reload
```

Renaming a Cisco IP switch

You might need to rename a Cisco IP switch to provide consistent naming throughout your configuration.

About this task

- In the examples in this task, the switch name is changed from `myswitch` to `IP_switch_A_1`.
- [Enable console logging](#) before performing this task.

Steps

1. Enter global configuration mode:

```
configure terminal
```

The following example shows the configuration mode prompt. Both prompts show the switch name of `myswitch`.

```
myswitch# configure terminal  
myswitch(config)#
```

2. Rename the switch:

```
switchname new-switch-name
```

If you are renaming both switches in the fabric, use the same command on each switch.

The CLI prompt changes to reflect the new name:

```
myswitch(config)# switchname IP_switch_A_1  
IP_switch_A_1(config)#
```

3. Exit configuration mode:

exit

The top-level switch prompt is displayed:

```
IP_switch_A_1(config)# exit
IP_switch_A_1#
```

4. Copy the current running configuration to the startup configuration file:

copy running-config startup-config

5. Verify that the switch name change is visible from the ONTAP cluster prompt.

Note that the new switch name is shown, and the old switch name (myswitch) does not appear.

a. Enter advanced privilege mode, pressing **y** when prompted:

set -privilege advanced

b. Display the attached devices:

network device-discovery show

c. Return to admin privilege mode:

set -privilege admin

The following example shows that the switch appears with the new name, IP_switch_A_1:

```
cluster_A::storage show> set advanced
```

Warning: These advanced commands are potentially dangerous; use them only when directed to do so by NetApp personnel.

```
Do you want to continue? {y|n}: y
```

```
cluster_A::storage show*> network device-discovery show
```

Node/ Protocol Platform	Local Port	Discovered Device	Interface	

node_A_2/cdp				
	e0M	LF01-410J53.mycompany.com (SAL18516DZY)	Ethernet125/1/28	N9K-
C9372PX				
	e1a	IP_switch_A_1 (FOC21211RBU)	Ethernet1/2	N3K-
C3232C				
	e1b	IP_switch_A_1 (FOC21211RBU)	Ethernet1/10	N3K-
C3232C				
.				
.			Ethernet1/18	N9K-
C9372PX				
node_A_1/cdp				
	e0M	LF01-410J53.mycompany.com (SAL18516DZY)	Ethernet125/1/26	N9K-
C9372PX				
	e0a	IP_switch_A_2 (FOC21211RB5)	Ethernet1/1	N3K-
C3232C				
	e0b	IP_switch_A_2 (FOC21211RB5)	Ethernet1/9	N3K-
C3232C				
	e1a	IP_switch_A_1 (FOC21211RBU)		
.				
.				
.				

16 entries were displayed.

Add, remove, or change ISL ports nondisruptively on Cisco IP switches

You might need to add, remove, or change ISL ports on Cisco IP switches. You can convert dedicated ISL ports to shared ISL ports, or change the speed of ISL ports on a Cisco IP switch.

About this task

If you are converting dedicated ISL ports to shared ISL ports, ensure the new ports meet the [Requirements for shared ISL ports](#).

You must complete all the steps on both switches to ensure ISL connectivity.

The following procedure assumes you are replacing a 10-Gb ISL connected at switch port Eth1/24/1 with two 100-Gb ISLs that are connected to switch ports 17 and 18.



If you are using a Cisco 9336C-FX2 switch in a shared configuration connecting NS224 shelves, changing the ISLs might require a new RCF file. You do not require a new RCF file if your current and new ISL speed is 40Gbps and 100Gbps. All other changes to ISL speed requires a new RCF file. For example, changing the ISL speed from 40Gbps to 100Gbps does not require a new RCF file, but changing the ISL speed from 10Gbps to 40Gbps requires a new RCF file.

Before you begin

Refer to the **Switches** section of the [NetApp Hardware Universe](#) to verify the supported transceivers.

[Enable console logging](#) before performing this task.

Steps

1. Disable the ISL ports of the ISLs on both switches in the fabric that you want to change.



You only need to disable the current ISL ports if you are moving them to a different port, or the speed of the ISL is changing. If you are adding an ISL port with the same speed as the existing ISLs, go to Step 3.

You must enter only one configuration command for each line and press Ctrl-Z after you have entered all the commands, as shown in the following example:

```
switch_A_1# conf t
switch_A_1(config)# int eth1/24/1
switch_A_1(config-if)# shut
switch_A_1(config-if)#
switch_A_1#

switch_B_1# conf t
switch_B_1(config)# int eth1/24/1
switch_B_1(config-if)# shut
switch_B_1(config-if)#
switch_B_1#
```

2. Remove the existing cables and transceivers.

3. Change the ISL port as required.



If you are using Cisco 9336C-FX2 switches in a shared configuration connecting NS224 shelves, and you need to upgrade the RCF file and apply the new configuration for the new ISL ports, follow the steps to [upgrade the RCF files on MetroCluster IP switches](#).

Option	Step
To change the speed of an ISL port...	Cable the new ISLs to the designated ports according to their speeds. You must ensure that these ISL ports for your switch are listed in the <i>MetroCluster IP Installation and Configuration</i> .
To add an ISL...	Insert QFSPs into the ports you are adding as ISL ports. Ensure they are listed in the <i>MetroCluster IP Installation and Configuration</i> and cable them accordingly.

4. Enable all ISL ports (if not enabled) on both switches in the fabric beginning with the following command:

```
switch_A_1# conf t
```

You must enter only one configuration command per line and press Ctrl-Z after you have entered all the commands:

```
switch_A_1# conf t
switch_A_1(config)# int eth1/17
switch_A_1(config-if)# no shut
switch_A_1(config-if)# int eth1/18
switch_A_1(config-if)# no shut
switch_A_1(config-if)#
switch_A_1#
switch_A_1# copy running-config startup-config

switch_B_1# conf t
switch_B_1(config)# int eth1/17
switch_B_1(config-if)# no shut
switch_B_1(config-if)# int eth1/18
switch_B_1(config-if)# no shut
switch_B_1(config-if)#
switch_B_1#
switch_B_1# copy running-config startup-config
```

5. Verify that the ISLs and port channels for the ISLs are established between both switches:

```
switch_A_1# show int brief
```

You should see the ISL interfaces in the command output as shown in the following example:


```

Switch_A_1# show interface brief
-----
-----
Ethernet          VLAN    Type Mode   Status Reason          Speed
Port
Interface
Ch #
-----
-----
Eth1/17           1       eth  access down   XCVR not inserted
auto(D) --
Eth1/18           1       eth  access down   XCVR not inserted
auto(D) --
-----
-----
Port-channel      VLAN    Type Mode   Status Reason
Speed  Protocol
Interface
-----
-----
Po10              1       eth  trunk  up     none
a-100G(D) lacp
Po11              1       eth  trunk  up     none
a-100G(D) lacp

```

6. Repeat the procedure for fabric 2.

Identifying storage in a MetroCluster IP configuration

If you need to replace a drive or shelf module, you first need to identify the location.

Identification of local and remote shelves

When you view shelf information from a MetroCluster site, all remote drives are on 0m, the virtual iSCSI host adapter. This means that the drives are accessed via the MetroCluster IP interfaces. All other drives are local.

After identifying whether a shelf is remote (on 0m), you can further identify the drive or shelf by the serial number or, depending on shelf ID assignments in your configuration, by shelf ID.



In MetroCluster IP configurations running ONTAP 9.4, the shelf ID is not required to be unique between the MetroCluster sites. This includes both internal shelves (0) and external shelves. The serial number is consistent when viewed from any node on either MetroCluster site.

Shelf IDs should be unique within the disaster recovery (DR) group except for the internal shelf.

With the drive or shelf module identified, you can replace the component using the appropriate procedure.

Example of sysconfig -a output

The following example uses the `sysconfig -a` command to show the devices on a node in the MetroCluster IP configuration. This node has the following shelves and devices attached:

- slot 0: Internal drives (local drives)
- slot 3: External shelf ID 75 and 76 (local drives)
- slot 0: Virtual iSCSI host adapter 0m (remote drives)

```
node_A_1> run local sysconfig -a

NetApp Release R9.4:  Sun Mar 18 04:14:58 PDT 2018
System ID: 1111111111 (node_A_1); partner ID: 2222222222 (node_A_2)
System Serial Number: serial-number (node_A_1)
.
.
.
slot 0: NVMe Disks
      0      : NETAPP  X4001S172A1T9NTE NA01 1831.1GB 4160B/sect
(S3NBNX0J500528)
      1      : NETAPP  X4001S172A1T9NTE NA01 1831.1GB 4160B/sect
(S3NBNX0J500735)
      2      : NETAPP  X4001S172A1T9NTE NA01 1831.1GB 4160B/sect
(S3NBNX0J501165)
.
.
.
slot 3: SAS Host Adapter 3a (PMC-Sierra PM8072 rev. C, SAS, <UP>)
MFG Part Number:  Microsemi Corp. 110-03801 rev. A0
Part number:      111-03801+A0
Serial number:    7A1063AF14B
Date Code:        20170320
Firmware rev:     03.08.09.00
Base WWN:         5:0000d1:702e69e:80
Phy State:        [12] Enabled, 12.0 Gb/s
                  [13] Enabled, 12.0 Gb/s
                  [14] Enabled, 12.0 Gb/s
                  [15] Enabled, 12.0 Gb/s
Mini-SAS HD Vendor:  Molex Inc.
Mini-SAS HD Part Number:  112-00436+A0
Mini-SAS HD Type:        Passive Copper (unequalized) 0.5m ID:00
Mini-SAS HD Serial Number: 614130640
                        75.0 : NETAPP  X438_S1633400AMD NA04 381.3GB 520B/sect
(S20KNYAG501805)
```

75.1 : NETAPP X438_S1633400AMD NA04 381.3GB 520B/sect
(S20KNYAG502050)
75.2 : NETAPP X438_PHM2400MCTO NA04 381.3GB 520B/sect
(25M0A03WT2KA)
75.3 : NETAPP X438_S1633400AMD NA04 381.3GB 520B/sect
(S20KNYAG501793)
75.4 : NETAPP X438_S1633400AMD NA04 381.3GB 520B/sect
(S20KNYAG502158)

.
. .
.

Shelf 75: DS224-12 Firmware rev. IOM12 A: 0220 IOM12 B: 0220
Shelf 76: DS224-12 Firmware rev. IOM12 A: 0220 IOM12 B: 0220

slot 3: SAS Host Adapter 3c (PMC-Sierra PM8072 rev. C, SAS, <UP>)

MFG Part Number: Microsemi Corp. 110-03801 rev. A0

Part number: 111-03801+A0

Serial number: 7A1063AF14B

Date Code: 20170320

Firmware rev: 03.08.09.00

Base WWN: 5:0000d1:702e69e:88

Phy State: [0] Enabled, 12.0 Gb/s

[1] Enabled, 12.0 Gb/s

[2] Enabled, 12.0 Gb/s

[3] Enabled, 12.0 Gb/s

Mini-SAS HD Vendor: Molex Inc.

Mini-SAS HD Part Number: 112-00436+A0

Mini-SAS HD Type: Passive Copper (unequalized) 0.5m ID:00

Mini-SAS HD Serial Number: 614130691

75.0 : NETAPP X438_S1633400AMD NA04 381.3GB 520B/sect
(S20KNYAG501805)

75.1 : NETAPP X438_S1633400AMD NA04 381.3GB 520B/sect
(S20KNYAG502050)

75.2 : NETAPP X438_PHM2400MCTO NA04 381.3GB 520B/sect
(25M0A03WT2KA)

75.3 : NETAPP X438_S1633400AMD NA04 381.3GB 520B/sect
(S20KNYAG501793)

.
. .
.

Shelf 75: DS224-12 Firmware rev. IOM12 A: 0220 IOM12 B: 0220

Shelf 76: DS224-12 Firmware rev. IOM12 A: 0220 IOM12 B: 0220

slot 3: SAS Host Adapter 3d (PMC-Sierra PM8072 rev. C, SAS, <UP>)

MFG Part Number: Microsemi Corp. 110-03801 rev. A0

```

Part number:      111-03801+A0
Serial number:   7A1063AF14B
Date Code:      20170320
Firmware rev:   03.08.09.00
Base WWN:       5:0000d1:702e69e:8c
Phy State:      [4] Enabled, 12.0 Gb/s
                 [5] Enabled, 12.0 Gb/s
                 [6] Enabled, 12.0 Gb/s
                 [7] Enabled, 12.0 Gb/s

Mini-SAS HD Vendor:      Molex Inc.
Mini-SAS HD Part Number: 112-00436+A0
Mini-SAS HD Type:        Passive Copper (unequalized) 0.5m ID:01
Mini-SAS HD Serial Number: 614130690
                        75.0 : NETAPP   X438_S1633400AMD NA04 381.3GB 520B/sect
(S20KNYAG501805)
                        75.1 : NETAPP   X438_S1633400AMD NA04 381.3GB 520B/sect
(S20KNYAG502050)
                        75.2 : NETAPP   X438_PHM2400MCTO NA04 381.3GB 520B/sect
(25M0A03WT2KA)
.
.
.
Shelf 75: DS224-12  Firmware rev. IOM12 A: 0220  IOM12 B: 0220
Shelf 76: DS224-12  Firmware rev. IOM12 A: 0220  IOM12 B: 0220

slot 4: Quad 10 Gigabit Ethernet Controller X710 SFP+
.
.
.
slot 0: Virtual iSCSI Host Adapter 0m
        0.0 : NETAPP   X4001S172A1T9NTE NA01 1831.1GB 4160B/sect
(S3NBNX0J500690)
        0.1 : NETAPP   X4001S172A1T9NTE NA01 1831.1GB 4160B/sect
(S3NBNX0J500571)
        0.2 : NETAPP   X4001S172A1T9NTE NA01 1831.1GB 4160B/sect
(S3NBNX0J500323)
        0.3 : NETAPP   X4001S172A1T9NTE NA01 1831.1GB 4160B/sect
(S3NBNX0J500724)
        0.4 : NETAPP   X4001S172A1T9NTE NA01 1831.1GB 4160B/sect
(S3NBNX0J500734)
        0.5 : NETAPP   X4001S172A1T9NTE NA01 1831.1GB 4160B/sect
(S3NBNX0J500598)
        0.12 : NETAPP  X4001S172A1T9NTE NA01 1831.1GB 4160B/sect
(S3NBNX0J501094)
        0.13 : NETAPP  X4001S172A1T9NTE NA01 1831.1GB 4160B/sect
(S3NBNX0J500519)

```

```
.  
. .  
Shelf 0: FS4483PSM3E  Firmware rev. PSM3E A: 0103  PSM3E B: 0103  
Shelf 35: DS224-12   Firmware rev. IOM12 A: 0220  IOM12 B: 0220  
Shelf 36: DS224-12   Firmware rev. IOM12 A: 0220  IOM12 B: 0220  
  
node_A_1::>
```

Adding shelves to a MetroCluster IP using shared Storage MetroCluster switches

You might need to add NS224 shelves to a MetroCluster using shared Storage MetroCluster switches.

Starting from ONTAP 9.10.1, you can add NS224 shelves from a MetroCluster using the shared Storage / MetroCluster switches. You can add more than one shelf at a time.

Before you begin

- Nodes must be running ONTAP 9.9.1 or later.
- All currently connected NS224 shelves must be attached to the same switches as the MetroCluster (shared Storage / MetroCluster switch configuration).
- This procedure cannot be used to convert a configuration with directly connected NS224 shelves or NS224 shelves attached to dedicated Ethernet switches to a configuration using shared Storage / MetroCluster switches.
- [Enable console logging](#) before performing this task.

Sending a custom AutoSupport message prior to maintenance

Before performing the maintenance, you should issue an AutoSupport message to notify NetApp technical support that maintenance is underway. Informing technical support that maintenance is underway prevents them from opening a case on the assumption that a disruption has occurred.

About this task

This task must be performed on each MetroCluster site.

Steps

1. To prevent automatic support case generation, send an Autosupport message to indicate the upgrade is underway.
 - a. Issue the following command:

```
system node autosupport invoke -node * -type all -message "Maint=10h Adding  
or Removing NS224 shelves" _
```

This example specifies a 10 hour maintenance window. You might want to allow additional time, depending on your plan.

If the maintenance is completed before the time has elapsed, you can invoke an AutoSupport message

indicating the end of the maintenance period:

```
system node autosupport invoke -node * -type all -message MAINT=end
```

- b. Repeat the command on the partner cluster.

Verifying the health of the MetroCluster configuration

You must verify the health and connectivity of the MetroCluster configuration prior to performing the transition.

Steps

1. Verify the operation of the MetroCluster configuration in ONTAP:

- a. Check whether the system is multipathed:

```
node run -node node-name sysconfig -a
```

- b. Check for any health alerts on both clusters:

```
system health alert show
```

- c. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

- d. Perform a MetroCluster check:

```
metrocluster check run
```

- e. Display the results of the MetroCluster check:

```
metrocluster check show
```

- f. Run Config Advisor.

[NetApp Downloads: Config Advisor](#)

- g. After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

2. Verify that the cluster is healthy:

```
cluster show -vserver Cluster
```

```
cluster_A::> cluster show -vserver Cluster
Node           Health Eligibility  Epsilon
-----
node_A_1      true   true           false
node_A_2      true   true           false

cluster_A::>
```

3. Verify that all cluster ports are up:

```
network port show -ipSPACE cluster
```

```
cluster_A::> network port show -ipSPACE cluster

Node: node_A_1-old

Port          IPspace      Broadcast Domain  Link  MTU  Speed(Mbps)  Health
-----
e0a           Cluster      Cluster           up    9000  auto/10000   healthy
e0b           Cluster      Cluster           up    9000  auto/10000   healthy

Node: node_A_2-old

Port          IPspace      Broadcast Domain  Link  MTU  Speed(Mbps)  Health
-----
e0a           Cluster      Cluster           up    9000  auto/10000   healthy
e0b           Cluster      Cluster           up    9000  auto/10000   healthy

4 entries were displayed.

cluster_A::>
```

4. Verify that all cluster LIFs are up and operational:

```
network interface show -vserver Cluster
```

Each cluster LIF should display true for Is Home and have a Status Admin/Oper of up/up

```
cluster_A::> network interface show -vserver cluster
```

Current Is	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Port
Home					
-----	-----	-----	-----	-----	-----

Cluster	node_A_1-old_clus1	up/up	169.254.209.69/16	node_A_1	e0a
true	node_A_1-old_clus2	up/up	169.254.49.125/16	node_A_1	e0b
true	node_A_2-old_clus1	up/up	169.254.47.194/16	node_A_2	e0a
true	node_A_2-old_clus2	up/up	169.254.19.183/16	node_A_2	e0b
true					

```
4 entries were displayed.
```

```
cluster_A::>
```

5. Verify that auto-revert is enabled on all cluster LIFs:

```
network interface show -vserver Cluster -fields auto-revert
```



```

cluster_A::> network interface show -vserver Cluster -fields auto-revert

          Logical
Vserver  Interface      Auto-revert
-----  -
Cluster
          node_A_1-old_clus1
                        true
          node_A_1-old_clus2
                        true
          node_A_2-old_clus1
                        true
          node_A_2-old_clus2
                        true

          4 entries were displayed.

cluster_A::>

```

Applying the new RCF file to the switches



If your switch is already correctly configured, you can skip these next sections and go directly to [Configuring MACsec encryption on Cisco 9336C switches](#), if applicable or to [Connecting the new NS224 shelf](#).

- You must change the switch configuration to add shelves.
- You should review the cabling details at [Platform port assignments](#).
- You must use the **RcfFileGenerator** tool to create the RCF file for your configuration. The [RcfFileGenerator](#) also provides a per-port cabling overview for each switch. Make sure that you choose the correct number of shelves. There are additional files created along with the RCF file that provide a detailed cabling layout matching your specific options. Use this cabling overview to verify your cabling when cabling the new shelves.

Upgrading RCF files on MetroCluster IP switches

If you are installing new switch firmware, you must install the switch firmware before upgrading the RCF file.

This procedure disrupts traffic on the switch where the RCF file is upgraded. Traffic will resume once the new RCF file is applied.

Steps

1. Verify the health of the configuration.
 - a. Verify that the MetroCluster components are healthy:

```
metrocluster check run
```

```
cluster_A::*> metrocluster check run
```

The operation runs in the background.

- a. After the `metrocluster check run` operation completes, run `metrocluster check show` to view the results.

After approximately five minutes, the following results are displayed:

```
-----
::*> metrocluster check show

Component          Result
-----
nodes              ok
lifs               ok
config-replication ok
aggregates        ok
clusters           ok
connections        not-applicable
volumes            ok
7 entries were displayed.
```

- b. To check the status of the running MetroCluster check operation, use the command:
metrocluster operation history show -job-id 38

- c. Verify that there are no health alerts:
system health alert show

2. Prepare the IP switches for the application of the new RCF files.

Resetting the Cisco IP switch to factory defaults

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

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

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.

```

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

```

The final set of prompts completes the configuration:

The following configuration will be applied:

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

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

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 is copied from SFTP server 10.10.99.99 to the local bootflash:

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

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

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

```

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	Image	Running-Version (pri:alt)	New-Version	Upg-Required
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. Repeat these steps on the remaining three IP switches in the MetroCluster IP configuration.

Configuring MACsec encryption on Cisco 9336C switches

If desired, 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.

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](#)

Enabling 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.

1. Enter the 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
```

Disabling Cisco MACsec Encryption

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



If you disable encryption, you must also delete your keys.

1. Enter the 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, policy and fallback-keychain configured on the interface to remove the MACsec configuration: `no macsec keychain keychain-name policy policy-name fallback-keychain keychain-name`

```
IP_switch_A_1(config-if)# no macsec keychain kc2 policy abc fallback-
keychain fb_kc2
```

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
```

Configuring a MACsec key chain and keys

For details on configuring a MACsec key chain, see the Cisco documentation for your switch.

Connecting the new NS224 shelf

Steps

1. Install the rail mount kit that came with your shelf by using the installation flyer that came in the kit box.
2. Install and secure the shelf onto the support brackets and rack or cabinet by using the installation flyer.
3. Connect the power cords to the shelf, secure them in with the power cord retainer, and then connect the power cords to different power sources for resiliency.

A shelf powers up when connected to a power source; it does not have power switches. When functioning correctly, a power supply's bicolored LED illuminates green.

4. Set the shelf ID to a number that is unique within the HA pair and across the configuration.
5. Connect the shelf ports in the following order:
 - a. Connect NSM-A, e0a to the switch (Switch-A1 or Switch-B1)
 - b. Connect NSM-B, e0a to the switch (Switch-A2 or Switch-B2)
 - c. Connect NSM-A, e0b to the switch (Switch-A1 or Switch-B1)
 - d. Connect NSM-B, e0b to the switch (Switch-A2 or Switch-B2)
6. Use the cabling layout generated from the **RcfFileGenerator** tool to cable the shelf to the appropriate ports.

Once the new shelf is cabled correctly, ONTAP automatically detects it on the network.

Configure end-to-end encryption in a MetroCluster IP configuration

Beginning with ONTAP 9.15.1, you can configure end-to-end encryption to encrypt back-end traffic, such as NVlog and storage replication data, between the sites in a MetroCluster IP configuration.

About this task

- You must be a cluster administrator to perform this task.
- Before you can configure end-to-end encryption, you must [Configure external key management](#).
- Review the supported systems and minimum ONTAP release required to configure end-to-end encryption in a MetroCluster IP configuration:

Minimum ONTAP release	Supported systems
ONTAP 9.15.1	<ul style="list-style-type: none">• AFF A400• FAS8300• FAS8700

Enable end-to-end encryption

Perform the following steps to enable end-to-end encryption.

Steps

1. Verify the health of the MetroCluster configuration.
 - a. Verify that the MetroCluster components are healthy:

```
metrocluster check run
```

```
cluster_A::*> metrocluster check run
```

The operation runs in the background.

- b. After the `metrocluster check run` operation completes, run:

```
metrocluster check show
```

After approximately five minutes, the following results are displayed:

```
cluster_A:::*> metrocluster check show
```

Component	Result
nodes	ok
lifs	ok
config-replication	ok
aggregates	ok
clusters	ok
connections	not-applicable
volumes	ok

7 entries were displayed.

- c. Check the status of the running MetroCluster check operation:

```
metrocluster operation history show -job-id <id>
```

- d. Verify that there are no health alerts:

```
system health alert show
```

2. Verify that external key management is configured on both clusters:

```
security key-manager external show-status
```

3. Enable end-to-end encryption for each DR group:

```
metrocluster modify -is-encryption-enabled true -dr-group-id  
<dr_group_id>
```

Example

```
cluster_A:::*> metrocluster modify -is-encryption-enabled true -dr-group  
-id 1  
Warning: Enabling encryption for a DR Group will secure NVLog and  
Storage  
          replication data sent between MetroCluster nodes and have an  
impact on  
          performance. Do you want to continue? {y|n}: y  
[Job 244] Job succeeded: Modify is successful.
```

Repeat this step for each DR group in the configuration.

4. Verify that end-to-end encryption is enabled:

```
metrocluster node show -fields is-encryption-enabled
```

Example

```
cluster_A::*> metrocluster node show -fields is-encryption-enabled

dr-group-id cluster      node      configuration-state is-encryption-
enabled
-----
1           cluster_A    node_A_1  configured         true
1           cluster_A    node_A_2  configured         true
1           cluster_B    node_B_1  configured         true
1           cluster_B    node_B_2  configured         true
4 entries were displayed.
```

Disable end-to-end encryption

Perform the following steps to disable end-to-end encryption.

Steps

1. Verify the health of the MetroCluster configuration.
 - a. Verify that the MetroCluster components are healthy:

```
metrocluster check run
```

```
cluster_A::*> metrocluster check run
```

The operation runs in the background.

- b. After the `metrocluster check run` operation completes, run:

```
metrocluster check show
```

After approximately five minutes, the following results are displayed:

```
cluster_A:::*> metrocluster check show
```

Component	Result
nodes	ok
lifs	ok
config-replication	ok
aggregates	ok
clusters	ok
connections	not-applicable
volumes	ok

7 entries were displayed.

- c. Check the status of the running MetroCluster check operation:

```
metrocluster operation history show -job-id <id>
```

- d. Verify that there are no health alerts:

```
system health alert show
```

2. Verify that external key management is configured on both clusters:

```
security key-manager external show-status
```

3. Disable end-to-end encryption on each DR group:

```
metrocluster modify -is-encryption-enabled false -dr-group-id  
<dr_group_id>
```

Example

```
cluster_A:::*> metrocluster modify -is-encryption-enabled false -dr-group  
-id 1  
[Job 244] Job succeeded: Modify is successful.
```

Repeat this step for each DR group in the configuration.

4. Verify that end-to-end encryption is disabled:


```
metrocluster node show -fields is-encryption-enabled
```

Example

```
cluster_A::*> metrocluster node show -fields is-encryption-enabled


dr-group-id cluster      node      configuration-state is-encryption-
enabled
-----
1           cluster_A    node_A_1  configured         false
1           cluster_A    node_A_2  configured         false
1           cluster_B    node_B_1  configured         false
1           cluster_B    node_B_2  configured         false
4 entries were displayed.
```

Power off and power on a single site in a MetroCluster IP configuration

If you need to perform site maintenance or relocate a single site in a MetroCluster IP configuration, you must know how to power off and power on the site.

If you need to relocate and reconfigure a site (for example, if you need to expand from a four-node to an eight-node cluster), you cannot complete these tasks at the same time. This procedure only covers the steps that are required to perform site maintenance or to relocate a site without changing its configuration.

The following diagram shows a MetroCluster configuration. Cluster_B is powered off for maintenance.



Power off a MetroCluster site

You must power off a site and all of the equipment before site maintenance or relocation can begin.

About this task

All the commands in the following steps are issued from the site that remains powered on.

Steps

1. Before you begin, check that any non-mirrored aggregates at the site are offline.
2. Verify the operation of the MetroCluster configuration in ONTAP:

- a. Check whether the system is multipathed:

```
node run -node node-name sysconfig -a
```

- b. Check for any health alerts on both clusters:

```
system health alert show
```

- c. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

- d. Perform a MetroCluster check:

```
metrocluster check run
```

- e. Display the results of the MetroCluster check:

```
metrocluster check show
```

- f. Check for any health alerts on the switches (if present):

```
storage switch show
```

- g. Run Config Advisor.

[NetApp Downloads: Config Advisor](#)

- h. After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

3. From the site you want to remain up, implement the switchover:

```
metrocluster switchover
```

```
cluster_A::*> metrocluster switchover
```

The operation can take several minutes to complete.

4. Monitor and verify the completion of the switchover:

```
metrocluster operation show
```

```
cluster_A::*> metrocluster operation show
  Operation: Switchover
  Start time: 10/4/2012 19:04:13
  State: in-progress
  End time: -
  Errors:

cluster_A::*> metrocluster operation show
  Operation: Switchover
  Start time: 10/4/2012 19:04:13
  State: successful
  End time: 10/4/2012 19:04:22
  Errors: -
```

5. If you have a MetroCluster IP configuration running ONTAP 9.6 or later, wait for the disaster site plexes to come online and the healing operations to automatically complete.

In MetroCluster IP configurations running ONTAP 9.5 or earlier, the disaster site nodes do not automatically boot to ONTAP and the plexes remain offline.

6. Move any volumes and LUNs that belong to unmirrored aggregates offline.
 - a. Move the volumes offline.

```
cluster_A::* volume offline <volume name>
```

- b. Move the LUNs offline.

```
cluster_A::* lun offline lun_path <lun_path>
```

7. Move unmirrored aggregates offline: `storage aggregate offline`

```
cluster_A*::> storage aggregate offline -aggregate <aggregate-name>
```

8. Depending on your configuration and ONTAP version, identify and move offline affected plexes that are located at the disaster site (Cluster_B).

You should move the following plexes offline:

- Non-mirrored plexes residing on disks located at the disaster site.

If you do not move the non-mirrored plexes at the disaster site offline, an outage might occur when the disaster site is later powered off.

- Mirrored plexes residing on disks located at the disaster site for aggregate mirroring. After they are moved offline, the plexes are inaccessible.

a. Identify the affected plexes.

Plexes that are owned by nodes at the surviving site consist of Pool1 disks. Plexes that are owned by nodes at the disaster site consist of Pool0 disks.

```

Cluster_A::> storage aggregate plex show -fields aggregate,status,is-
online,Plex,pool
aggregate      plex  status          is-online pool
-----
Node_B_1_aggr0 plex0 normal,active true      0
Node_B_1_aggr0 plex1 normal,active true      1

Node_B_2_aggr0 plex0 normal,active true      0
Node_B_2_aggr0 plex5 normal,active true      1

Node_B_1_aggr1 plex0 normal,active true      0
Node_B_1_aggr1 plex3 normal,active true      1

Node_B_2_aggr1 plex0 normal,active true      0
Node_B_2_aggr1 plex1 normal,active true      1

Node_A_1_aggr0 plex0 normal,active true      0
Node_A_1_aggr0 plex4 normal,active true      1

Node_A_1_aggr1 plex0 normal,active true      0
Node_A_1_aggr1 plex1 normal,active true      1

Node_A_2_aggr0 plex0 normal,active true      0
Node_A_2_aggr0 plex4 normal,active true      1

Node_A_2_aggr1 plex0 normal,active true      0
Node_A_2_aggr1 plex1 normal,active true      1
14 entries were displayed.

Cluster_A::>

```

The affected plexes are those that are remote to cluster A. The following table shows whether the disks are local or remote relative to cluster A:

Node	Disks in pool	Should the disks be set offline?	Example of plexes to be moved offline
Node_A_1 and Node_A_2	Disks in pool 0	No. Disks are local to cluster A.	-
	Disks in pool 1	Yes. Disks are remote to cluster A.	Node_A_1_aggr0/plex4 Node_A_1_aggr1/plex1 Node_A_2_aggr0/plex4 Node_A_2_aggr1/plex1

Node_B_1 and Node_B_2	Disks in pool 0	Yes. Disks are remote to cluster A.	Node_B_1_aggr1/plex0 Node_B_1_aggr0/plex0 Node_B_2_aggr0/plex0 Node_B_2_aggr1/plex0
	Disks in pool 1	No. Disks are local to cluster A.	-

b. Move the affected plexes offline:

```
storage aggregate plex offline
```

```
storage aggregate plex offline -aggregate Node_B_1_aggr0 -plex plex0
```



Perform this step for all plexes that have disks that are remote to Cluster_A.

9. Persistently offline the ISL switch ports according to the switch type.

10. Halt the nodes by running the following command on each node:

```
node halt -inhibit-takeover true -skip-lif-migration true -node <node-name>
```

11. Power off the equipment at the disaster site.

You must power off the following equipment in the order shown:

- Storage controllers - the storage controllers should currently be at the `LOADER` prompt, you must power them off completely.
- MetroCluster IP switches
- Storage shelves

Relocating the powered-off site of the MetroCluster

After the site is powered off, you can begin maintenance work. The procedure is the same whether the MetroCluster components are relocated within the same data center or relocated to a different data center.

- The hardware should be cabled in the same way as the previous site.
- If the Inter-Switch Link (ISL) speed, length, or number has changed, they all need to be reconfigured.

Steps

1. Verify that the cabling for all components is carefully recorded so that it can be correctly reconnected at the new location.
2. Physically relocate all the hardware, storage controllers, IP switches, FibreBridges, and storage shelves.
3. Configure the ISL ports and verify the intersite connectivity.
 - a. Power on the IP switches.



Do **not** power up any other equipment.

4. Use tools on the switches (as they are available) to verify the intersite connectivity.



You should only proceed if the links are correctly configured and stable.

5. Disable the links again if they are found to be stable.

Powering on the MetroCluster configuration and returning to normal operation

After maintenance has been completed or the site has been moved, you must power on the site and reestablish the MetroCluster configuration.

About this task

All the commands in the following steps are issued from the site that you power on.

Steps

1. Power on the switches.

You should power on the switches first. They might have been powered on during the previous step if the site was relocated.

- a. Reconfigure the Inter-Switch Link (ISL) if required or if this was not completed as part of the relocation.
 - b. Enable the ISL if fencing was completed.
 - c. Verify the ISL.
2. Power on the storage controllers and wait until you see the `LOADER` prompt. The controllers must not be fully booted.

If auto boot is enabled, press `Ctrl+C` to stop the controllers from automatically booting.

3. Power on the shelves, allowing enough time for them to power on completely.
4. Verify that the storage is visible.
 - a. Verify that the storage is visible from the surviving site. Bring the offline plexes back online to restart the resync operation and reestablish the SyncMirror.
 - b. Verify that the local storage is visible from the node in Maintenance mode:

```
disk show -v
```

5. Reestablish the MetroCluster configuration.

Follow the instructions in [Verifying that your system is ready for a switchback](#) to perform healing and switchback operations according to your MetroCluster configuration.

Powering off an entire MetroCluster IP configuration

You must power off the entire MetroCluster IP configuration and all of the equipment before maintenance or relocation can begin.



Beginning with ONTAP 9.8, the **storage switch** command is replaced with **system switch**. The following steps show the **storage switch** command, but if you are running ONTAP 9.8 or later, the **system switch** command is preferred.

1. Verify the MetroCluster configuration from both sites in the MetroCluster configuration.

a. Confirm that the MetroCluster configuration and operational mode are normal.

```
metrocluster show
```

b. Run the following command:

```
metrocluster interconnect show
```

c. Confirm connectivity to the disks by entering the following command on any one of the MetroCluster nodes:

```
run local sysconfig -v
```

d. Run the following command:

```
storage port show
```

e. Run the following command:

```
storage switch show
```

f. Run the following command:

```
network interface show
```

g. Run the following command:

```
network port show
```

h. Run the following command:

```
network device-discovery show
```

i. Perform a MetroCluster check:

```
metrocluster check run
```

j. Display the results of the MetroCluster check:

```
metrocluster check show
```

k. Run the following command:

```
metrocluster configuration-settings interface show
```

2. If necessary, disable AUSO by modifying the AUSO Failure Domain to

```
auso-disabled
```



In a MetroCluster IP configuration, the AUSO Failure Domain is already set to 'auso-disabled' unless the configuration is configured with ONTAP Mediator.

3. Verify the change using the command

```
metrocluster operation show
```



```
cluster_A_site_A::*> metrocluster operation show
  Operation: modify
    State: successful
  Start Time: 4/25/2020 20:20:36
  End Time: 4/25/2020 20:20:36
  Errors: -
```

4. Halt the nodes:

halt

```
system node halt -node node1_SiteA -inhibit-takeover true -ignore-quorum
-warnings true
```

5. Power off the following equipment at the site:

- Storage controllers
- MetroCluster IP switches
- Storage shelves

6. Wait for thirty minutes and then power on all storage shelves, MetroCluster IP switches, and storage controllers.

7. After the controllers are powered on, verify the MetroCluster configuration from both sites.

To verify the configuration, repeat step 1.

8. Perform power cycle checks.

- a. Verify that all sync-source SVMs are online:

vserver show

- b. Start any sync-source SVMs that are not online:

vserver start

Copyright information

Copyright © 2024 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.