



Maintain the MetroCluster components

ONTAP MetroCluster

NetApp
March 13, 2026

This PDF was generated from https://docs.netapp.com/us-en/ontap-metrocluster/maintain/concept_where_to_find_procedures_for_mcc_maintenance_tasks.html on March 13, 2026. Always check docs.netapp.com for the latest.

Table of Contents

- Maintain the MetroCluster components 1
 - Learn about MetroCluster maintenance 1
 - Prepare for maintenance tasks 1
 - Maintenance procedures for different types of MetroCluster configurations 1
 - All other maintenance procedures 1
 - Prepare for MetroCluster maintenance 2
 - Enable console logging before performing maintenance tasks 2
 - Remove ONTAP Mediator or Tiebreaker monitoring before performing maintenance tasks 3
 - MetroCluster failure and recovery scenarios 4
 - Using the Interoperability Matrix Tool to find MetroCluster information 5
- Maintenance procedures for MetroCluster FC configurations 5
 - Modify a switch or ATTO bridge IP address for health monitoring 6
 - FC-to-SAS bridge maintenance 7
 - FC switch maintenance and replacement 66
 - Replacing a shelf nondisruptively in a fabric-attached MetroCluster configuration 115
 - Hot add storage to a MetroCluster FC configuration 121
 - Hot-removing storage from a MetroCluster FC configuration 143
 - Power off and power on a single site in a MetroCluster FC configuration 146
 - Powering off an entire MetroCluster FC configuration 160
- Maintenance procedures for MetroCluster IP configurations 162
 - IP switch maintenance and replacement 162
 - Identifying storage in a MetroCluster IP configuration 188
 - Adding shelves to a MetroCluster IP using shared Storage MetroCluster switches 192
 - Configure end-to-end encryption in a MetroCluster IP configuration 208
 - Power off and power on a single site in a MetroCluster IP configuration 212
 - Power off an entire MetroCluster IP configuration 219
- Maintenance procedures for all MetroCluster configurations 220
 - Replacing a shelf nondisruptively in a stretch MetroCluster configuration 220
 - When to migrate root volumes to a new destination 223
 - Moving a metadata volume in MetroCluster configurations 223
 - Renaming a cluster in MetroCluster configurations 226
 - Verify the health of a MetroCluster configuration 228
 - Where to find additional information 230

Maintain the MetroCluster components

Learn about MetroCluster maintenance

Learn how to prepare for MetroCluster maintenance tasks and choose the correct maintenance procedure for your configuration.

Prepare for maintenance tasks

Review the information in [Prepare for MetroCluster maintenance](#) before performing any maintenance procedures.



You must enable console logging and remove ONTAP Mediator or Tiebreaker monitoring before you perform maintenance tasks.

Maintenance procedures for different types of MetroCluster configurations

- If you have a MetroCluster IP configuration, review the procedures in [Maintenance procedures for MetroCluster IP configurations](#).
- If you have a MetroCluster FC configuration, review the procedures in [Maintenance procedures for MetroCluster FC configurations](#).
- If you cannot find the procedure in the specific section for your configuration, review the procedures in [Maintenance procedures for all MetroCluster configurations](#).

All other maintenance procedures

The following table provides links to procedures related to MetroCluster maintenance that are not located in the three sections listed above:

Component	MetroCluster type (FC or IP)	Task	Procedure
ONTAP software	Both	ONTAP software upgrade	Upgrade, revert, or downgrade

Controller module	Both	FRU replacement (including replacement controller modules, PCIe cards, FC-VI card, and so on)	ONTAP Hardware Systems Documentation
		 <p>Moving a storage controller module or NVRAM card among the MetroCluster storage systems is not supported.</p>	
		Upgrade and expansion	MetroCluster Upgrade and Expansion
		Transition from FC to IP connectivity	Transition from MetroCluster FC to MetroCluster IP
Drive shelf	FC	All other shelf maintenance procedures. The standard procedures can be used.	Maintain DS460C DS224C and DS212C disk shelves
	IP	<p>All shelf maintenance procedures. The standard procedures can be used.</p> <p>If adding shelves for an unmirrored aggregate, see Considerations when using unmirrored aggregates</p>	Maintain DS460C DS224C and DS212C disk shelves

Prepare for MetroCluster maintenance

Enable console logging before performing maintenance tasks

Enable console logging on your devices before performing maintenance tasks.

NetApp strongly recommends that you enable console logging on the devices that you are using and take the following actions before performing maintenance procedures:

- Leave AutoSupport enabled during maintenance.

- Trigger a maintenance AutoSupport message before and after maintenance to disable case creation for the duration of the maintenance activity.

See the Knowledge Base article [How to suppress automatic case creation during scheduled maintenance windows](#).

- Enable session logging for any CLI session. For instructions on how to enable session logging, review the "Logging Session Output" section in the Knowledge Base article [How to configure PuTTY for optimal connectivity to ONTAP systems](#).

Remove ONTAP Mediator or Tiebreaker monitoring before performing maintenance tasks

Before performing maintenance tasks, you must remove monitoring if the MetroCluster configuration is monitored with the Tiebreaker or Mediator utility.

Maintenance tasks include upgrading the controller platform, upgrading ONTAP, and performing a negotiated switchover and switchback.

Steps

1. Collect the output for the following command:

```
storage iscsi-initiator show
```

2. Remove the existing MetroCluster configuration from Tiebreaker, Mediator, or other software that can initiate switchover.

If you are using...	Use this procedure...
Tiebreaker	Removing MetroCluster Configurations in the <i>MetroCluster Tiebreaker Installation and Configuration content</i>
Mediator	Issue the following command from the ONTAP prompt: <pre>metrocluster configuration-settings mediator remove</pre>
Third-party applications	Refer to the product documentation.

3. After completing maintenance of the MetroCluster configuration, you can resume monitoring with the Tiebreaker or Mediator utility.

If you are using...	Use this procedure
Tiebreaker	Adding MetroCluster configurations in the <i>MetroCluster Tiebreaker Installation and Configuration</i> section.

If you are using...	Use this procedure
Mediator	Configure ONTAP Mediator from a MetroCluster IP configuration in the <i>MetroCluster IP Installation and Configuration</i> section.
Third-party applications	Refer to the product documentation.

MetroCluster failure and recovery scenarios

You should be aware of how the MetroCluster configuration responds to different failure events.



For additional information about recovery from node failures, see the section "Choosing the correct recovery procedure" in the [Recover from a disaster](#).

Event	Impact	Recovery
Single node failure	A failover is triggered.	The configuration recovers through a local takeover. RAID is not impacted. Review system messages and replace failed FRUs as necessary. ONTAP Hardware Systems Documentation
Two nodes fail at one site	Two nodes will fail only if automated switchover is enabled in the MetroCluster Tiebreaker software.	Manual unplanned switchover (USO) if automated switchover in MetroCluster Tiebreaker software is not enabled. ONTAP Hardware Systems Documentation
MetroCluster IP interface—failure of one port	The system is degraded. Additional port failure impacts HA mirroring.	The second port is used. Health Monitor generates an alert if the physical link to the port is broken. Review system messages and replace failed FRUs as necessary. ONTAP Hardware Systems Documentation

MetroCluster IP interface—failure of both ports	HA capability is impacted. RAID SyncMirror of the node stops syncing.	Immediate manual recovery is required as there is no HA takeover. Review system messages and replace failed FRUs as necessary. ONTAP Hardware Systems Documentation
Failure of one MetroCluster IP switch	No impact. Redundancy is provided through the second network.	Replace the failed switch as necessary. Replacing an IP switch
Failure of two MetroCluster IP switches that are in the same network	No impact. Redundancy is provided through the second network.	Replace the failed switch as necessary. Replacing an IP switch
Failure of two MetroCluster IP switches that are at one site	RAID SyncMirror of the node stops syncing. HA capability is impacted and the cluster goes out of quorum.	Replace the failed switch as necessary. Replacing an IP switch
Failure of two MetroCluster IP switches that are at different sites and not on the same network (diagonal failure)	RAID SyncMirror of the node stops syncing.	RAID SyncMirror of the node stops syncing. Cluster and HA capability are not impacted. Replace the failed switch as necessary. Replacing an IP switch

Using the Interoperability Matrix Tool to find MetroCluster information

When setting up the MetroCluster configuration, you can use the Interoperability Tool to ensure you are using supported software and hardware versions.

[NetApp Interoperability Matrix Tool](#)

After opening the Interoperability Matrix, you can use the Storage Solution field to select your MetroCluster solution.

You use the **Component Explorer** to select the components and ONTAP version to refine your search.

You can click **Show Results** to display the list of supported configurations that match the criteria.

Maintenance procedures for MetroCluster FC configurations

Modify a switch or ATTO bridge IP address for health monitoring

After modifying the IP addresses of MetroCluster FC back-end switches and ATTO bridges, you must replace the old health monitoring IP addresses with the new values.

- [Modify a switch IP address](#)
- [Modify an ATTO bridge IP address](#)

Modify a switch IP address

Replace the old health monitoring IP address of a MetroCluster FC back-end switch.

Before you begin

Refer to the switch vendor's documentation for your switch model to change the IP address on the switch before changing the health monitoring IP address.

Steps

1. Run the `::> storage switch show` command and in the output, note the switches that are reporting errors.
2. Remove the switch entries with old IP addresses:

```
::> storage switch remove -name switch_name
```

3. Add the switches with new IP addresses:

```
::> storage switch add -name switch_name -address new_IP_address -managed-by  
in-band
```

4. Verify the new IP addresses and confirm that there are no errors:

```
::> storage switch show
```

5. If required, refresh the entries:

```
::> set advanced
```

```
::*> storage switch refresh
```

```
::*> set admin
```

Modify an ATTO bridge IP address

Replace the old health monitoring IP address of an ATTO bridge.

Steps

1. Run the `::> storage bridge show` command and in the output, note the ATTO bridges that are reporting errors.
2. Remove the ATTO bridge entries with old IP addresses:

```
::> storage bridge remove -name ATTO_bridge_name
```

3. Add the ATTO bridges with new IP addresses:

```
::> storage bridge add -name ATTO_bridge_name -address new_IP_address -managed  
-by in-band
```

4. Verify the new IP addresses and confirm that there are no errors:

```
::> storage bridge show
```

5. If required, refresh the entries:

```
::> set advanced
```

```
::*> storage bridge refresh
```

```
::*> set admin
```

FC-to-SAS bridge maintenance

Support for FibreBridge 7600N bridges in MetroCluster configurations

The FibreBridge 7600N bridge is supported on ONTAP 9.5 and later as a replacement for the FibreBridge 7500N or 6500N bridge or when adding new storage to the MetroCluster configuration. The zoning requirements and restrictions regarding use of the bridge's FC ports are the same as that of the FibreBridge 7500N bridge.

[NetApp Interoperability Matrix Tool](#)



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

Use case	Zoning changes needed?	Restrictions	Procedure
Replacing a single FibreBridge 7500N bridge with a single FibreBridge 7600N bridge	No	The FibreBridge 7600N bridge must be configured exactly the same as the FibreBridge 7500N bridge.	Hot-swapping a FibreBridge 7500N with a 7600N bridge
Replacing a single FibreBridge 6500N bridge with a single FibreBridge 7600N bridge	No	The FibreBridge 7600N bridge must be configured exactly the same as the FibreBridge 6500N bridge.	Hot-swapping a FibreBridge 6500N bridge with a FibreBridge 7600N or 7500N bridge

Adding new storage through adding a new pair of FibreBridge 7600N bridges	<p>Yes</p> <p>You must add storage zones for each of the FC ports of the new bridges.</p>	<p>You must have available ports on the FC switch fabric (in a fabric-attached MetroCluster configuration) or on the storage controllers (in a stretch MetroCluster configuration). Each pair of FibreBridge 7500N or 7600N bridges can support up to four stacks.</p>	<p>Hot-adding a stack of SAS disk shelves and bridges to a MetroCluster system</p>
---	---	--	--

Support for FibreBridge 7500N bridges in MetroCluster configurations

The FibreBridge 7500N bridge is supported as a replacement for the FibreBridge 6500N bridge or for when adding new storage to the MetroCluster configuration. The supported configurations have zoning requirements and restrictions regarding use of the bridge's FC ports and stack and storage shelf limits.



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

Use case	Zoning changes needed?	Restrictions	Procedure
Replacing a single FibreBridge 6500N bridge with a single FibreBridge 7500N bridge	No	<p>The FibreBridge 7500N bridge must be configured exactly the same as the FibreBridge 6500N bridge, using a single FC port and attaching to a single stack. The second FC port on the FibreBridge 7500N must not be used.</p>	<p>Hot-swapping a FibreBridge 6500N bridge with a FibreBridge 7600N or 7500N bridge</p>
Consolidating multiple stacks by replacing multiple pairs of FibreBridge 6500N bridges with a single pair of FibreBridge 7500N bridges	Yes	<p>In this case, you take the FibreBridge 6500N bridges out of service and replace them with a single pair of FibreBridge 7500N bridges. Each pair of FibreBridge 7500N or 7600N bridges can support up to four stacks.</p> <p>At the end of the procedure, both the top and bottom of the stacks must be connected to corresponding ports on the FibreBridge 7500N bridges.</p>	<p>Replacing a pair of FibreBridge 6500N bridges with 7600N or 7500N bridges</p>

Use case	Zoning changes needed?	Restrictions	Procedure
Adding new storage through adding a new pair of FibreBridge 7500N bridges	Yes You must add storage zones for each of the FC ports of the new bridges.	You must have available ports on the FC switch fabric (in a fabric-attached MetroCluster configuration) or on the storage controllers (in a stretch MetroCluster configuration). Each pair of FibreBridge 7500N or 7600N bridges can support up to four stacks.	Hot-adding a stack of SAS disk shelves and bridges to a MetroCluster system

Enabling IP port access on the FibreBridge 7600N bridge if necessary

If you are using an ONTAP version prior to 9.5, or otherwise plan to use out-of-band access to the FibreBridge 7600N bridge using telnet or other IP port protocols and services (FTP, ExpressNAV, ICMP, or QuickNAV), you can enable the access services via the console port.

Unlike the ATTO FibreBridge 7500N bridge, the FibreBridge 7600N bridge is shipped with all IP port protocols and services disabled.

Beginning with ONTAP 9.5, *in-band management* of the bridges is supported. This means the bridges can be configured and monitored from the ONTAP CLI via the FC connection to the bridge. Physical access to the bridge via the bridge Ethernet ports is not required and the bridge user interfaces are not required.

Beginning with ONTAP 9.8, *in-band management* of the bridges is supported by default and out-of-band SNMP management is deprecated.

This task is required if you are **not** using in-band management to manage the bridges. In this case, you need to configure the bridge via the Ethernet management port.

Steps

1. Access the bridge's console interface by connecting a serial cable to the serial port on the FibreBridge 7600N bridge.
2. Using the console, enable the access services, and then save the configuration:

```
set closeport none
```

```
saveconfiguration
```

The `set closeport none` command enables all access services on the bridge.

3. Disable a service, if desired, by issuing the `set closeport` and repeating the command as necessary until all desired services are disabled:

```
set closeport service
```

The `set closeport` command disables a single service at a time.

service can specify one of the following:

- expressnav
- ftp
- icmp
- quicknav
- snmp
- telnet

You can check whether a specific protocol is enabled or disabled by using the `get closeport` command.

4. If you are enabling SNMP, you must also issue the `set SNMP enabled` command:

```
set SNMP enabled
```

SNMP is the only protocol that requires a separate enable command.

5. Save the configuration:

```
saveconfiguration
```

Updating firmware on a FibreBridge bridge

The procedure for updating the bridge firmware depends on your bridge model and ONTAP version.

About this task

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

Updating firmware on FibreBridge 7600N or 7500N bridges on configurations running ONTAP 9.4 and later

You might need to update the firmware on your FibreBridge bridges to ensure that you have the latest features or to resolve possible issues. This procedure should be used for FibreBridge 7600N or 7500N bridges on configurations running ONTAP 9.4 and later.

- The MetroCluster configuration must be operating normally.
- All of the FibreBridge bridges in the MetroCluster configuration must be up and operating.
- All of the storage paths must be available.
- You need the admin password and access to an HTTP, FTP, or Trivial File Transfer Protocol (TFTP) server.
- You must be using a supported firmware version.

[NetApp Interoperability Matrix Tool](#)

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

- You can use this task only on FibreBridge 7600N or 7500N bridges in configurations running ONTAP 9.4 or later.

- You must perform this task on each FibreBridge bridge in the MetroCluster configuration, so that all of the bridges are running the same firmware version.



This procedure is nondisruptive and takes approximately 30 minutes to complete.



Beginning with ONTAP 9.8, the `system bridge` command replaces the `storage bridge`. The following steps show the `system bridge` command, but if you're running a version earlier than ONTAP 9.8, you should use the `storage bridge` command.

Steps

1. Invoke an AutoSupport message indicating the start of the maintenance:

```
system node autosupport invoke -node * -type all -message MAINT=maintenance-  
window-in-hours
```

“maintenance-window-in-hours” specifies the length of the maintenance window, with a maximum of 72 hours. 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
```

2. Go to the ATTO FibreBridge page and select the appropriate firmware for your bridge.

[ATTO FibreBridge Firmware Download Page](#)

3. Review the Caution/MustRead and End User Agreement, and click the check box to indicate acceptance and proceed.
4. Place the firmware file in a network location that is network accessible to the controller modules.

You can enter the commands in the remaining steps from the console of either controller module.

5. Change to the advanced privilege level:

```
set -privilege advanced
```

You must respond with “y” when prompted to continue into advanced mode and see the advanced mode prompt (*>).

6. Update the bridge firmware.

Beginning in ONTAP 9.16.1, you can use credentials to update the bridge firmware if they are required by the server to download the firmware package.

If credentials are not required:

- a. Update the bridge firmware:

```
system bridge firmware update -bridge <name> -uri <URL-of-firmware-  
package>
```

Example

```
cluster_A> system bridge firmware update -bridge bridge_A_1a -uri  
http://192.168.132.97/firmware.ZBD
```

If credentials are required:

- a. Update the bridge firmware and specify the required user name:

```
system bridge firmware update -bridge <name> -uri <URL-of-  
firmware-package> -username <name>
```

- b. Enter the password when prompted in the output, as shown in the following example:

Example

```
cluster_A> system bridge firmware update -bridge bridge_A_1a -uri  
http://192.168.132.97/firmware.ZBD -username abc  
  
(system bridge)  
  
Enter the password:  
  
[Job 70] Job is queued: System bridge firmware update job.
```

7. Return to the admin privilege level:

```
set -privilege admin
```

8. Verify that the firmware upgrade is complete:

```
job show -name "<job_name>"
```

The following example shows that the job “system bridge firmware update” is still running:

```
cluster_A> job show -name "system bridge firmware update"
Owning
```

Job ID	Name	Vserver	Node	State
2246	job-name	cluster_A	node_A_1	Running

Description: System bridge firmware update job

After approximately 10 minutes, the new firmware is fully installed and the job state will be Success:

```
cluster_A> job show -name "system bridge firmware update"
```

Job ID	Name	Vserver	Node	State
2246	System bridge firmware update	cluster_A	node_A_1	Success

Description: System bridge firmware update job

9. Complete the steps according to whether in-band management is enabled and which version of ONTAP your system is running:

- If you are running ONTAP 9.4, in-band management is not supported and the command must be issued from the bridge console:
 - i. Run the `flashimages` command on the console of the bridge and confirm that the correct firmware versions are displayed.



The example shows that primary flash image shows the new firmware image, while the secondary flash image shows the old image.

```
flashimages

;Type Version
;=====
Primary 3.16 001H
Secondary 3.15 002S
Ready.
```

- i. Reboot the bridge by running the `firmwarerestart` command from the bridge.

- If you are running ONTAP 9.5 or later, in-band management is supported and the command can be issued from the cluster prompt:
- ii. Run the system bridge `run-cli -name <bridge_name> -command FlashImages` command.



The example shows that primary flash image shows the new firmware image, while the secondary flash image shows the old image.

```
cluster_A> system bridge run-cli -name ATTO_7500N_IB_1 -command
FlashImages

[Job 2257]

;Type          Version
;=====
Primary 3.16 001H
Secondary 3.15 002S
Ready.

[Job 2257] Job succeeded.
```

- iii. If necessary, restart the bridge:

```
system bridge run-cli -name ATTO_7500N_IB_1 -command FirmwareRestart
```



Beginning with ATTO firmware version 2.95 the bridge will restart automatically and this step is not required.

- 10. Verify that the bridge restarted correctly:

```
sysconfig
```

The system should be cabled for multipath high availability (both controllers have access through the bridges to the disk shelves in each stack).

```
cluster_A> node run -node cluster_A-01 -command sysconfig
NetApp Release 9.6P8: Sat May 23 16:20:55 EDT 2020
System ID: 1234567890 (cluster_A-01); partner ID: 0123456789 (cluster_A-
02)
System Serial Number: 200012345678 (cluster_A-01)
System Rev: A4
System Storage Configuration: Quad-Path HA
```

- 11. Verify that the FibreBridge firmware was updated:

```
system bridge show -fields fw-version,symbolic-name
```

```
cluster_A> system bridge show -fields fw-version,symbolic-name
name fw-version symbolic-name
-----
ATTO_20000010affeaffe 3.10 A06X bridge_A_1a
ATTO_20000010affeaffae 3.10 A06X bridge_A_1b
ATTO_20000010affeaffff 3.10 A06X bridge_A_2a
ATTO_20000010affeafffa 3.10 A06X bridge_A_2b
4 entries were displayed.
```

12. Verify the partitions are updated from the bridge's prompt:

```
flashimages
```

The primary flash image displays the new firmware image, while the secondary flash image displays the old image.

```
Ready.
flashimages

;Type          Version
;=====
  Primary      3.16 001H
  Secondary     3.15 002S

Ready.
```

13. Repeat steps 5 to 10 to ensure that both flash images are updated to the same version.

14. Verify that both flash images are updated to the same version.

```
flashimages
```

The output should show the same version for both partitions.

```
Ready.
flashimages

;Type          Version
;=====
  Primary      3.16 001H
  Secondary     3.16 001H

Ready.
```

- Repeat steps 5 to 13 on the next bridge until all of the bridges in the MetroCluster configuration have been updated.

Replacing a single FC-to-SAS bridge

You can nondisruptively replace a bridge with a same model bridge or with a new model bridge.

Before you begin

You need the admin password and access to an FTP or SCP server.

About this task

This procedure is nondisruptive and takes approximately 60 minutes to complete.

This procedure uses the bridge CLI to configure and manage a bridge, and to update the bridge firmware and the ATTO QuickNAV utility to configure the bridge Ethernet management 1 port. You can use other interfaces if they meet the requirements.

[Requirements for using other interfaces to configure and manage FibreBridge bridges](#)

Related information

[Replacing a pair of FibreBridge 6500N bridges with 7600N or 7500N bridges](#)

Verifying storage connectivity

Before replacing bridges, you should verify bridge and storage connectivity. Familiarizing yourself with the command output enables you to subsequently confirm connectivity after making configuration changes.

About this task

You can issue these commands from the admin prompt of any of the controller modules in the MetroCluster configuration at the site undergoing maintenance.

Steps

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

```
run local sysconfig -v
```

The output shows the disks attached to the initiator ports on the controller, and identifies the shelves connected to the FC-to-SAS bridges:

```
node_A_1> run local sysconfig -v
NetApp Release 9.3.2X18: Sun Dec 13 01:23:24 PST 2017
System ID: 4068741258 (node_A_1); partner ID: 4068741260 (node_B_1)
System Serial Number: 940001025471 (node_A_1)
System Rev: 70
System Storage Configuration: Multi-Path HA**<=== Configuration should
be multi-path HA**
.
.
.
```

```

slot 0: FC Host Adapter 0g (QLogic 8324 rev. 2, N-port, <UP>)**<===
Initiator port**
    Firmware rev:      7.5.0
    Flash rev:         0.0.0
    Host Port Id:      0x60130
    FC Node Name:      5:00a:098201:bae312
    FC Port Name:      5:00a:098201:bae312
    SFP Vendor:        UTILITIES CORP.
    SFP Part Number:   FTLF8529P3BCVAN1
    SFP Serial Number: URQ0Q9R
    SFP Capabilities:  4, 8 or 16 Gbit
    Link Data Rate:    16 Gbit
    Switch Port:       brcd6505-fcs40:1
**<List of disks visible to port\>**
    ID      Vendor  Model          FW      Size
brcd6505-fcs29:12.126L1527  : NETAPP  X302_HJUPI01TSSM NA04
847.5GB (1953525168 512B/sect)
brcd6505-fcs29:12.126L1528  : NETAPP  X302_HJUPI01TSSA NA02
847.5GB (1953525168 512B/sect)
.
.
.
**<List of FC-to-SAS bridges visible to port\>**
FC-to-SAS Bridge:
brcd6505-fcs40:12.126L0      : ATTO    FibreBridge6500N 1.61
FB6500N102980
brcd6505-fcs42:13.126L0     : ATTO    FibreBridge6500N 1.61
FB6500N102980
brcd6505-fcs42:6.126L0      : ATTO    FibreBridge6500N 1.61
FB6500N101167
brcd6505-fcs42:7.126L0      : ATTO    FibreBridge6500N 1.61
FB6500N102974
.
.
.
**<List of storage shelves visible to port\>**
brcd6505-fcs40:12.shelf6: DS4243  Firmware rev. IOM3 A: 0200
IOM3 B: 0200
brcd6505-fcs40:12.shelf8: DS4243  Firmware rev. IOM3 A: 0200
IOM3 B: 0200
.
.
.

```

Hot-swapping a bridge with a replacement bridge of the same model

You can hot-swap a failed bridge with another bridge of the same model.

About this task

If you will be using in-band management of the bridge rather than IP management, the steps for configuring the Ethernet port and IP settings can be skipped, as noted in the relevant steps.



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

Steps

1. If the old bridge is accessible, you can retrieve the configuration information.

If...	Then...
You are using IP management	Connect to the old bridge with a Telnet connection and copy the output of the bridge configuration.
You are using in-band management	Use the ONTAP CLI to retrieve the configuration information with the following commands: <pre>storage bridge run-cli -name <i>bridge-name</i> -command "info"</pre> <pre>storage bridge run-cli -name <i>bridge-name</i> -command "sasportlist"</pre>

- a. Enter the command:

```
storage bridge run-cli -name bridge_A1 -command "info"
```

```
info

Device Status           = Good
Unsaved Changes         = None
Device                  = "FibreBridge 7500N"
Serial Number           = FB7500N100000
Device Version          = 3.10
Board Revision          = 7
Build Number            = 007A
Build Type               = Release
Build Date              = "Aug 20 2019" 11:01:24
Flash Revision          = 0.02
Firmware Version        = 3.10
BCE Version (FPGA 1)    = 15
BAU Version (FPGA 2)    = 33
```

```
User-defined name      = "bridgeA1"
World Wide Name       = 20 00 00 10 86 A1 C7 00
MB of RAM Installed   = 512
FC1 Node Name        = 20 00 00 10 86 A1 C7 00
FC1 Port Name        = 21 00 00 10 86 A1 C7 00
FC1 Data Rate        = 16Gb
FC1 Connection Mode   = ptp
FC1 FW Revision      = 11.4.337.0
FC2 Node Name        = 20 00 00 10 86 A1 C7 00
FC2 Port Name        = 22 00 00 10 86 A1 C7 00
FC2 Data Rate        = 16Gb
FC2 Connection Mode   = ptp
FC2 FW Revision      = 11.4.337.0
SAS FW Revision      = 3.09.52
MP1 IP Address        = 10.10.10.10
MP1 IP Subnet Mask    = 255.255.255.0
MP1 IP Gateway        = 10.10.10.1
MP1 IP DHCP           = disabled
MP1 MAC Address       = 00-10-86-A1-C7-00
MP2 IP Address        = 0.0.0.0 (disabled)
MP2 IP Subnet Mask    = 0.0.0.0
MP2 IP Gateway        = 0.0.0.0
MP2 IP DHCP           = enabled
MP2 MAC Address       = 00-10-86-A1-C7-01
SNMP                  = enabled
SNMP Community String = public
PS A Status           = Up
PS B Status           = Up
Active Configuration  = NetApp
```

Ready.

b. Enter the command:

```
storage bridge run-cli -name bridge_A1 -command "sasportlist"
```

SASPortList

	;Connector	PHY	Link	Speed	SAS Address
	;=====				
Device	A	1	Up	6Gb	5001086000a1c700
Device	A	2	Up	6Gb	5001086000a1c700
Device	A	3	Up	6Gb	5001086000a1c700
Device	A	4	Up	6Gb	5001086000a1c700
Device	B	1	Disabled	12Gb	5001086000a1c704
Device	B	2	Disabled	12Gb	5001086000a1c704
Device	B	3	Disabled	12Gb	5001086000a1c704
Device	B	4	Disabled	12Gb	5001086000a1c704
Device	C	1	Disabled	12Gb	5001086000a1c708
Device	C	2	Disabled	12Gb	5001086000a1c708
Device	C	3	Disabled	12Gb	5001086000a1c708
Device	C	4	Disabled	12Gb	5001086000a1c708
Device	D	1	Disabled	12Gb	5001086000a1c70c
Device	D	2	Disabled	12Gb	5001086000a1c70c
Device	D	3	Disabled	12Gb	5001086000a1c70c
Device	D	4	Disabled	12Gb	5001086000a1c70c

2. If the bridge is in a fabric-attached MetroCluster configuration, disable all of the switch ports that connect to the bridge FC port or ports.
3. From the ONTAP cluster prompt, remove the bridge undergoing maintenance from health monitoring:
 - a. Remove the bridge:
`storage bridge remove -name bridge-name`
 - b. View the list of monitored bridges and confirm that the removed bridge is not present:
`storage bridge show`
4. Properly ground yourself.
5. Power down the ATTO bridge and remove the power cables connected to the bridge.
6. Disconnect the cables that are connected to the old bridge.

You should make note of the port to which each cable was connected.

7. Remove the old bridge from the rack.
8. Install the new bridge into the rack.
9. Reconnect the power cord and, if configuring for IP access to the bridge, a shielded Ethernet cable.



You must not reconnect the SAS or FC cables at this time.

10. Connect the bridge to a power source, and then turn it on.

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

11. If configuring for in-band management, connect a cable from FibreBridge RS-232 serial port to the serial (COM) port on a personal computer.

The serial connection will be used for initial configuration, and then in-band management via ONTAP and the FC ports can be used to monitor and manage the bridge.

12. If configuring for IP management, configure the Ethernet management 1 port for each bridge by following the procedure in section 2.0 of the *ATTO FibreBridge Installation and Operation Manual* for your bridge model.

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

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

13. Configure the bridge.

If you retrieved the configuration information from the old bridge, use the information to configure the new bridge.

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

The *ATTO FibreBridge Installation and Operation Manual* for your bridge model has the most current information on available commands and how to use them.



Do not configure time synchronization on ATTO FibreBridge 7600N or 7500N. The time synchronization for ATTO FibreBridge 7600N or 7500N is set to the cluster time after the bridge is discovered by ONTAP. It is also synchronized periodically once a day. The time zone used is GMT and is not changeable.

- a. If configuring for IP management, configure the IP settings of the bridge.

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

If using the CLI, you must run the following commands:

```
set ipaddress mp1 _ip-address
```

```
set ipsubnetmask mp1 subnet-mask
```

```
set ipgateway mp1 x.x.x.x
```

```
set ipdhcp mp1 disabled
```

```
set ethernetspeed mp1 1000
```

- b. Configure the bridge name.

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

Example bridge names for one stack group on each site:

- bridge_A_1a
- bridge_A_1b
- bridge_B_1a
- bridge_B_1b

If using the CLI, you must run the following command:

```
set bridgename bridgename
```

c. If running ONTAP 9.4 or earlier, enable SNMP on the bridge:

```
set SNMP enabled
```

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

14. Configure the bridge FC ports.

a. Configure the data rate/speed of the bridge FC ports.

The supported FC data rate depends on your model bridge.

- The FibreBridge 7600N bridge supports up to 32, 16, or 8 Gbps.
- The FibreBridge 7500N bridge supports up to 16, 8, or 4 Gbps.



The FCDataRate speed you select is limited to the maximum speed supported by both the bridge and the switch to which the bridge port connects. Cabling distances must not exceed the limitations of the SFPs and other hardware.

If using the CLI, you must run the following command:

```
set FCDataRate port-number port-speed
```

b. If you are configuring a FibreBridge 7500N, configure the connection mode that the port uses to "ptp".



The FCConnMode setting is not required when configuring a FibreBridge 7600N bridge.

If using the CLI, you must run the following command:

```
set FCConnMode port-number ptp
```

c. If you are configuring a FibreBridge 7600N or 7500N bridge, you must configure or disable the FC2 port.

- If you are using the second port, you must repeat the previous substeps for the FC2 port.
- If you are not using the second port, then you must disable the port:

```
FCPortDisable port-number
```

- d. If you are configuring a FibreBridge 7600N or 7500N bridge, disable the unused SAS ports:

```
SASPortDisable sas-port
```



SAS ports A through D are enabled by default. You must disable the SAS ports that are not being used. If only SAS port A is used, then SAS ports B, C, and D must be disabled.

15. Secure access to the bridge and save the bridge's configuration.

- a. From the controller prompt check the status of the bridges: `storage bridge show`

The output shows which bridge is not secured.

- b. Check the status of the unsecured bridge's ports:

```
info
```

The output shows the status of Ethernet ports MP1 and MP2.

- c. If Ethernet port MP1 is enabled, run the following command:

```
set EthernetPort mp1 disabled
```



If Ethernet port MP2 is also enabled, repeat the previous substep for port MP2.

- d. Save the bridge's configuration.

You must run the following commands:

```
SaveConfiguration
```

```
FirmwareRestart
```

You are prompted to restart the bridge.

16. Connect the FC cables to the same ports on the new bridge.

17. Update the FibreBridge firmware on each bridge.

If the new bridge is the same type as the partner bridge, upgrade to the same firmware as the partner bridge. If the new bridge is a different type to the partner bridge, upgrade to the latest firmware supported by the bridge and version of ONTAP. See [Updating firmware on a FibreBridge bridge](#)

18. Reconnect the SAS cables to the same ports on the new bridge.

You must replace the cables connecting the bridge to the top or bottom of the shelf stack. The FibreBridge 7600N and 7500N bridges require mini-SAS cables for these connections.



Wait at least 10 seconds before connecting the port. The SAS cable connectors are keyed; when oriented correctly into a SAS port, the connector clicks into place and the disk shelf SAS port LNK LED illuminates green. For disk shelves, you insert a SAS cable connector with the pull tab oriented down (on the underside of the connector). For controllers, the orientation of SAS ports can vary depending on the platform model; therefore, the correct orientation of the SAS cable connector varies.

19. Verify that each bridge can see all of the disk drives and disk shelves to which the bridge is connected.

If you are using the...	Then...
ATTO ExpressNAV GUI	<p>a. In a supported web browser, enter the IP address of the bridge in the browser box.</p> <p>You are brought to the ATTO FibreBridge homepage, which has a link.</p> <p>b. Click the link, and then enter your user name and the password that you designated when you configured the bridge.</p> <p>The ATTO FibreBridge status page appears with a menu to the left.</p> <p>c. Click Advanced in the menu.</p> <p>d. View the connected devices:</p> <pre>sastargets</pre> <p>e. Click Submit.</p>
Serial port connection	<p>View the connected devices:</p> <pre>sastargets</pre>

The output shows the devices (disks and disk shelves) to which the bridge is connected. The output lines are sequentially numbered so that you can quickly count the devices.



If the text response truncated appears at the beginning of the output, you can use Telnet to connect to the bridge, and then view all of the output by using the `sastargets` command.

The following output shows that 10 disks are connected:

```

Tgt VendorID ProductID      Type SerialNumber
  0 NETAPP    X410_S15K6288A15 DISK 3QP1CLE300009940UHJV
  1 NETAPP    X410_S15K6288A15 DISK 3QP1ELF600009940V1BV
  2 NETAPP    X410_S15K6288A15 DISK 3QP1G3EW00009940U2M0
  3 NETAPP    X410_S15K6288A15 DISK 3QP1EWMP00009940U1X5
  4 NETAPP    X410_S15K6288A15 DISK 3QP1FZLE00009940G8YU
  5 NETAPP    X410_S15K6288A15 DISK 3QP1FZLF00009940TZKZ
  6 NETAPP    X410_S15K6288A15 DISK 3QP1CEB400009939MGXL
  7 NETAPP    X410_S15K6288A15 DISK 3QP1G7A900009939FNNT
  8 NETAPP    X410_S15K6288A15 DISK 3QP1FY0T00009940G8PA
  9 NETAPP    X410_S15K6288A15 DISK 3QP1FXW600009940VERQ

```

20. Verify that the command output shows that the bridge is connected to all of the appropriate disks and disk shelves in the stack.

If the output is...	Then...
Correct	Repeat Step 19 for each remaining bridge.
Not correct	<ul style="list-style-type: none"> a. Check for loose SAS cables or correct the SAS cabling by repeating Step 18. b. Repeat Step 19.

21. If the bridge is in a fabric-attached MetroCluster configuration, re-enable the FC switch port that you disabled at the beginning of this procedure.

This should be the port that connects to the bridge.

22. From the system console of both controller modules, verify that all of the controller modules have access through the new bridge to the disk shelves (that is, that the system is cabled for Multipath HA):

```
run local sysconfig
```



It might take up to a minute for the system to complete discovery.

If the output does not indicate Multipath HA, you must correct the SAS and FC cabling because not all of the disk drives are accessible through the new bridge.

The following output states that the system is cabled for Multipath HA:

```
NetApp Release 8.3.2: Tue Jan 26 01:41:49 PDT 2016
System ID: 1231231231 (node_A_1); partner ID: 4564564564 (node_A_2)
System Serial Number: 700000123123 (node_A_1); partner Serial Number:
700000456456 (node_A_2)
System Rev: B0
System Storage Configuration: Multi-Path HA
System ACP Connectivity: NA
```



When the system is not cabled as Multipath HA, restarting a bridge might cause loss of access to the disk drives and result in a multi-disk panic.

23. If running ONTAP 9.4 or earlier, verify that the bridge is configured for SNMP.

If you are using the bridge CLI, run the following command:

```
get snmp
```

24. From the ONTAP cluster prompt, add the bridge to health monitoring:

- a. Add the bridge, using the command for your version of ONTAP:

ONTAP version	Command
9.5 and later	<code>storage bridge add -address 0.0.0.0 -managed-by in-band -name <i>bridge-name</i></code>
9.4 and earlier	<code>storage bridge add -address <i>bridge-ip-address</i> -name <i>bridge-name</i></code>

- b. Verify that the bridge has been added and is properly configured:

```
storage bridge show
```

It might take as long as 15 minutes to reflect all data because of the polling interval. The ONTAP health monitor can contact and monitor the bridge if the value in the “Status” column is “ok”, and other information, such as the worldwide name (WWN), is displayed.

The following example shows that the FC-to-SAS bridges are configured:

```

controller_A_1::> storage bridge show

Bridge                Symbolic Name Is Monitored  Monitor Status
Vendor Model          Bridge WWN
-----
-----
ATTO_10.10.20.10  atto01          true           ok           Atto
FibreBridge 7500N  20000010867038c0
ATTO_10.10.20.11  atto02          true           ok           Atto
FibreBridge 7500N  20000010867033c0
ATTO_10.10.20.12  atto03          true           ok           Atto
FibreBridge 7500N  20000010867030c0
ATTO_10.10.20.13  atto04          true           ok           Atto
FibreBridge 7500N  2000001086703b80

4 entries were displayed

controller_A_1::>

```

25. 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.

Related information

[In-band management of the FC-to-SAS bridges](#)

Hot-swapping a FibreBridge 7500N with a 7600N bridge

You can hot-swap a FibreBridge 7500N bridge with a 7600N bridge.

About this task

If you will be using in-band management of the bridge rather than IP management, the steps for configuring the Ethernet port and IP settings can be skipped, as noted in the relevant steps.



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

Steps

1. If the bridge is in a fabric-attached MetroCluster configuration, disable all of the switch ports that connect to the bridge FC port or ports.
2. From the ONTAP cluster prompt, remove the bridge undergoing maintenance from health monitoring:
 - a. Remove the bridge:

```
storage bridge remove -name bridge-name
```
 - b. View the list of monitored bridges and confirm that the removed bridge is not present:

```
storage bridge show
```
3. Properly ground yourself.
4. Remove the power cables connected to the bridge to power down the bridge.
5. Disconnect the cables that are connected to the old bridge.

You should make note of the port to which each cable was connected.

6. Remove the old bridge from the rack.
7. Install the new bridge into the rack.
8. Reconnect the power cord and shielded Ethernet cable.



You must not reconnect the SAS or FC cables at this time.

9. Connect the bridge to a power source, and then turn it on.

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

10. If configuring for in-band management, connect a cable from FibreBridge RS-232 serial port to the serial (COM) port on a personal computer.

The serial connection will be used for initial configuration, and then in-band management via ONTAP and the FC ports can be used to monitor and manage the bridge.

11. If configuring for in-band management, connect a cable from FibreBridge RS-232 serial port to the serial (COM) port on a personal computer.

The serial connection will be used for initial configuration, and then in-band management via ONTAP and the FC ports can be used to monitor and manage the bridge.

12. If configuring for IP management, configure the Ethernet management 1 port for each bridge by following the procedure in section 2.0 of the *ATTO FibreBridge Installation and Operation Manual* for your bridge model.

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

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

13. Configure the bridges.

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

The *ATTO FibreBridge Installation and Operation Manual* for your bridge model has the most current information on available commands and how to use them.



Do not configure time synchronization on FibreBridge 7600N. The time synchronization for FibreBridge 7600N is set to the cluster time after the bridge is discovered by ONTAP. It is also synchronized periodically once a day. The time zone used is GMT and is not changeable.

a. If configuring for IP management, configure the IP settings of the bridge.

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

If using the CLI, you must run the following commands:

```
set ipaddress mp1 ip-address

set ipsubnetmask mp1 subnet-mask

set ipgateway mp1 x.x.x.x

set ipdhcp mp1 disabled

set ethernetspeed mp1 1000
```

b. Configure the bridge name.

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

Example bridge names for one stack group on each site:

- bridge_A_1a
- bridge_A_1b
- bridge_B_1a
- bridge_B_1b

If using the CLI, you must run the following command:

```
set bridgename bridgename
```

c. If running ONTAP 9.4 or earlier, enable SNMP on the bridge:

```
set SNMP enabled
```

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

14. Configure the bridge FC ports.

a. Configure the data rate/speed of the bridge FC ports.

The supported FC data rate depends on your model bridge.

- The FibreBridge 7600N bridge supports up to 32, 16, or 8 Gbps.
- The FibreBridge 7500N bridge supports up to 16, 8, or 4 Gbps.



The FCDataRate speed you select is limited to the maximum speed supported by both the bridge and the FC port of the controller module or switch to which the bridge port connects. Cabling distances must not exceed the limitations of the SFPs and other hardware.

If using the CLI, you must run the following command:

```
set FCDataRate port-number port-speed
```

b. You must configure or disable the FC2 port.

- If you are using the second port, you must repeat the previous substeps for the FC2 port.
- If you are not using the second port, then you must disable the unused port:

```
FCPortDisable port-number
```

The following example shows the disabling of FC port 2:

```
FCPortDisable 2  
  
Fibre Channel Port 2 has been disabled.
```

c. Disable the unused SAS ports:

```
SASPortDisable sas-port
```



SAS ports A through D are enabled by default. You must disable the SAS ports that are not being used.

If only SAS port A is used, then SAS ports B, C, and D must be disabled. The following example shows disabling of SAS port B. You must similarly disable SAS ports C and D:

```
SASPortDisable b
```

```
SAS Port B has been disabled.
```

15. Secure access to the bridge and save the bridge's configuration.

- a. From the controller prompt check the status of the bridges:

```
storage bridge show
```

The output shows which bridge is not secured.

- b. Check the status of the unsecured bridge's ports:

```
info
```

The output shows the status of Ethernet ports MP1 and MP2.

- c. If Ethernet port MP1 is enabled, run the following command:

```
set EthernetPort mp1 disabled
```



If Ethernet port MP2 is also enabled, repeat the previous substep for port MP2.

- d. Save the bridge's configuration.

You must run the following commands:

```
SaveConfiguration
```

```
FirmwareRestart
```

You are prompted to restart the bridge.

16. Connect the FC cables to the same ports on the new bridge.

17. Update the FibreBridge firmware on each bridge.

[Update firmware on a FibreBridge bridge](#)

18. Reconnect the SAS cables to the same ports on the new bridge.



Wait at least 10 seconds before connecting the port. The SAS cable connectors are keyed; when oriented correctly into a SAS port, the connector clicks into place and the disk shelf SAS port LNK LED illuminates green. For disk shelves, you insert a SAS cable connector with the pull tab oriented down (on the underside of the connector). For controllers, the orientation of SAS ports can vary depending on the platform model; therefore, the correct orientation of the SAS cable connector varies.

19. Verify that each bridge can see all of the disk drives and disk shelves to which the bridge is connected:

```
sastargets
```

The output shows the devices (disks and disk shelves) to which the bridge is connected. The output lines are sequentially numbered so that you can quickly count the devices.

The following output shows that 10 disks are connected:

```

Tgt VendorID ProductID      Type      SerialNumber
  0 NETAPP    X410_S15K6288A15 DISK      3QP1CLE300009940UHJV
  1 NETAPP    X410_S15K6288A15 DISK      3QP1ELF600009940V1BV
  2 NETAPP    X410_S15K6288A15 DISK      3QP1G3EW00009940U2M0
  3 NETAPP    X410_S15K6288A15 DISK      3QP1EWMP00009940U1X5
  4 NETAPP    X410_S15K6288A15 DISK      3QP1FZLE00009940G8YU
  5 NETAPP    X410_S15K6288A15 DISK      3QP1FZLF00009940TZKZ
  6 NETAPP    X410_S15K6288A15 DISK      3QP1CEB400009939MGXL
  7 NETAPP    X410_S15K6288A15 DISK      3QP1G7A900009939FNNT
  8 NETAPP    X410_S15K6288A15 DISK      3QP1FY0T00009940G8PA
  9 NETAPP    X410_S15K6288A15 DISK      3QP1FXW600009940VERQ

```

- Verify that the command output shows that the bridge is connected to all of the appropriate disks and disk shelves in the stack.

If the output is...	Then...
Correct	Repeat the previous step for each remaining bridge.
Not correct	<ol style="list-style-type: none"> Check for loose SAS cables or correct the SAS cabling by repeating Step 18. Repeat the previous step.

- If the bridge is in a fabric-attached MetroCluster configuration, reenabling the FC switch port that you disabled at the beginning of this procedure.

This should be the port that connects to the bridge.

- From the system console of both controller modules, verify that all of the controller modules have access through the new bridge to the disk shelves (that is, that the system is cabled for Multipath HA):

```
run local sysconfig
```



It might take up to a minute for the system to complete discovery.

If the output does not indicate Multipath HA, you must correct the SAS and FC cabling because not all of the disk drives are accessible through the new bridge.

The following output states that the system is cabled for Multipath HA:

```
NetApp Release 8.3.2: Tue Jan 26 01:41:49 PDT 2016
System ID: 1231231231 (node_A_1); partner ID: 4564564564 (node_A_2)
System Serial Number: 700000123123 (node_A_1); partner Serial Number:
700000456456 (node_A_2)
System Rev: B0
System Storage Configuration: Multi-Path HA
System ACP Connectivity: NA
```



When the system is not cabled as Multipath HA, restarting a bridge might cause loss of access to the disk drives and result in a multi-disk panic.

23. If running ONTAP 9.4 or earlier, verify that the bridge is configured for SNMP.

If you are using the bridge CLI, run the following command:

```
get snmp
```

24. From the ONTAP cluster prompt, add the bridge to health monitoring:

- a. Add the bridge, using the command for your version of ONTAP:

ONTAP version	Command
9.5 and later	<code>storage bridge add -address 0.0.0.0 -managed-by in-band -name <i>bridge-name</i></code>
9.4 and earlier	<code>storage bridge add -address <i>bridge-ip-address</i> -name <i>bridge-name</i></code>

- b. Verify that the bridge has been added and is properly configured:

```
storage bridge show
```

It might take as long as 15 minutes to reflect all data because of the polling interval. The ONTAP health monitor can contact and monitor the bridge if the value in the “Status” column is “ok”, and other information, such as the worldwide name (WWN), is displayed.

The following example shows that the FC-to-SAS bridges are configured:

```

controller_A_1::> storage bridge show

Bridge                Symbolic Name Is Monitored  Monitor Status
Vendor Model          Bridge WWN
-----
-----
ATTO_10.10.20.10  atto01          true           ok           Atto
FibreBridge 7500N    20000010867038c0
ATTO_10.10.20.11  atto02          true           ok           Atto
FibreBridge 7500N    20000010867033c0
ATTO_10.10.20.12  atto03          true           ok           Atto
FibreBridge 7500N    20000010867030c0
ATTO_10.10.20.13  atto04          true           ok           Atto
FibreBridge 7500N    2000001086703b80

4 entries were displayed

controller_A_1::>

```

25. 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.

Related information

[In-band management of the FC-to-SAS bridges](#)

Hot-swapping a FibreBridge 6500N bridge with a FibreBridge 7600N or 7500N bridge

You can hot-swap a FibreBridge 6500N bridge with a FibreBridge 7600N or 7500N bridge to replace a failed bridge or upgrade your bridge in a fabric-attached or a bridge-attached MetroCluster configuration.

About this task

- This procedure is for hot-swapping a single FibreBridge 6500N bridge with single FibreBridge 7600N or 7500N bridge.
- When you hot-swap a FibreBridge 6500N bridge with a FibreBridge 7600N or 7500N bridge, you must use only one FC port and one SAS port on the FibreBridge 7600N or 7500N bridge.
- If you will be using in-band management of the bridge rather than IP management, the steps for configuring the Ethernet port and IP settings can be skipped, as noted in the relevant steps.



If you are hot-swapping both FibreBridge 6500N bridges in a pair, you must use the [Consolidate Multiple Storage Stacks](#) procedure for zoning instructions. By replacing both FibreBridge 6500N bridges on the bridge, you can take advantage of the additional ports on the FibreBridge 7600N or 7500N bridge.



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

Steps

1. Do one of the following:
 - If the failed bridge is in a fabric-attached MetroCluster configuration, disable the switch port that connects to the bridge FC port.
 - If the failed bridge is in a stretch MetroCluster configuration, use either one of the available FC ports.
2. From the ONTAP cluster prompt, remove the bridge undergoing maintenance from health monitoring:
 - a. Remove the bridge:

```
storage bridge remove -name bridge-name
```

- b. View the list of monitored bridges and confirm that the removed bridge is not present:

```
storage bridge show
```

3. Properly ground yourself.
4. Turn off the power switch of the bridge.
5. Disconnect the cables connected from the shelf to the FibreBridge 6500N bridge ports and power cables.

You should make note of the ports that each cable was connected to.

6. Remove the FibreBridge 6500N bridge that you need to replace from the rack.
7. Install the new FibreBridge 7600N or 7500N bridge into the rack.
8. Reconnect the power cord and, if necessary, the shielded Ethernet cable.



Do not reconnect the SAS or FC cables at this time.

9. If configuring for in-band management, connect a cable from FibreBridge RS-232 serial port to the serial (COM) port on a personal computer.

The serial connection will be used for initial configuration, and then in-band management via ONTAP and the FC ports can be used to monitor and manage the bridge.

10. If configuring for IP management, connect the Ethernet management 1 port on each bridge to your network by using an Ethernet cable.

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

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

11. If configuring for IP management, configure the Ethernet management 1 port for each bridge by following the procedure in section 2.0 of the *ATTO FibreBridge Installation and Operation Manual* for your bridge model.

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

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

12. Configure the bridge.

If you retrieved the configuration information from the old bridge, use the information to configure the new bridge.

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

The *ATTO FibreBridge Installation and Operation Manual* for your bridge model has the most current information on available commands and how to use them.



Do not configure time synchronization on ATTO FibreBridge 7600N or 7500N. The time synchronization for ATTO FibreBridge 7600N or 7500N is set to the cluster time after the bridge is discovered by ONTAP. It is also synchronized periodically once a day. The time zone used is GMT and is not changeable.

- a. If configuring for IP management, configure the IP settings of the bridge.

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

If using the CLI, you must run the following commands:

```
set ipaddress mp1 ip-address
```

```
set ipsubnetmask mp1 subnet-mask
```

```
set ipgateway mp1 x.x.x.x
```

```
set ipdhcp mp1 disabled
set ethernetspeed mp1 1000
```

b. Configure the bridge name.

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

Example bridge names for one stack group on each site:

- bridge_A_1a
- bridge_A_1b
- bridge_B_1a
- bridge_B_1b

If using the CLI, you must run the following command:

```
set bridgename bridgename
```

c. If running ONTAP 9.4 or earlier, enable SNMP on the bridge:

```
set SNMP enabled
```

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

13. Configure the bridge FC ports.

a. Configure the data rate/speed of the bridge FC ports.

The supported FC data rate depends on your model bridge.

- The FibreBridge 7600N bridge supports up to 32, 16, or 8 Gbps.
- The FibreBridge 7500N bridge supports up to 16, 8, or 4 Gbps.
- The FibreBridge 6500N bridge supports up to 8, 4, or 2 Gbps.



The FCDataRate speed you select is limited to the maximum speed supported by both the bridge and the switch to which the bridge port connects. Cabling distances must not exceed the limitations of the SFPs and other hardware.

If using the CLI, you must run the following command:

```
set FCDataRate port-number port-speed
```

b. If you are configuring a FibreBridge 7500N or 6500N bridge, configure the connection mode that the port uses to ptp.



The FCConnMode setting is not required when configuring a FibreBridge 7600N bridge.

If using the CLI, you must run the following command:

```
set FCConnMode port-number ptp
```

c. If you are configuring a FibreBridge 7600N or 7500N bridge, you must configure or disable the FC2 port.

- If you are using the second port, you must repeat the previous substeps for the FC2 port.
- If you are not using the second port, then you must disable the port:

```
FCPortDisable port-number
```

d. If you are configuring a FibreBridge 7600N or 7500N bridge, disable the unused SAS ports:

```
SASPortDisable sas-port
```



SAS ports A through D are enabled by default. You must disable the SAS ports that are not being used. If only SAS port A is used, then SAS ports B, C, and D must be disabled.

14. Secure access to the bridge and save the bridge's configuration.

a. From the controller prompt check the status of the bridges:

```
storage bridge show
```

The output shows which bridge is not secured.

b. Check the status of the unsecured bridge's ports:

```
info
```

The output shows the status of Ethernet ports MP1 and MP2.

c. If Ethernet port MP1 is enabled, run the following command:

```
set EthernetPort mp1 disabled
```



If Ethernet port MP2 is also enabled, repeat the previous substep for port MP2.

d. Save the bridge's configuration.

You must run the following commands:

```
SaveConfiguration
```

```
FirmwareRestart
```

You are prompted to restart the bridge.

15. Turn on Health Monitoring for the FibreBridge 7600N or 7500N bridge.

16. Connect the FC cables to the Fibre Channel 1 ports on the new bridge.

You must cable the FC port to the same switch or controller port that the FibreBridge 6500N bridge had been connected to.

17. Update the FibreBridge firmware on each bridge.

If the new bridge is the same type as the partner bridge, upgrade to the same firmware as the partner bridge. If the new bridge is a different type to the partner bridge, upgrade to the latest firmware and version of ONTAP supported by the bridge.

Update firmware on a FibreBridge bridge

18. Reconnect the SAS cables to the SAS A ports on the new bridge.

The SAS port must be cabled to the same shelf port that the FibreBridge 6500N bridge had been connected to.



Do not force a connector into a port. The mini-SAS cables are keyed; when oriented correctly into a SAS port, the SAS cable clicks into place and the disk shelf SAS port LNK LED illuminates green. For disk shelves, you insert a SAS cable connector with the pull tab oriented down (on the underside of the connector). For controllers, the orientation of SAS ports can vary depending on the platform model; therefore, the correct orientation of the SAS cable connector varies.

19. Verify that the bridge can detect all of the disk drives and disk shelves it is connected to.

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

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

Tgt	VendorID	ProductID	Type	SerialNumber
0	NETAPP	X410_S15K6288A15	DISK	3QP1CLE300009940UHJV
1	NETAPP	X410_S15K6288A15	DISK	3QP1ELF600009940V1BV
2	NETAPP	X410_S15K6288A15	DISK	3QP1G3EW00009940U2M0
3	NETAPP	X410_S15K6288A15	DISK	3QP1EWMP00009940U1X5
4	NETAPP	X410_S15K6288A15	DISK	3QP1FZLE00009940G8YU
5	NETAPP	X410_S15K6288A15	DISK	3QP1FZLF00009940TZKZ
6	NETAPP	X410_S15K6288A15	DISK	3QP1CEB400009939MGXL
7	NETAPP	X410_S15K6288A15	DISK	3QP1G7A900009939FNNT
8	NETAPP	X410_S15K6288A15	DISK	3QP1FY0T00009940G8PA
9	NETAPP	X410_S15K6288A15	DISK	3QP1FXW600009940VERQ



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

- Verify that the command output shows that the bridge is connected to all of the necessary disks and disk shelves in the stack.

If the output is...	Then...
Correct	Repeat the previous step for each remaining bridge.
Not correct	<ol style="list-style-type: none"> Check for loose SAS cables or correct the SAS cabling by repeating Step 18. Repeat the previous step for each remaining bridge.

- Reenable the FC switch port that connects to the bridge.
- Verify that all controllers have access through the new bridge to the disk shelves (that the system is cabled for Multipath HA), at the system console of both controllers:

```
run local sysconfig
```



It might take up to a minute for the system to complete discovery.

For example, the following output shows that the system is cabled for Multipath HA:

```
NetApp Release 8.3.2: Tue Jan 26 01:23:24 PST 2016
System ID: 1231231231 (node_A_1); partner ID: 4564564564 (node_A_2)
System Serial Number: 700000123123 (node_A_1); partner Serial Number:
700000456456 (node_A_2)
System Rev: B0
System Storage Configuration: Multi-Path HA
System ACP Connectivity: NA
```

If the command output indicates that the configuration is mixed-path or single-path HA, you must correct

the SAS and FC cabling because not all disk drives are accessible through the new bridge.



When the system is not cabled as Multipath HA, restarting a bridge might cause loss of access to the disk drives and result in a multi-disk panic.

23. From the ONTAP cluster prompt, add the bridge to health monitoring:

a. Add the bridge, using the command for your version of ONTAP:

ONTAP version	Command
9.5 and later	<code>storage bridge add -address 0.0.0.0 -managed-by in-band -name <i>bridge-name</i></code>
9.4 and earlier	<code>storage bridge add -address <i>bridge-ip-address</i> -name <i>bridge-name</i></code>

b. Verify that the bridge has been added and is properly configured:

```
storage bridge show
```

It might take as long as 15 minutes to reflect all data because of the polling interval. The ONTAP health monitor can contact and monitor the bridge if the value in the “Status” column is “ok”, and other information, such as the worldwide name (WWN), is displayed.

The following example shows that the FC-to-SAS bridges are configured:

```
controller_A_1::> storage bridge show

Bridge                Symbolic Name Is Monitored  Monitor Status
Vendor Model          Bridge WWN
-----
-----
ATTO_10.10.20.10  atto01          true          ok          Atto
FibreBridge 7500N  20000010867038c0
ATTO_10.10.20.11  atto02          true          ok          Atto
FibreBridge 7500N  20000010867033c0
ATTO_10.10.20.12  atto03          true          ok          Atto
FibreBridge 7500N  20000010867030c0
ATTO_10.10.20.13  atto04          true          ok          Atto
FibreBridge 7500N  2000001086703b80

4 entries were displayed

controller_A_1::>
```

24. 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.

25. After replacing the part, return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return & Replacements](#) page for further information.

Related information

[In-band management of the FC-to-SAS bridges](#)

Replacing a pair of FibreBridge 6500N bridges with 7600N or 7500N bridges

To take advantage of the additional FC2 port on the FibreBridge 7600N or 7500N bridges and reduce rack utilization, you can nondisruptively replace 6500N bridges and consolidate up to four storage stacks behind a single pair of FibreBridge 7600N or 7500N bridges.

Before you begin

You need the admin password and access to an FTP or SCP server.

About this task

You should use this procedure if:

- You are replacing a pair of FibreBridge 6500N bridges with FibreBridge 7600N or 7500N bridges.

After the replacement, both bridges in the pair must be the same model.

- You previously replaced a single FibreBridge 6500N bridge with a 7600N or 7500N bridge and are now replacing the second bridge in the pair.

- You have a pair of FibreBridge 7600N or 7500N bridges with available SAS ports and you are consolidating SAS storage stacks that are currently connected using FibreBridge 6500N bridges.

This procedure is nondisruptive and takes approximately two hours to complete.

Related information

[Replacing a single FC-to-SAS bridge](#)

Verifying storage connectivity

Before replacing bridges, you should verify bridge and storage connectivity. Familiarizing yourself with the command output enables you to subsequently confirm connectivity after making configuration changes.

You can issue these commands from the admin prompt of any of the controller modules in the MetroCluster configuration at the site undergoing maintenance.

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

```
run local sysconfig -v
```

The output shows the disks attached to the initiator ports on the controller, and identifies the shelves connected to the FC-to-SAS bridges:

```
node_A_1> run local sysconfig -v
NetApp Release 9.3.2X18: Sun Dec 13 01:23:24 PST 2017
System ID: 4068741258 (node_A_1); partner ID: 4068741260 (node_B_1)
System Serial Number: 940001025471 (node_A_1)
System Rev: 70
System Storage Configuration: Multi-Path HA**<=== Configuration should
be multi-path HA**
.
.
.
slot 0: FC Host Adapter 0g (QLogic 8324 rev. 2, N-port, <UP>)**<===
Initiator port**
    Firmware rev:      7.5.0
    Flash rev:         0.0.0
    Host Port Id:      0x60130
    FC Node Name:      5:00a:098201:bae312
    FC Port Name:      5:00a:098201:bae312
    SFP Vendor:        UTILITIES CORP.
    SFP Part Number:   FTLF8529P3BCVAN1
    SFP Serial Number: URQ0Q9R
    SFP Capabilities:  4, 8 or 16 Gbit
    Link Data Rate:    16 Gbit
    Switch Port:       brcd6505-fcs40:1
**<List of disks visible to port\>**
    ID      Vendor  Model          FW      Size
    brcd6505-fcs29:12.126L1527  : NETAPP  X302_HJUPI01TSSM NA04
```

```

847.5GB (1953525168 512B/sect)
      brcd6505-fcs29:12.126L1528      : NETAPP      X302_HJUPI01TSSA NA02
847.5GB (1953525168 512B/sect)
      .
      .
      .
      **<List of FC-to-SAS bridges visible to port\>**
      FC-to-SAS Bridge:
      brcd6505-fcs40:12.126L0          : ATTO          FibreBridge6500N 1.61
FB6500N102980
      brcd6505-fcs42:13.126L0          : ATTO          FibreBridge6500N 1.61
FB6500N102980
      brcd6505-fcs42:6.126L0           : ATTO          FibreBridge6500N 1.61
FB6500N101167
      brcd6505-fcs42:7.126L0          : ATTO          FibreBridge6500N 1.61
FB6500N102974
      .
      .
      .
      **<List of storage shelves visible to port\>**
      brcd6505-fcs40:12.shelf6: DS4243  Firmware rev. IOM3 A: 0200
IOM3 B: 0200
      brcd6505-fcs40:12.shelf8: DS4243  Firmware rev. IOM3 A: 0200
IOM3 B: 0200
      .
      .
      .

```

Hot-swapping FibreBridge 6500N bridges to create a pair of FibreBridge 7600N or 7500N bridges

To hot-swap one or two FibreBridge 6500N bridges to create a configuration with a pair of FibreBridge 7600N or 7500N bridges, you must replace the bridges one at a time and follow the correct cabling procedure. The new cabling is different from the original cabling.


About this task

You can also use this procedure if the following conditions are true:

- You are replacing a pair of FibreBridge 6500N bridges that are both connected to the same stack of SAS storage.
- You previously replaced one FibreBridge 6500N bridge in the pair, and your storage stack is configured with one FibreBridge 6500N bridge and one FibreBridge 7600N or 7500N bridge.


In this case, you should start with the step below to hot-swap the bottom FibreBridge 6500N bridge with a FibreBridge 7600N or 7500N bridge.

The following diagram shows an example of the initial configuration, in which four FibreBridge 6500N bridges are connecting two SAS storage stacks:



Steps

1. Using the following guidelines, hot-swap the top FibreBridge 6500N bridge with a FibreBridge 7600N or 7500N bridge using the procedure in [Hot-swapping a FibreBridge 6500N bridge with a FibreBridge 7600N or 7500N bridge](#):
 - Connect the FibreBridge 7600N or 7500N bridge FC1 port to the switch or controller.
This is the same connection that was made to the FibreBridge 6500N bridge FC1 port.
 - Do not connect the FibreBridge 7600N or 7500N bridge FC2 port at this time. The following diagram shows that bridge_A_1a has been replaced and is now a FibreBridge 7600N or 7500N bridge:



2. Confirm connectivity to the bridge-connected disks and that the new FibreBridge 7500N is visible in the configuration:

```
run local sysconfig -v
```

```
node_A_1> run local sysconfig -v
NetApp Release 9.3.2X18: Sun Dec 13 01:23:24 PST 2015
System ID: 0536872165 (node_A_1); partner ID: 0536872141 (node_B_1)
System Serial Number: 940001025465 (node_A_1)
System Rev: 70
System Storage Configuration: Multi-Path HA**<=== Configuration should
be multi-path HA**
.
```

```

.
.
slot 0: FC Host Adapter 0g (QLogic 8324 rev. 2, N-port, <UP>)**<===
Initiator port**
    Firmware rev:      7.5.0
    Flash rev:         0.0.0
    Host Port Id:      0x60100
    FC Node Name:      5:00a:098201:bae312
    FC Port Name:      5:00a:098201:bae312
    SFP Vendor:        FINISAR CORP.
    SFP Part Number:   FTLF8529P3BCVAN1
    SFP Serial Number: URQ0R1R
    SFP Capabilities:  4, 8 or 16 Gbit
    Link Data Rate:    16 Gbit
    Switch Port:       brcd6505-fcs40:1
**<List of disks visible to port\>**
    ID      Vendor  Model                      FW      Size
brcd6505-fcs40:12.126L1527  : NETAPP  X302_HJUPI01TSSM NA04
847.5GB (1953525168 512B/sect)
brcd6505-fcs40:12.126L1528  : NETAPP  X302_HJUPI01TSSA NA02
847.5GB (1953525168 512B/sect)
.
.
.
**<List of FC-to-SAS bridges visible to port\>**
FC-to-SAS Bridge:
brcd6505-fcs40:12.126L0      : ATTO    FibreBridge7500N A30H
FB7500N100104**<===**
brcd6505-fcs42:13.126L0     : ATTO    FibreBridge6500N 1.61
FB6500N102980
brcd6505-fcs42:6.126L0     : ATTO    FibreBridge6500N 1.61
FB6500N101167
brcd6505-fcs42:7.126L0     : ATTO    FibreBridge6500N 1.61
FB6500N102974
.
.
.
**<List of storage shelves visible to port\>**
brcd6505-fcs40:12.shelf6: DS4243  Firmware rev. IOM3 A: 0200
IOM3 B: 0200
brcd6505-fcs40:12.shelf8: DS4243  Firmware rev. IOM3 A: 0200
IOM3 B: 0200
.
.
.


```

3. Using the following guidelines, hot-swap the bottom FibreBridge 6500N bridge with a FibreBridge 7600N or 7500N bridge using the procedure in [Hot-swapping a FibreBridge 6500N bridge with a FibreBridge 7600N or 7500N bridge](#):

- Connect the FibreBridge 7600N or 7500N bridge FC2 port to the switch or controller.

This is the same connection that was made to the FibreBridge 6500N bridge FC1 port.


- Do not connect the FibreBridge 7600N or 7500N bridge FC1 port at this time.



4. Confirm connectivity to the bridge-connected disks:

```
run local sysconfig -v
```

The output shows the disks attached to the initiator ports on the controller, and identifies the shelves connected to the FC-to-SAS bridges:



```

node_A_1> run local sysconfig -v
NetApp Release 9.3.2X18: Sun Dec 13 01:23:24 PST 2015
System ID: 0536872165 (node_A_1); partner ID: 0536872141 (node_B_1)
System Serial Number: 940001025465 (node_A_1)
System Rev: 70
System Storage Configuration: Multi-Path HA**<=== Configuration should
be multi-path HA**
.
.
.
slot 0: FC Host Adapter 0g (QLogic 8324 rev. 2, N-port, <UP>)**<===
Initiator port**
    Firmware rev:      7.5.0
    Flash rev:         0.0.0
    Host Port Id:      0x60100
    FC Node Name:      5:00a:098201:bae312
    FC Port Name:      5:00a:098201:bae312
    SFP Vendor:        FINISAR CORP.
    SFP Part Number:   FTLF8529P3BCVAN1
    SFP Serial Number: URQ0R1R
    SFP Capabilities:  4, 8 or 16 Gbit
    Link Data Rate:    16 Gbit
    Switch Port:       brcd6505-fcs40:1
**<List of disks visible to port\>**
    ID      Vendor  Model          FW      Size
brcd6505-fcs40:12.126L1527      : NETAPP      X302_HJUPI01TSSM NA04
847.5GB (1953525168 512B/sect)
brcd6505-fcs40:12.126L1528      : NETAPP      X302_HJUPI01TSSA NA02
847.5GB (1953525168 512B/sect)
.
.
.
**<List of FC-to-SAS bridges visible to port\>**
FC-to-SAS Bridge:
brcd6505-fcs40:12.126L0          : ATTO        FibreBridge7500N A30H
FB7500N100104
brcd6505-fcs42:13.126L0          : ATTO        FibreBridge7500N A30H
FB7500N100104
.
.
.
**<List of storage shelves visible to port\>**
brcd6505-fcs40:12.shelf6: DS4243  Firmware rev. IOM3 A: 0200
IOM3 B: 0200
brcd6505-fcs40:12.shelf8: DS4243  Firmware rev. IOM3 A: 0200
IOM3 B: 0200

```

•
•
•

Cabling the bridge SAS ports when consolidating storage behind FibreBridge 7600N or 7500N bridges

When consolidating multiple SAS storage stacks behind a single pair of FibreBridge 7600N or 7500N bridges with available SAS ports, you must move the top and bottom SAS cables to the new bridges.

About this task

The FibreBridge 6500N bridge SAS ports use QSFP connectors. The FibreBridge 7600N or 7500N bridge SAS ports use mini-SAS connectors.



If you insert a SAS cable into the wrong port, when you remove the cable from a SAS port, you must wait at least 120 seconds before plugging the cable into a different SAS port. If you fail to do so, the system will not recognize that the cable has been moved to another port.




Wait at least 10 seconds before connecting the port. The SAS cable connectors are keyed; when oriented correctly into a SAS port, the connector clicks into place and the disk shelf SAS port LNK LED illuminates green. For disk shelves, you insert a SAS cable connector with the pull tab oriented down (on the underside of the connector).

Steps


1. Remove the cable that connects the SAS A port of the top FibreBridge 6500N bridge to the top SAS shelf, being sure to note the SAS port on the storage shelf to which it connects.

The cable is shown in blue in the following example:




2. Using a cable with a mini-SAS connector, connect the same SAS port on the storage shelf to the SAS B port of the top FibreBridge 7600N or 7500N bridge.

The cable is shown in blue in the following example:




- Remove the cable that connects the SAS A port of the bottom FibreBridge 6500N bridge to the top SAS shelf, being sure to note the SAS port on the storage shelf to which it connects.

This cable is shown in green in the following example:



- Using a cable with a mini-SAS connector, connect the same SAS port on the storage shelf to the SAS B port of the bottom FibreBridge 7600N or 7500N bridge.

This cable is shown in green in the following example:



5. Confirm connectivity to the bridge-connected disks:

```
run local sysconfig -v
```

The output shows the disks attached to the initiator ports on the controller, and identifies the shelves connected to the FC-to-SAS bridges:

```
node_A_1> run local sysconfig -v
NetApp Release 9.3.2X18: Sun Dec 13 01:23:24 PST 2015
System ID: 0536872165 (node_A_1); partner ID: 0536872141 (node_B_1)
System Serial Number: 940001025465 (node_A_1)
System Rev: 70
System Storage Configuration: Multi-Path HA**<=== Configuration should
be multi-path HA**
.
.
.
slot 0: FC Host Adapter 0g (QLogic 8324 rev. 2, N-port, <UP>)**<===
Initiator port**
    Firmware rev:      7.5.0
    Flash rev:         0.0.0
    Host Port Id:      0x60100
    FC Node Name:      5:00a:098201:bae312
    FC Port Name:      5:00a:098201:bae312
    SFP Vendor:        FINISAR CORP.
    SFP Part Number:   FTLF8529P3BCVAN1
    SFP Serial Number: URQ0R1R
    SFP Capabilities:  4, 8 or 16 Gbit
    Link Data Rate:    16 Gbit
    Switch Port:       brcd6505-fcs40:1
```


```

**<List of disks visible to port\>**
      ID      Vendor   Model                               FW      Size
      brcd6505-fcs40:12.126L1527    : NETAPP  X302_HJUPI01TSSM NA04
847.5GB (1953525168 512B/sect)
      brcd6505-fcs40:12.126L1528    : NETAPP  X302_HJUPI01TSSA NA02
847.5GB (1953525168 512B/sect)
      .
      .
      .
**<List of FC-to-SAS bridges visible to port\>**
FC-to-SAS Bridge:
      brcd6505-fcs40:12.126L0        : ATTO      FibreBridge7500N A30H
FB7500N100104
      brcd6505-fcs42:13.126L0        : ATTO      FibreBridge7500N A30H
FB7500N100104
      .
      .
      .
**<List of storage shelves visible to port\>**
      brcd6505-fcs40:12.shelf6: DS4243  Firmware rev. IOM3 A: 0200
IOM3 B: 0200
      brcd6505-fcs40:12.shelf8: DS4243  Firmware rev. IOM3 A: 0200
IOM3 B: 0200
      .
      .
      .

```

6. Remove the old FibreBridge 6500N bridges that are no longer connected to the SAS storage.
7. Wait two minutes for the system to recognize the changes.
8. If the system was cabled incorrectly, remove the cable, correct the cabling, and then reconnect the correct cable.
9. If necessary, repeat the preceding steps to move up to two additional SAS stacks behind the new FibreBridge 7600N or 7500N bridges, using SAS ports C and then D.

Each SAS stack must be connected to the same SAS port on the top and bottom bridge. For example, if the top connection of the stack is connected to the top bridge SAS B port, the bottom connection must be connected to the SAS B port of the bottom bridge.



Updating zoning when adding FibreBridge 7600N or 7500N bridges to a configuration

The zoning must be changed when you are replacing FibreBridge 6500N bridges with FibreBridge 7600N or 7500N bridges and using both FC ports on the FibreBridge 7600N or 7500N bridges. The required changes depend on whether you are running a version of ONTAP earlier than 9.1 or 9.1 and later.

Updating zoning when adding FibreBridge 7500N bridges to a configuration (prior to ONTAP 9.1)


The zoning must be changed when you are replacing FibreBridge 6500N bridges with FibreBridge 7500N bridges and using both FC ports on the FibreBridge 7500N bridges. Each zone can have no more than four initiator ports. The zoning you use depends on whether you are running ONTAP prior to version 9.1 or 9.1 and later.

About this task

The specific zoning in this task is for versions of ONTAP prior to version 9.1.

The zoning changes are required to avoid issues with ONTAP, which requires that no more than four FC initiator ports can have a path to a disk. After recabling to consolidate the shelves, the existing zoning would result in each disk being reachable by eight FC ports. You must change the zoning to reduce the initiator ports in each zone to four.

The following diagram shows the zoning on site_A before the changes:



Steps

1. Update the storage zones for the FC switches by removing half of the initiator ports from each existing zone and creating new zones for the FibreBridge 7500N FC2 ports.

The zones for the new FC2 ports will contain the initiator ports removed from the existing zones. In the diagrams, these zones are shown with dashed lines.


For details about the zoning commands, see the FC switch sections of the [Fabric-attached MetroCluster installation and configuration](#) or [Stretch MetroCluster installation and configuration](#).

The following examples show the storage zones and the ports in each zone before and after the consolidation. The ports are identified by *domain, port* pairs.

- Domain 5 consists of switch FC_switch_A_1.
- Domain 6 consists of switch FC_switch_A_2.
- Domain 7 consists of switch FC_switch_B_1.
- Domain 8 consists of switch FC_switch_B_2.

Before or after consolidation	Zone	Domains and ports	Colors in diagrams (The diagrams only show Site A)
Zones before the consolidation. There is a zone for each FC port on the four FibreBridge 6500N bridges.	STOR_A_1a-FC1	5,1; 5,2; 5,4; 5,5; 7,1; 7,2; 7,4; 7,5; 5,6	Purple + dashed purple + blue
	STOR_A_1b-FC1	6,1; 6,2; 6,4; 6,5; 8,1; 8,2; 8,4; 8,5; 6,6	Brown + dashed brown + green
	STOR_A_2a-FC1	5,1; 5,2; 5,4; 5,5; 7,1; 7,2; 7,4; 7,5; 5,7	Purple + dashed purple + red
	STOR_A_2b-FC1	6,1; 6,2; 6,4; 6,5; 8,1; 8,2; 8,4; 8,5; 6,7	Brown + dashed brown + orange
Zones after the consolidation. There is a zone for each FC port on the two FibreBridge 7500N bridges.	STOR_A_1a-FC1	7,1; 7,4; 5,1; 5,4; 5,6	Purple + blue
	STOR_A_1b-FC1	7,2; 7,5; 5,2; 5,5; 5,7	Dashed purple + red
	STOR_A_1a-FC2	8,1; 8,4; 6,1; 6,4; 6,6	Brown + green
	STOR_A_1b-FC2	8,2; 8,5; 6,2; 6,5; 6,7	Dashed brown + orange

The following diagram shows zoning at site_A after the consolidation:



Updating zoning when adding FibreBridge 7600N or 7500N bridges to a configuration (ONTAP 9.1 and later)

The zoning must be changed when you are replacing FibreBridge 6500N bridges with FibreBridge 7600N or 7500N bridges and using both FC ports on the FibreBridge 7600N or 7500N bridges. Each zone can have no more than four initiator ports.

About this task

- This task applies to ONTAP 9.1 and later.
- FibreBridge 7600N bridges are supported in ONTAP 9.6 and later.
- The specific zoning in this task is for ONTAP 9.1 and later.
- The zoning changes are required to avoid issues with ONTAP, which requires that no more than four FC initiator ports can have a path to a disk.

After recabling to consolidate the shelves, the existing zoning would result in each disk being reachable by eight FC ports. You must change the zoning to reduce the initiator ports in each zone to four.

Step

1. Update the storage zones for the FC switches by removing half of the initiator ports from each existing zone and creating new zones for the FibreBridge 7600N or 7500N FC2 ports.

The zones for the new FC2 ports will contain the initiator ports removed from the existing zones.

Refer to the FC switch section of [Fabric-attached MetroCluster installation and configuration](#) for details about the zoning commands.

Cabling the second bridge FC port when adding FibreBridge 7600N or 7500N bridges to a configuration


To provide multiple paths to the storage stacks, you can cable the second FC port on each FibreBridge 7600N or 7500N bridge when you have added the FibreBridge 7600N or 7500N bridge to your configuration.

Before you begin


The zoning must have been adjusted to provide zones for the second FC ports.

Steps

1. Cable the FC2 port of the top bridge to the correct port on FC_switch_A_2.



2. Cable the FC1 port of the bottom bridge to the correct port on FC_switch_A_1.



3. Confirm connectivity to the bridge-connected disks:

```
run local sysconfig -v
```

The output shows the disks attached to the initiator ports on the controller, and identifies the shelves connected to the FC-to-SAS bridges:

```
node_A_1> run local sysconfig -v
NetApp Release 9.3.2X18: Sun Dec 13 01:23:24 PST 2015
System ID: 0536872165 (node_A_1); partner ID: 0536872141 (node_B_1)
System Serial Number: 940001025465 (node_A_1)
System Rev: 70
System Storage Configuration: Multi-Path HA**<=== Configuration should
```

```

be multi-path HA**
.
.
.
slot 0: FC Host Adapter 0g (QLogic 8324 rev. 2, N-port, <UP>)**<===
Initiator port**
    Firmware rev:      7.5.0
    Flash rev:         0.0.0
    Host Port Id:      0x60100
    FC Node Name:      5:00a:098201:bae312
    FC Port Name:      5:00a:098201:bae312
    SFP Vendor:        FINISAR CORP.
    SFP Part Number:   FTLF8529P3BCVAN1
    SFP Serial Number: URQ0R1R
    SFP Capabilities:  4, 8 or 16 Gbit
    Link Data Rate:    16 Gbit
    Switch Port:       brcd6505-fcs40:1
**<List of disks visible to port\>**
    ID      Vendor  Model          FW      Size
    brcd6505-fcs40:12.126L1527  : NETAPP  X302_HJUPI01TSSM NA04
847.5GB (1953525168 512B/sect)
    brcd6505-fcs40:12.126L1528  : NETAPP  X302_HJUPI01TSSA NA02
847.5GB (1953525168 512B/sect)
.
.
.
**<List of FC-to-SAS bridges visible to port\>**
FC-to-SAS Bridge:
    brcd6505-fcs40:12.126L0      : ATTO      FibreBridge7500N A30H
FB7500N100104
    brcd6505-fcs42:13.126L0      : ATTO      FibreBridge7500N A30H
FB7500N100104
.
.
.
**<List of storage shelves visible to port\>**
    brcd6505-fcs40:12.shelf6: DS4243  Firmware rev. IOM3 A: 0200
IOM3 B: 0200
    brcd6505-fcs40:12.shelf8: DS4243  Firmware rev. IOM3 A: 0200
IOM3 B: 0200
.
.
.

```

Disabling unused SAS ports on the FC-to-SAS bridges

After making cabling changes to the bridge, you should disable any unused SAS ports on FC-to-SAS bridges to avoid health monitor alerts related to the unused ports.

Steps

1. Disable unused SAS ports on the top FC-to-SAS bridge:
 - a. Log in to the bridge CLI.
 - b. Disable any unused ports.



If you have configured an ATTO 7500N bridge, then all of the SAS ports (A through D) are enabled by default, and you must disable the SAS ports that are not being used:

```
SASPortDisable sas port
```

If SAS ports A and B are used, then SAS ports C and D must be disabled. In the following example, the unused SAS ports C and D are disabled:

```
Ready. *
SASPortDisable C

SAS Port C has been disabled.

Ready. *
SASPortDisable D

SAS Port D has been disabled.

Ready. *
```

- c. Save the bridge configuration:
SaveConfiguration

The following example shows that SAS ports C and D have been disabled. Note that the asterisk no longer appears, indicating that the configuration has been saved.

```
Ready. *
SaveConfiguration

Ready.
```

2. Repeat the previous step on the bottom FC-to-SAS bridge.

Requirements for using other interfaces to configure and manage FibreBridge bridges

You can use the combination of a serial port, Telnet, and FTP to manage the FibreBridge bridges instead of the recommended management interfaces. Your system must meet the

requirements for the applicable interface before you install the bridges.

You can use a serial port or Telnet to configure the bridge and Ethernet management 1 port, and to manage the bridge. You can use FTP to update the bridge firmware.



The *ATTO FibreBridge Installation and Operation Manual* for your model bridge has more information about management interfaces.

You can access this document on the ATTO web site by using the link provided on the ATTO FibreBridge Description page.

Serial port

When using the serial port to configure and manage a bridge, and to configure the Ethernet management 1 port, your system must meet the following requirements:

- A serial cable (which connects from the bridge serial port to a serial (COM) port on the computer you are using for setup)

The bridge serial port is RJ-45 and has the same pin-out as the controllers.

- A terminal emulation program such as Hyperterminal, Teraterm, or PuTTY to access the console

The terminal program should be capable of logging screen output to a file.

Telnet

When using Telnet to configure and manage a bridge, your system must meet the following requirements:

- A serial cable (which connects from the bridge serial port to a serial (COM) port on the computer you are using for setup)

The bridge serial port is RJ-45 and has the same pin-out as the controllers.

- (Recommended) A non-default user name and password (for accessing the bridge)
- A terminal emulation program such as Hyperterminal, Teraterm, or PuTTY to access the console

The terminal program should be capable of logging screen output to a file.

- An IP address, subnet mask, and gateway information for the Ethernet management 1 port on each bridge

FTP

When using FTP to update bridge firmware, your system must meet the following requirements:

- A standard Ethernet cable (which connects from the bridge Ethernet management 1 port to your network)
- (Recommended) A non-default user name and password (for accessing the bridge)

Hot-replacing a failed power supply module

When there is a change in status of a power supply module to the bridge, you can remove and install the power supply module.

You can view the change in status of a power supply module through the LEDs on the bridge. You can also view the status of power supply modules via ExpressNAV GUI and the bridge CLI, via serial port, or via Telnet.

- This procedure is NDO (non-disruptive) and takes approximately 15 minutes to complete.
- You need the admin password and access to an FTP or SCP server.



The *ATTO FibreBridge Installation and Operation Manual* for your model bridge has more information about management interfaces.

You can access this and other content on the ATTO web site by using the link provided on the ATTO FibreBridge Description page.

In-band management of the FC-to-SAS bridges

Beginning with ONTAP 9.5 with FibreBridge 7500N or 7600N bridges, in-band management of the bridges is supported as an alternative to IP management of the bridges. Beginning with ONTAP 9.8, out-of-band management is deprecated.



About this task

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

When using in-band management, the bridges can be managed and monitored from the ONTAP CLI via the FC connection to the bridge. Physical access to the bridge via the bridge Ethernet ports is not required, reducing the security vulnerability of the bridge.

The availability of in-band management of the bridges depends on the version of ONTAP:

- Beginning with ONTAP 9.8, bridges are managed via in-band connections by default and out-of-band management of the bridges via SNMP is deprecated.
- ONTAP 9.5 through 9.7: Either in-band management or out-of-band SNMP management is supported.
- Prior to ONTAP 9.5, only out-of-band SNMP management is supported.

Bridge CLI commands can be issued from the ONTAP interface `storage bridge run-cli -name bridge-name -command bridge-command-name command` at the ONTAP interface.



Using in-band management with IP access disabled is recommended to improve security by limiting physical connectivity the bridge.

Related information

[Hot-swapping a bridge with a replacement bridge of the same model](#)

[Hot-swapping a FibreBridge 7500N with a 7600N bridge](#)

[Hot-swapping a FibreBridge 6500N bridge with a FibreBridge 7600N or 7500N bridge](#)

[Hot-adding a stack of SAS disk shelves and bridges](#)

Managing a FibreBridge bridge from ONTAP

Beginning with ONTAP 9.5, you can use the ONTAP CLI to pass FibreBridge commands to the bridge and display the results of those commands.

About this task



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

Steps

1. Run the applicable FibreBridge command within the `storage bridge run-cli` command:

```
storage bridge run-cli -name bridge-name -command "command-text"
```

The following command runs the FibreBridge `SASPortDisable` command from the ONTAP prompt to disable SAS port b on the bridge:

```
cluster_A::> storage bridge run-cli -name "SASPortDisable b"

SAS Port B has been disabled.
Ready
cluster_A::>
```

Securing or unsecuring the FibreBridge bridge

To easily disable potentially unsecure Ethernet protocols on a bridge, beginning with ONTAP 9.5 you can secure the bridge. This disables the bridge's Ethernet ports. You can also reenablen Ethernet access.

- Securing the bridge disables telnet and other IP port protocols and services (FTP, ExpressNAV, ICMP, or QuickNAV) on the bridge.
- This procedure uses out-of-band management using the ONTAP prompt, which is available beginning with ONTAP 9.5.

You can issue the commands from the bridge CLI if you are not using out-of-band management.

- The **`unsecurebridge`** command can be used to reenablen the Ethernet ports.
- In ONTAP 9.7 and earlier, running the **`securebridge`** command on the ATTO FibreBridge might not update the bridge status correctly on the partner cluster. If this occurs, run the **`securebridge`** command from the partner cluster.



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

Steps

1. From the ONTAP prompt of the cluster containing the bridge, secure or unsecure the bridge.

The following command secures bridge_A_1:

```
cluster_A> storage bridge run-cli -bridge bridge_A_1 -command
securebridge
```

The following command unsecures bridge_A_1:

```
cluster_A> storage bridge run-cli -bridge bridge_A_1 -command
unsecurebridge
```

2. From the ONTAP prompt of the cluster containing the bridge, save the bridge configuration:

storage bridge run-cli -bridge *bridge-name* -command saveconfiguration

The following command secures bridge_A_1:

```
cluster_A> storage bridge run-cli -bridge bridge_A_1 -command
saveconfiguration
```

3. From the ONTAP prompt of the cluster containing the bridge, restart the bridge's firmware:

storage bridge run-cli -bridge *bridge-name* -command firmwarerestart

The following command secures bridge_A_1:

```
cluster_A> storage bridge run-cli -bridge bridge_A_1 -command
firmwarerestart
```

FC switch maintenance and replacement

Upgrade or downgrade the firmware on a Brocade FC switch

To upgrade or downgrade the firmware on a Brocade FC switch, you must use the Brocade-specific commands to disable the switch, perform and verify the firmware change, and reboot and reenale the switch.

About this task

Confirm that you have checked and performed the following tasks for your configuration:

- Verify that your new firmware version is compatible with your solution. See the [Hardware Universe](#) for more information.
- You have the firmware files.

- The system is properly cabled.
- All paths to the storage shelves are available.
- The disk shelf stacks are stable.
- The FC switch fabric is healthy.
- No failed components are present in the system.
- The system is operating normally.
- You have the admin password and access to an FTP or SCP server.
- Console logging is enabled.

[Enable console logging](#)

The switch fabric is disabled during a firmware upgrade or downgrade, and the MetroCluster configuration relies on the second fabric to continue operation.

Beginning in Fabric OS 9.0.1, SNMPv2 is not supported on Brocade switches. If you upgrade to Fabric OS 9.0.1 or later, you must use SNMPv3 for health monitoring. For more information, see [Configuring SNMPv3 in a MetroCluster configuration](#).

If you are upgrading to Fabric OS v 9.2.x or later, you must have a Brocade TruFOS certificate installed, refer to [Brocade Fabric OS Software Upgrade Guide, 9.2.x](#) for more information.

This task must be performed on each of the switch fabrics in succession so that all switches are running the same firmware version.



This procedure is nondisruptive and takes approximately one hour to complete.

Steps

1. Log in to each of the switches in the fabric.

The examples in the following steps use the switch `FC_switch_A_1`.

2. Disable each of the switches in the fabric:

`switchCfgPersistentDisable`

If this command is not available, then run the `switchDisable` command.

```
FC_switch_A_1:admin> switchCfgPersistentDisable
```

3. Download the desired firmware version:

`firmwareDownload`

When prompted for the file name, you must specify the subdirectory or relative path to the firmware file.

You can run the `firmwareDownload` command at the same time on both switches, but you must allow the firmware to download and commit properly before moving to the next step.

```
FC_switch_A_1:admin> firmwaredownload
Server Name or IP Address: 10.64.203.188
User Name: test
File Name: v7.3.1b
Network Protocol(1-auto-select, 2-FTP, 3-SCP, 4-SFTP, 5-HTTP) [1]: 2
Password:
Server IP: 10.64.203.188, Protocol IPv4
Checking system settings for firmwaredownload...
System settings check passed.
```

4. Verify that the firmware was downloaded and committed to both partitions:

firmwareShow

The following example shows that the firmware download is complete as both images are updated:

```
FC_switch_A_1:admin> firmwareShow
Appl      Primary/Secondary Versions
-----
FOS       v7.3.1b
          v7.3.1b
```

5. Reboot the switches:

reboot

Some firmware versions automatically perform an haReboot operation after the firmware download is finished. The reboot in this step is required even if the haReboot has been performed.

```
FC_switch_A_1:admin> reboot
```

6. Check whether the new firmware is for an intermediate firmware level or for a final specified release.

If the download is for the intermediate firmware level, then perform the previous two steps until the specified release is installed.

7. Enable the switches:

switchCfgPersistentEnable

If this command is not available, then the switch should be in the `enabled` state after the `reboot` command is executed.

```
FC_switch_A_1:admin> switchCfgPersistentEnable
```

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

switchShow

```
FC_switch_A_1:admin> switchShow
```

9. Verify that the buffer usage information for a port group or all of the port groups in the switch is displayed properly:

portbuffershow

```
FC_switch_A_1:admin> portbuffershow
```

10. Verify that the current configuration of a port is displayed properly:

portcfgshow

```
FC_switch_A_1:admin> portcfgshow
```

Verify the port settings, such as speed, mode, trunking, encryption, and compression, in the Inter-Switch Link (ISL) output. Verify that the port settings were not affected by the firmware download.

11. 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.

12. Wait 15 minutes before repeating this procedure for the second switch fabric.

Upgrading or downgrading the firmware on a Cisco FC switch

To upgrade or downgrade the firmware on a Cisco FC switch you must use the Cisco-specific commands to disable the switch, perform and verify the upgrade, and reboot and reenable the switch.

About this task

Confirm that you have checked and performed the following tasks for your configuration:

- The system is properly cabled.
- All paths to the storage shelves are available.
- The disk shelf stacks are stable.
- The FC switch fabric are healthy.
- All components in the system are healthy.
- The system is operating normally.
- You have the admin password and access to an FTP or SCP server.
- Console logging is enabled.

[Enable console logging](#)

The switch fabric is disabled during the firmware upgrade or downgrade and the MetroCluster configuration relies on the second fabric to continue operation.

You must repeat this task on each of the switch fabrics in succession to ensure that all switches are running the same firmware version.

You must have the firmware files.



This procedure is nondisruptive and takes approximately one hour to complete.

Steps

1. Log in to each of the switches in the fabric.

In the examples, the switches are called FC_switch_A_1 and FC_switch_B_1.

2. Determine whether there is enough space in the bootflash directory on each switch:

```
dir bootflash
```

If not, delete the unwanted firmware files by using the `delete bootflash:file_name` command.

3. Copy the kickstart and system files to the switches:

```
copy source_filetarget_file
```

In the following example, the kickstart file (`m9200-s2ek9-kickstart-mz.5.2.1.bin`) and the system file (`m9200-s2ek9-mz.5.2.1.bin`) are located on the FTP server 10.10.10.55 in the `/firmware/` path.

The following example shows the commands issued on FC_switch_A_1:

```
FC_switch_A_1# copy ftp://10.10.10.55/firmware/m9200-s2ek9-kickstart-
mz.5.2.1.bin bootflash:m9200-s2ek9-kickstart-mz.5.2.1.bin
FC_switch_A_1# copy ftp://10.10.10.55/firmware/m9200-s2ek9-mz.5.2.1.bin
bootflash:m9200-s2ek9-mz.5.2.1.bin
```

4. Disable all of the VSANs on both of the switches in this fabric.

Use the following procedure to disable the VSANs:

- a. Open the config terminal:

```
config t
```

- b. Enter: **vsan database**

- c. Check the state of the VSANs:

```
show vsan
```

All VSANs must be active.

- d. Suspend the VSANs:

```
vsan vsan-num suspend
```

Example: vsan 10 suspend

- e. Check the state of the VSANs again:

```
show vsan
```

All VSANs must be suspended.

- f. Exit the config terminal:

```
end
```

- g. Save the configuration.

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

The following example displays the output for FC_switch_A_1:

```
FC_switch_A_1# config t
Enter configuration commands, one per line. End with CNTL/Z.
FC_switch_A_1(config)# vsan database
FC_switch_A_1(config-vsan-db)# show vsan
vsan 1 information
      name:VSAN0001  state:active
      interoperability mode:default
```

```

        loadbalancing:src-id/dst-id/oxid
        operational state:up

vsan 30 information
    name:MC1_FCVI_2_30  state:active
    interoperability mode:default
    loadbalancing:src-id/dst-id
    operational state:up

vsan 40 information
    name:MC1_STOR_2_40  state:active
    interoperability mode:default
    loadbalancing:src-id/dst-id/oxid
    operational state:up

vsan 70 information
    name:MC2_FCVI_2_70  state:active
    interoperability mode:default
    loadbalancing:src-id/dst-id
    operational state:up

vsan 80 information
    name:MC2_STOR_2_80  state:active
    interoperability mode:default
    loadbalancing:src-id/dst-id/oxid
    operational state:up

vsan 4079:evfp_isolated_vsan

vsan 4094:isolated_vsan

FC_switch_A_1(config-vsan-db)# vsan 1 suspend
FC_switch_A_1(config-vsan-db)# vsan 30 suspend
FC_switch_A_1(config-vsan-db)# vsan 40 suspend
FC_switch_A_1(config-vsan-db)# vsan 70 suspend
FC_switch_A_1(config-vsan-db)# vsan 80 suspend
FC_switch_A_1(config-vsan-db)# end
FC_switch_A_1#
FC_switch_A_1# show vsan
vsan 1 information
    name:VSAN0001  state:suspended
    interoperability mode:default
    loadbalancing:src-id/dst-id/oxid
    operational state:down

vsan 30 information

```

```

name:MC1_FCVI_2_30 state:suspended
interoperability mode:default
loadbalancing:src-id/dst-id
operational state:down

vsan 40 information
name:MC1_STOR_2_40 state:suspended
interoperability mode:default
loadbalancing:src-id/dst-id/oxid
operational state:down

vsan 70 information
name:MC2_FCVI_2_70 state:suspended
interoperability mode:default
loadbalancing:src-id/dst-id
operational state:down

vsan 80 information
name:MC2_STOR_2_80 state:suspended
interoperability mode:default
loadbalancing:src-id/dst-id/oxid
operational state:down

vsan 4079:evfp_isolated_vsan

vsan 4094:isolated_vsan

```

5. Install the desired firmware on the switches:

```

install all system bootflash:systemfile_name kickstart
bootflash:kickstartfile_name

```

The following example shows the commands issued on FC_switch_A_1:

```

FC_switch_A_1# install all system bootflash:m9200-s2ek9-mz.5.2.1.bin
kickstart bootflash:m9200-s2ek9-kickstart-mz.5.2.1.bin
Enter Yes to confirm the installation.

```

6. Check the version of the firmware on each switch to make sure the correct version was installed:

```

show version

```

7. Enable all of the VSANs on both of the switches in this fabric.

Use the following procedure to enable the VSANs:

a. Open the config terminal:

config t

b. Enter: **vsan database**

c. Check the state of the VSANs:

show vsan

The VSANs must be suspended.

d. Activate the VSANs:

no vsan vsan-num suspend

Example: no vsan 10 suspend

e. Check the state of the VSANs again:

show vsan

All VSANs must be active.

f. Exit the config terminal:

end

g. Save the configuration:

copy running-config startup-config

The following example displays the output for FC_switch_A_1:

```
FC_switch_A_1# config t
Enter configuration commands, one per line.  End with CNTL/Z.
FC_switch_A_1(config)# vsan database
FC_switch_A_1(config-vsan-db)# show vsan
vsan 1 information
    name:VSAN0001  state:suspended
    interoperability mode:default
    loadbalancing:src-id/dst-id/oxid
    operational state:down

vsan 30 information
    name:MC1_FCVI_2_30  state:suspended
    interoperability mode:default
    loadbalancing:src-id/dst-id
    operational state:down

vsan 40 information
    name:MC1_STOR_2_40  state:suspended
    interoperability mode:default
```

```
loadbalancing:src-id/dst-id/oxid
operational state:down

vsan 70 information
  name:MC2_FCVI_2_70  state:suspended
  interoperability mode:default
  loadbalancing:src-id/dst-id
  operational state:down

vsan 80 information
  name:MC2_STOR_2_80  state:suspended
  interoperability mode:default
  loadbalancing:src-id/dst-id/oxid
  operational state:down

vsan 4079:evfp_isolated_vsan

vsan 4094:isolated_vsan

FC_switch_A_1(config-vsan-db)# no vsan 1 suspend
FC_switch_A_1(config-vsan-db)# no vsan 30 suspend
FC_switch_A_1(config-vsan-db)# no vsan 40 suspend
FC_switch_A_1(config-vsan-db)# no vsan 70 suspend
FC_switch_A_1(config-vsan-db)# no vsan 80 suspend
FC_switch_A_1(config-vsan-db)#
FC_switch_A_1(config-vsan-db)# show vsan
vsan 1 information
  name:VSAN0001  state:active
  interoperability mode:default
  loadbalancing:src-id/dst-id/oxid
  operational state:up

vsan 30 information
  name:MC1_FCVI_2_30  state:active
  interoperability mode:default
  loadbalancing:src-id/dst-id
  operational state:up

vsan 40 information
  name:MC1_STOR_2_40  state:active
  interoperability mode:default
  loadbalancing:src-id/dst-id/oxid
  operational state:up

vsan 70 information
  name:MC2_FCVI_2_70  state:active
```

```

interoperability mode:default
loadbalancing:src-id/dst-id
operational state:up

vsan 80 information
  name:MC2_STOR_2_80 state:active
  interoperability mode:default
  loadbalancing:src-id/dst-id/oxid
  operational state:up

vsan 4079:evfp_isolated_vsan

vsan 4094:isolated_vsan

FC_switch_A_1(config-vsan-db)# end
FC_switch_A_1#

```

8. 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.

9. Repeat this procedure for the second switch fabric.

Upgrading to new Brocade FC switches

If you are upgrading to new Brocade FC switches, you must replace the switches in the first fabric, verify that the MetroCluster configuration is fully operational, and then replace the switches in the second fabric.

- The MetroCluster configuration must be healthy and in normal operation.
- The MetroCluster switch fabrics consist of four Brocade switches.

The illustrations in the following steps show current switches.


- The switches must be running the most recent supported firmware.

[NetApp Interoperability Matrix Tool](#)

- This procedure is nondisruptive and takes approximately two hours to complete.
- You need the admin password and access to an FTP or SCP server.
- [Enable console logging](#) before performing this task.

The switch fabrics are upgraded one at a time.

At the end of this procedure, all four switches will be upgraded to new switches.




Steps


1. Disable the first switch fabric:

```
FC_switch_A_1:admin> switchCfgPersistentDisable
```

```
FC_switch_A_1:admin> switchCfgPersistentDisable
```



2. Replace the old switches at one MetroCluster site.
 - a. Uncable and remove the disabled switch.
 - b. Install the new switch in the rack.



- c. Disable the new switches by running the following command on both switches:

```
switchCfgPersistentDisable
```

```
FC_switch_A_1:admin> switchCfgPersistentDisable
```


- d. Cable the new switch using the recommended port assignments.

[Port assignments for FC switches](#)

- e. Repeat these substeps at the partner MetroCluster site to replace the second switch in the first switch

fabric.

Both switches in fabric 1 have been replaced.



3. Power up the new switches and let them boot up.
4. Configure the Brocade FC switches using one of the following procedures:

[Configure Brocade FC switches with RCF files](#)

[Configure the Brocade FC switches manually](#)

5. Save the switch configuration:

```
cfgSave
```

6. Wait 10 minutes to allow the configuration to stabilize.
7. Confirm connectivity to the disks by entering the following command on any one of the MetroCluster nodes:

```
run local sysconfig -v
```

The output shows the disks attached to the initiator ports on the controller, and identifies the shelves connected to the FC-to-SAS bridges:

```
node_A_1> run local sysconfig -v
NetApp Release 9.3.2X18: Sun Dec 13 01:23:24 PST 2017
System ID: 4068741258 (node_A_1); partner ID: 4068741260 (node_B_1)
System Serial Number: 940001025471 (node_A_1)
System Rev: 70
System Storage Configuration: Multi-Path HA**<=== Configuration should
be multi-path HA**
.
.
.
```

```

slot 0: FC Host Adapter 0g (QLogic 8324 rev. 2, N-port, <UP>)**<===
Initiator port**
    Firmware rev:      7.5.0
    Flash rev:         0.0.0
    Host Port Id:      0x60130
    FC Node Name:      5:00a:098201:bae312
    FC Port Name:      5:00a:098201:bae312
    SFP Vendor:        UTILITIES CORP.
    SFP Part Number:   FTLF8529P3BCVAN1
    SFP Serial Number: URQ0Q9R
    SFP Capabilities:  4, 8 or 16 Gbit
    Link Data Rate:    16 Gbit
    Switch Port:       brcd6505-fcs40:1
**<List of disks visible to port\>**
    ID      Vendor  Model          FW      Size
    brcd6505-fcs29:12.126L1527  : NETAPP  X302_HJUPI01TSSM NA04
847.5GB (1953525168 512B/sect)
    brcd6505-fcs29:12.126L1528  : NETAPP  X302_HJUPI01TSSA NA02
847.5GB (1953525168 512B/sect)
    .
    .
    .
**<List of FC-to-SAS bridges visible to port\>**
FC-to-SAS Bridge:
    brcd6505-fcs40:12.126L0      : ATTO      FibreBridge6500N 1.61
FB6500N102980
    brcd6505-fcs42:13.126L0      : ATTO      FibreBridge6500N 1.61
FB6500N102980
    brcd6505-fcs42:6.126L0       : ATTO      FibreBridge6500N 1.61
FB6500N101167
    brcd6505-fcs42:7.126L0       : ATTO      FibreBridge6500N 1.61
FB6500N102974
    .
    .
    .
**<List of storage shelves visible to port\>**
    brcd6505-fcs40:12.shelf6: DS4243  Firmware rev. IOM3 A: 0200
IOM3 B: 0200
    brcd6505-fcs40:12.shelf8: DS4243  Firmware rev. IOM3 A: 0200
IOM3 B: 0200
    .
    .
    .

```

8. Returning to the switch prompt, verify the switch firmware version:

```
firmwareShow
```

The switches must be running the most recent supported firmware.

[NetApp Interoperability Matrix Tool](#)

9. Simulate a switchover operation:

- a. From any node's prompt, change to the advanced privilege level:

```
set -privilege advanced
```

You need to respond with "y" when prompted to continue into advanced mode and see the advanced mode prompt (*>).

- b. Perform the switchover operation with the `-simulate` parameter:


```
metrocluster switchover -simulate
```

- c. Return to the admin privilege level:

```
set -privilege admin
```

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

After repeating the steps, all four switches have been upgraded and the MetroCluster configuration is in normal operation.



Replacing a Brocade FC switch

You must use this Brocade-specific procedure to replace a failed switch.

About this task

You need the admin password and access to an FTP or SCP server.

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

In the following examples, FC_switch_A_1 is the healthy switch and FC_switch_B_1 is the impaired switch. The switch port usage in the examples is shown in the following table:

Port connections	Ports
FC-VI connections	0, 3
HBA connections	1, 2, 4, 5
FC-to-SAS bridge connections	6, 7
ISL connections	10, 11

The examples show two FC-to-SAS bridges. If you have more, you must disable and subsequently enable the additional ports.



This procedure is nondisruptive and takes approximately two hours to complete.

Your switch port usage should follow the recommended assignments.

- [Port assignments for FC switches](#)

Steps

1. Fence off the switch undergoing replacement by disabling the ISL ports on the healthy switch in the fabric and the FC-VI and HBA ports on the impaired switch (if the impaired switch is still operating):
 - a. Disable the ISL ports on the healthy switch for each port:

```
portcfgpersistentdisable port-number
```

```
FC_switch_A_1:admin> portcfgpersistentdisable 10  
FC_switch_A_1:admin> portcfgpersistentdisable 11
```

- b. If the impaired switch is still operational, disable the FC-VI and HBA ports on that switch for each port:

```
portcfgpersistentdisable port-number
```

```
FC_switch_B_1:admin> portcfgpersistentdisable 0  
FC_switch_B_1:admin> portcfgpersistentdisable 1  
FC_switch_B_1:admin> portcfgpersistentdisable 2  
FC_switch_B_1:admin> portcfgpersistentdisable 3  
FC_switch_B_1:admin> portcfgpersistentdisable 4  
FC_switch_B_1:admin> portcfgpersistentdisable 5
```

2. If the impaired switch is still operational, gather the output from the `switchshow` command.

```
FC_switch_B_1:admin> switchshow
  switchName: FC_switch_B_1
  switchType: 71.2
  switchState:Online
  switchMode: Native
  switchRole: Subordinate
  switchDomain:      2
  switchId:   fffc01
  switchWwn:  10:00:00:05:33:86:89:cb
  zoning:    OFF
  switchBeacon: OFF
```

3. Boot and preconfigure the new switch prior to physically installing it:

- a. Power up the new switch and let it boot up.
- b. Check the firmware version on the switch to confirm that it matches the version of the other FC switches:

```
firmwareShow
```

- c. Configure the new switch by following the Brocade procedures in [Configure the FC switches](#).



At this point, the new switch is not cabled to the MetroCluster configuration.

- d. Disable the FC-VI, HBA, and storage ports on the new switch, and the ports connected to the FC-SAS bridges.

```
FC_switch_B_1:admin> portcfgpersistentdisable 0
FC_switch_B_1:admin> portcfgpersistentdisable 1
FC_switch_B_1:admin> portcfgpersistentdisable 2
FC_switch_B_1:admin> portcfgpersistentdisable 3
FC_switch_B_1:admin> portcfgpersistentdisable 4
FC_switch_B_1:admin> portcfgpersistentdisable 5

FC_switch_B_1:admin> portcfgpersistentdisable 6
FC_switch_B_1:admin> portcfgpersistentdisable 7
```

4. Physically replace the switch:

- a. Power off the impaired FC switch.
- b. Power off the replacement FC switch.
- c. Uncable and remove the impaired switch, carefully noting which cables connected to which ports.
- d. Install the replacement switch in the rack.
- e. Cable the replacement switch exactly as the old switch was cabled.
- f. Power on the new FC switch.

5. To enable ISL encryption, refer to [Configure the Brocade FC switches manually](#).

If you are enabling ISL encryption, you need to complete the following tasks:

- Disable the virtual fabric
- Set the payload
- Set the authentication policy
- Enable ISL encryption on Brocade switches

6. Complete the configuration of the new switch:

a. Enable the ISLs:

```
portcfgpersistentenable port-number
```

```
FC_switch_B_1:admin> portcfgpersistentenable 10  
FC_switch_B_1:admin> portcfgpersistentenable 11
```

b. Verify the zoning configuration:

```
cfg show
```

c. On the replacement switch (FC_switch_B_1 in the example), verify that the ISLs are online:

```
switchshow
```

```
FC_switch_B_1:admin> switchshow  
switchName: FC_switch_B_1  
switchType: 71.2  
switchState:Online  
switchMode: Native  
switchRole: Principal  
switchDomain:      4  
switchId:   fffc03  
switchWwn:  10:00:00:05:33:8c:2e:9a  
zoning:      OFF  
switchBeacon: OFF  
  
Index Port Address Media Speed State Proto  
=====
```

...							
10	10	030A00 id	16G	Online	FC E-Port		
		10:00:00:05:33:86:89:cb	"FC_switch_A_1"				
11	11	030B00 id	16G	Online	FC E-Port		
		10:00:00:05:33:86:89:cb	"FC_switch_A_1"	(downstream)			
...							

- d. Enable the storage ports that connect to the FC bridges.

```
FC_switch_B_1:admin> portcfgpersistentenable 6
FC_switch_B_1:admin> portcfgpersistentenable 7
```

- e. Enable the storage, HBA, and FC-VI ports.

The following example shows the commands used to enable the ports connecting HBA adapters:

```
FC_switch_B_1:admin> portcfgpersistentenable 1
FC_switch_B_1:admin> portcfgpersistentenable 2
FC_switch_B_1:admin> portcfgpersistentenable 4
FC_switch_B_1:admin> portcfgpersistentenable 5
```

The following example shows the commands used to enable the ports connecting the FC-VI adapters:

```
FC_switch_B_1:admin> portcfgpersistentenable 0
FC_switch_B_1:admin> portcfgpersistentenable 3
```

7. Verify that the ports are online:

```
switchshow
```

8. 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](#).

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

Renaming a Brocade FC switch

You might need to rename a Brocade FC switch to ensure consistent naming throughout your configuration.

About this task

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

Steps

1. Persistently disable the switch or switches in one fabric:

```
switchcfgpersistentdisable
```

The following example shows the output for the **switchcfgpersistentdisable** command:

```
7840_FCIP_2:admin> switchcfgpersistentdisable
Switch's persistent state set to 'disabled'
2018/03/09-07:41:06, [ESM-2105], 146080, FID 128, INFO, 7840_FCIP_2, VE
Tunnel 24 is DEGRADED.
2018/03/09-07:41:06, [ESM-2104], 146081, FID 128, INFO, 7840_FCIP_2, VE
Tunnel 24 is OFFLINE.

7840_FCIP_2:admin>
```

2. Rename the switch or switches:

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

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

The following example shows the output for the **switchname new-switch-name** command:

```
7840_FCIP_2:admin> switchname FC_switch_1_B
Committing configuration...
Done.
Switch name has been changed.Please re-login into the switch for the
change to be applied.
2018/03/09-07:41:20, [IPAD-1002], 146082, FID 128, INFO, FC_switch_1_B,
Switch name has been successfully changed to FC_switch_1_B.
7840_FCIP_2:admin>
```

3. Reboot the switch or switches:

reboot

If you are renaming both switches in the fabric, reboot both switches. Once the reboot is complete, the switch is renamed in all places.

The following example shows the output for the **reboot** command:

```
7840_FCIP_2:admin> reboot
Warning: This command would cause the switch to reboot
and result in traffic disruption.
Are you sure you want to reboot the switch [y/n]?y
2018/03/09-07:42:08, [RAS-1007], 146083, CHASSIS, INFO, Brocade7840,
System is about to reload.
Rebooting! Fri Mar  9 07:42:11 CET 2018

Broadcast message from root (ttyS0) Fri Mar  9 07:42:11 2018...

The system is going down for reboot NOW !!
INIT: Switching to runlevel: 6
INIT:
2018/03/09-07:50:48, [ESM-1013], 146104, FID 128, INFO, FC_switch_1_B,
DP0 Configuration replay has completed.
2018/03/09-07:50:48, [ESM-1011], 146105, FID 128, INFO, FC_switch_1_B,
DP0 is ONLINE.

*** CORE FILES WARNING (03/09/18 - 08:00:00 ) ***
10248 KBytes in 1 file(s)
use "supportsave" command to upload

*** FFDC FILES WARNING (03/09/18 - 08:00:00 ) ***
520 KBytes in 1 file(s)
```

4. Persistently enable the switches: **switchcfgpersistentenable**

The following example shows the output for the **switchcfgpersistentenable** command:

```

FC_switch_1_B:admin> switchcfgpersistentenable
Switch's persistent state set to 'enabled'
FC_switch_1_B:admin>
FC_switch_1_B:admin>
FC_switch_1_B:admin> 2018/03/09-08:07:07, [ESM-2105], 146106, FID 128,
INFO, FC_switch_1_B, VE Tunnel 24 is DEGRADED.
2018/03/09-08:07:10, [ESM-2106], 146107, FID 128, INFO, FC_switch_1_B,
VE Tunnel 24 is ONLINE.

FC_switch_1_B:admin>

```

```

FC_switch_1_B:admin> switchshow
switchName:      FC_switch_1_B
switchType:      148.0
switchState:     Online
switchMode:      Native
switchRole:      Subordinate
switchDomain:     6
switchId:        fffc06
switchWwn:       10:00:50:eb:1a:9a:a5:79
zoning:          ON (CFG_FAB_2_RCF_9_3)
switchBeacon:    OFF
FC Router:       OFF
FC Router BB Fabric ID: 128
Address Mode:    0
HIF Mode:        OFF

```

Index	Port	Address	Media	Speed	State	Proto
0	0	060000	id	16G	Online	FC F-Port
		50:0a:09:81:06:a5:5a:08				
1	1	060100	id	16G	Online	FC F-Port
		50:0a:09:83:06:a5:5a:08				

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

storage switch show

The following example shows the output for the **storage switch show** command:

```

cluster_A::*> storage switch show
(storage switch show)
          Symbolic                               Is
Monitor
Switch      Name      Vendor  Model  Switch  WWN          Monitored
Status
-----
-----
Brocade_172.20.7.90
          RTP-FC01-510Q40
          Brocade Brocade7840
          1000c4f57c904bc8 true
ok
Brocade_172.20.7.91
          RTP-FC02-510Q40
          Brocade Brocade7840
          100050eb1a9aa579 true
ok
Brocade_172.20.7.92

```

Disabling encryption on Brocade FC switches

You might need to disable encryption on Brocade FC switches.

Steps

1. Send an AutoSupport message from both sites indicating the beginning of maintenance.

```
cluster_A::> autosupport invoke -node * -type all -message MAINT=4h
```

```
cluster_B::> autosupport invoke -node * -type all -message MAINT=4h
```

2. Verify the operation of the MetroCluster configuration from Cluster A.

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

metrocluster show

```
cluster_A::> metrocluster show
```

- b. Perform a MetroCluster check:

metrocluster check run

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

- c. Display the results of the MetroCluster check:

metrocluster check show

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

3. Check the status of both switches:

fabric show

```
switch_A_1:admin> fabric show
```

```
switch_B_1:admin> fabric show
```

4. Disable both switches:

switchdisable

```
switch_A_1:admin> switchdisable
```

```
switch_B_1:admin> switchdisable
```

5. Check the available paths for the nodes on each cluster:

sysconfig

```
cluster_A::> system node run -node node-name -command sysconfig -a
```

```
cluster_B::> system node run -node node-name -command sysconfig -a
```

As the switch fabric is now disabled, the System Storage Configuration should be Single-Path HA.

6. Check the aggregate status for both clusters.

```
cluster_A::> aggr status
```

```
cluster_B::> aggr status
```

System output should show the aggregates are mirrored and normal for both clusters:

```
mirrored,normal
```

7. Repeat the following substeps from the admin prompt on both switches.

a. Show which ports are encrypted:

portenccompshow

```
switch_A_1:admin> portenccompshow
```

b. Disable encryption on the encrypted ports:

portcfgencrypt - disable port-number

```
switch_A_1:admin> portcfgencrypt --disable 40  
switch_A_1:admin> portcfgencrypt --disable 41  
switch_A_1:admin> portcfgencrypt --disable 42  
switch_A_1:admin> portcfgencrypt --disable 43
```

c. Set the authentication type to all:

authUtil --set -a all

```
switch_A_1:admin> authUtil --set -a all
```

d. Set the authentication policy on the switch. to off:

authutil --policy -sw off

```
switch_A_1:admin> authutil --policy -sw off
```

e. Set the authentication Diffie-Hellman group to * :

authutil --set -g *

```
switch_A_1:admin> authUtil --set -g *
```

f. Delete the secret key database:

secAuthSecret --remove -all

```
switch_A_1:admin> secAuthSecret --remove -all
```

g. Confirm that encryption is disabled on the ports:

portenccompshow

```
switch_A_1:admin> portenccompshow
```

h. Enable the switch:

switchenable

```
switch_A_1:admin> switchenable
```

i. Confirm the status of the ISLs:

islshow

```
switch_A_1:admin> islshow
```

8. Check the available paths for the nodes on each cluster:

sysconfig

```
cluster_A::> system node run -node * -command sysconfig -a
```

```
cluster_B::> system node run -node * -command sysconfig -a
```

The system output should indicate that System Storage Configuration has changed back to Quad-Path HA.

9. Check the aggregate status for both clusters.

```
cluster_A::> aggr status
```

```
cluster_B::> aggr status
```

The system should show that the aggregates are mirrored and normal for both clusters as shown in the following system output:

```
mirrored,normal
```

10. Verify the operation of the MetroCluster configuration from Cluster A.

a. Perform a MetroCluster check:

metrocluster check run

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

b. Display the results of the MetroCluster check:

metrocluster check show

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

11. Send an AutoSupport message from both sites indicating the end of maintenance.

```
cluster_A::> autosupport invoke -node node-name -type all -message  
MAINT=END
```

```
cluster_B::> autosupport invoke -node node-name -type all -message  
MAINT=END
```

Change ISL properties, ISL ports, or the IOD/OOD configuration on a Brocade switch

You might need to add ISLs to a switch if you are adding or upgrading hardware such as additional or faster controllers or switches.

Before you begin

Ensure that the system is properly configured, that all fabric switches are operational, and that no errors exist.

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

If the equipment on the ISL link changes and the new link configuration no longer supports the current configuration---trunking and ordered delivery---then the fabric needs to be reconfigured for the correct routing policy: either in-order-deliver (IOD) or out-of-order-delivery (OOD).



To make changes to OOD from ONTAP software, use the following steps: [Configuring in-order delivery or out-of-order delivery of frames on ONTAP software](#)

Steps

1. Disable the FCVI and storage HBA ports:

```
portcfgpersistentdisable port number
```

By default the first 8 ports (ports 0 through 7) are used for FCVI and Storage HBA. The ports must be persistently disabled so that the ports remain disabled in the event of a switch reboot.

The following example shows ISL ports 0—7 being disabled on both switches:

```
Switch_A_1:admin> portcfgpersistentdisable 0-7
Switch_B_1:admin> portcfgpersistentdisable 0-7
```

2. Change the ISL ports as required.

Option	Step
To change the speed of an ISL port...	<p>Use the <code>portcfgspeed port number port speed</code> command on both switches on the fabric.</p> <p>In the following example, you change the ISL port speed from 40 Gbps to 16 Gbps:</p> <pre>brocade_switch_A_1:admin> portcfgspeed 40 16</pre> <p>You can verify that the speed has changed using the <code>switchshow</code> command:</p> <pre>brocade_switch_A_1:admin> switchshow</pre> <p>You should see the following output:</p> <pre> . . . 40 40 062800 id 16G No_Sync FC Disabled . . .</pre>
To change the distance of an ISL port...	Use the <code>portcfglongdistance port number port distance</code> command on both switches in the fabric.
To remove an ISL...	Disconnect the link.
To add an ISL...	Insert SFPs into the ports you are adding as ISL ports. Ensure that these ports are listed in Install a fabric-attached MetroCluster for the switch to which you are adding them.
To relocate an ISL...	Relocating an ISL is the same as removing and then adding an ISL. First, remove the ISL by disconnecting the link and then insert SFPs into the ports you are adding as ISL ports.



When you make changes to ISL ports you might also need to apply additional settings recommended by the WDM vendor. Refer to the WDM vendor documentation for guidance.

3. Reconfigure for out-of-order delivery (OOD) or in-order-delivery (IOD).



If the routing policies remain the same, you do not need to reconfigure and this step can be ignored. The ONTAP configuration needs to match the fabric configuration. If the fabric is configured for OOD, then ONTAP must also be configured for OOD. The same applies for IOD.

This step should be executed in the following scenarios:

- More than one ISL formed a trunk before the change, but after the change, trunking is no longer supported. In this case, you must configure the fabric for OOD.
- There is one ISL before the change and multiple ISLs after the change.
- If multiple ISLs form a trunk, configure the fabric for IOD. If multiple ISLs **cannot** form a trunk, configure the fabric for OOD.
- Persistently disable the switches using the `switchcfgpersistentdisable` command as shown in the following example:

```
Switch_A_1:admin> switchcfgpersistentdisable
Switch_B_1:admin> switchcfgpersistentdisable
```

- Configure the trunking mode for each ISL `portcfgtrunkport port number` as shown in the following table:

Scenario	Steps
Configure the ISL for trunking \((IOD\)	<p>Set the <code>portcfgtrunkport port number</code> to 1:</p> <pre>FC_switch_A_1:admin> portcfgtrunkport 20 1 FC_switch_A_1:admin> portcfgtrunkport 21 1 FC_switch_B_1:admin> portcfgtrunkport 20 1 FC_switch_B_1:admin> portcfgtrunkport 21 1</pre>
Configure the ISL for trunking \((OOD\)	<p>Set the <code>portcfgtrunkport port number</code> to 0:</p> <pre>FC_switch_A_1:admin> portcfgtrunkport 20 0 FC_switch_A_1:admin> portcfgtrunkport 21 0 FC_switch_B_1:admin> portcfgtrunkport 20 0 FC_switch_B_1:admin> portcfgtrunkport 21 0</pre>

- Configure the fabric for IOD or OOD as required.

Scenario	Steps
----------	-------

Configure the fabric for IOD	<p>Set the three settings of IOD, APT, and DLS using the <code>iodset</code>, <code>aptpolicy</code>, and <code>dlsreset</code> commands as shown in the following example:</p> <pre> Switch_A_1:admin> iodset Switch_A_1:admin> aptpolicy 1 Policy updated successfully. Switch_A_1:admin> dlsreset FC_switch_A_1:admin>portcfgtrunkport 40 1 FC_switch_A_1:admin>portcfgtrunkport 41 1 Switch_B_1:admin> iodset Switch_B_1:admin> aptpolicy 1 Policy updated successfully. Switch_B_1:admin> dlsreset FC_switch_B_1:admin>portcfgtrunkport 20 1 FC_switch_B_1:admin>portcfgtrunkport 21 1 </pre>
Configure the fabric for OOD	<p>Set the three settings of IOD, APT, and DLS using the <code>iodreset</code>, <code>aptpolicy</code>, and <code>dlset</code> commands as shown in the following example:</p> <pre> Switch_A_1:admin> iodreset Switch_A_1:admin> aptpolicy 3 Policy updated successfully. Switch_A_1:admin> dlset FC_switch_A_1:admin> portcfgtrunkport 40 0 FC_switch_A_1:admin> portcfgtrunkport 41 0 Switch_B_1:admin> iodreset Switch_B_1:admin> aptpolicy 3 Policy updated successfully. Switch_B_1:admin> dlset FC_switch_B_1:admin> portcfgtrunkport 40 0 FC_switch_B_1:admin> portcfgtrunkport 41 0 </pre>

c. Enable the switches persistently:

```
switchcfgpersistentenable
```

```

switch_A_1:admin>switchcfgpersistentenable
switch_B_1:admin>switchcfgpersistentenable

```

If this command does not exist, use the `switchenable` command as shown in the following example:

```
brocade_switch_A_1:admin>  
switchenable
```

- d. Verify the OOD settings using the `iodshow`, `aptpolicy`, and `dlsshow` commands as shown in the following example:

```
switch_A_1:admin> iodshow  
IOD is not set  
  
switch_A_1:admin> aptpolicy  
  
Current Policy: 3 0(ap)  
  
3 0(ap) : Default Policy  
1: Port Based Routing Policy  
3: Exchange Based Routing Policy  
0: AP Shared Link Policy  
1: AP Dedicated Link Policy  
command aptpolicy completed  
  
switch_A_1:admin> dlsshow  
DLS is set by default with current routing policy
```



You must run these commands on both switches.

- e. Verify the IOD settings using the `iodshow`, `aptpolicy`, and `dlsshow` commands as shown in the following example:

```

switch_A_1:admin> iodshow
IOD is set

switch_A_1:admin> aptpolicy
Current Policy: 1 0(ap)

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

switch_A_1:admin> dlsshow
DLS is not set

```



You must run these commands on both switches.

- Verify that the ISLs are online and trunked (if the linking equipment supports trunking) using the `islshow` and `trunkshow` commands.



If FEC is enabled, the deskew value of the last online port of the trunk group might show a difference of up to 36 although the cables are all of the same length.

Are ISLs trunked?	You see the following system output...
Yes	<p>If the ISLs are trunked, only a single ISL appears in the output for the <code>islshow</code> command. Either port 40 or 41 can appear depending on which is the trunk master. The output of <code>trunkshow</code> should one trunk with ID "1" listing both the physical ISLs on ports 40 and 41. In the following example the ports 40 and 41 are configured for use as an ISL:</p> <pre> switch_A_1:admin> islshow 1: 40-> 40 10:00:00:05:33:88:9c:68 2 switch_B_1 sp: 16.000G bw: 32.000G TRUNK CR_RECOV FEC switch_A_1:admin> trunkshow 1: 40-> 40 10:00:00:05:33:88:9c:68 2 deskew 51 MASTER 41-> 41 10:00:00:05:33:88:9c:68 2 deskew 15 </pre>

No

If the ISLs are not trunked, both ISLs appear separately in the outputs for `islshow` and `trunkshow`. Both commands list the ISLs with their ID of “1” and “2”. In the following example, the ports “40” and “41” are configured for use as an ISL:

```
switch_A_1:admin> islshow
1: 40-> 40 10:00:00:05:33:88:9c:68 2 switch_B_1 sp:
16.000G bw: 16.000G TRUNK CR_RECOV FEC
2: 41-> 41 10:00:00:05:33:88:9c:68 2 switch_B_1 sp:
16.000G bw: 16.000G TRUNK CR_RECOV FEC
switch_A_1:admin> trunkshow
1: 40-> 40 10:00:00:05:33:88:9c:68 2 deskew 51 MASTER
2: 41-> 41 10:00:00:05:33:88:9c:68 2 deskew 48 MASTER
```

5. Run the `spinfab` command on both switches to verify that the ISLs are healthy:

```
switch_A_1:admin> spinfab -ports 0/40 - 0/41
```

6. Enable the ports that were disabled in step 1:

```
portenable port number
```

The following example shows ISL ports “0” through “7” being enabled:

```
brocade_switch_A_1:admin> portenable 0-7
```

Replacing a Cisco FC switch

You must use Cisco-specific steps to replace a failed Cisco FC switch.

Before you begin

You need the admin password and access to an FTP or SCP server.

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

About this task

This procedure is nondisruptive and takes approximately two hours to complete.

In the examples in this procedure, `FC_switch_A_1` is the healthy switch and `FC_switch_B_1` is the impaired switch. The switch port usage in the examples is shown in the following table:

Role	Ports
FC-VI connections	1, 4

HBA connections	2, 3, 5, 6
FC-to-SAS bridge connections	7, 8
ISL connections	36, 40

The examples show two FC-to-SAS bridges. If you have more, you must disable and subsequently enable the additional ports.

Your switch port usage should follow the recommended assignments.

- [Port assignments for FC switches](#)

Steps

1. Disable the ISL ports on the healthy switch to fence off the impaired switch.

These steps are performed on the healthy switch.

- a. Enter configuration mode:

```
conf t
```

- b. Disable the ISL ports on the healthy switch with the `interface` and `shut` commands.

```
FC_switch_A_1# conf t
FC_switch_A_1(config)# interface fc1/36
FC_switch_A_1(config)# shut
FC_switch_A_1(config)# interface fc1/40
FC_switch_A_1(config)# shut
```

- c. Exit configuration mode and copy the configuration to the startup configuration.

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

2. Fence off the FC-VI and HBA ports on the impaired switch (if it is still running).

These steps are performed on the impaired switch.

- a. Enter configuration mode:

```
conf t
```

- b. If the impaired switch is still operational, disable the FC-VI and HBA ports on the impaired switch with the `interface` and `shut` commands.

```
FC_switch_B_1(config)# interface fc1/1
FC_switch_B_1(config)# shut
FC_switch_B_1(config)# interface fc1/4
FC_switch_B_1(config)# shut
FC_switch_B_1(config)# interface fc1/2-3
FC_switch_B_1(config)# shut
FC_switch_B_1(config)# interface fc1/5-6
FC_switch_B_1(config)# shut
```

- c. Exit configuration mode and copy the configuration to the startup configuration.

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

3. If the impaired switch is still operational, determine the WWN for the switch:

```
show wwn switch
```

```
FC_switch_B_1# show wwn switch
Switch WWN is 20:00:54:7f:ee:e3:86:50
FC_switch_B_1#
```

4. Boot and preconfigure the replacement switch, prior to physically installing it.

At this point the replacement switch is not cabled to the MetroCluster configuration. The ISL ports on the partner switch are disabled (in shut mode) and offline.

- a. Power on the replacement switch and let it boot up.
- b. Check the firmware version on the replacement switch to confirm that it matches the version of the other FC switches:

```
show version
```

- c. Configure the replacement switch as described in the *MetroCluster Installation and Configuration Guide*, skipping the “Configuring zoning on a Cisco FC switch” section.

[Fabric-attached MetroCluster installation and configuration](#)

You will configure zoning later in this procedure.

- d. Disable the FC-VI, HBA, and storage ports on the replacement switch.

```

FC_switch_B_1# conf t
FC_switch_B_1(config)# interface fc1/1
FC_switch_B_1(config)# shut
FC_switch_B_1(config)# interface fc1/4
FC_switch_B_1(config)# shut
FC_switch_B_1(config)# interface fc1/2-3
FC_switch_B_1(config)# shut
FC_switch_B_1(config)# interface fc1/5-6
FC_switch_B_1(config)# shut
FC_switch_B_1(config)# interface fc1/7-8
FC_switch_B_1(config)# shut
FC_switch_B_1# copy running-config startup-config
FC_switch_B_1#

```

5. Physically replace the impaired switch:

- a. Power off the impaired switch.
- b. Power off the replacement switch.
- c. Uncable and remove the impaired switch, carefully noting which cables connected to which ports.
- d. Install the replacement switch in the rack.
- e. Cable the replacement switch exactly as the impaired switch was cabled.
- f. Power on the replacement switch.

6. Enable the ISL ports on the replacement switch.

```

FC_switch_B_1# conf t
FC_switch_B_1(config)# interface fc1/36
FC_switch_B_1(config)# no shut
FC_switch_B_1(config)# end
FC_switch_B_1# copy running-config startup-config
FC_switch_B_1(config)# interface fc1/40
FC_switch_B_1(config)# no shut
FC_switch_B_1(config)# end
FC_switch_B_1#

```

7. Verify that the ISL ports on the replacement switch are up:

```
show interface brief
```

8. Adjust the zoning on the replacement switch to match the MetroCluster configuration:

- a. Distribute the zoning information from the healthy fabric.

In this example, FC_switch_B_1 has been replaced and the zoning information is retrieved from FC_switch_A_1:

```
FC_switch_A_1(config-zone)# zoneset distribute full vsan 10
FC_switch_A_1(config-zone)# zoneset distribute full vsan 20
FC_switch_A_1(config-zone)# end
```

- b. On the replacement switch, verify that the zoning information was properly retrieved from the healthy switch:

```
show zone
```

```
FC_switch_B_1# show zone
zone name FC-VI_Zone_1_10 vsan 10
  interface fc1/1 swwn 20:00:54:7f:ee:e3:86:50
  interface fc1/4 swwn 20:00:54:7f:ee:e3:86:50
  interface fc1/1 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/4 swwn 20:00:54:7f:ee:b8:24:c0

zone name STOR_Zone_1_20_25A vsan 20
  interface fc1/2 swwn 20:00:54:7f:ee:e3:86:50
  interface fc1/3 swwn 20:00:54:7f:ee:e3:86:50
  interface fc1/5 swwn 20:00:54:7f:ee:e3:86:50
  interface fc1/6 swwn 20:00:54:7f:ee:e3:86:50
  interface fc1/2 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/3 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/5 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/6 swwn 20:00:54:7f:ee:b8:24:c0

zone name STOR_Zone_1_20_25B vsan 20
  interface fc1/2 swwn 20:00:54:7f:ee:e3:86:50
  interface fc1/3 swwn 20:00:54:7f:ee:e3:86:50
  interface fc1/5 swwn 20:00:54:7f:ee:e3:86:50
  interface fc1/6 swwn 20:00:54:7f:ee:e3:86:50
  interface fc1/2 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/3 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/5 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/6 swwn 20:00:54:7f:ee:b8:24:c0
FC_switch_B_1#
```

- c. Find the WWNs of the switches.

In this example, the two switch WWNs are as follows:

- FC_switch_A_1: 20:00:54:7f:ee:b8:24:c0
- FC_switch_B_1: 20:00:54:7f:ee:c6:80:78

```
FC_switch_B_1# show wwn switch
Switch WWN is 20:00:54:7f:ee:c6:80:78
FC_switch_B_1#

FC_switch_A_1# show wwn switch
Switch WWN is 20:00:54:7f:ee:b8:24:c0
FC_switch_A_1#
```

a. Remove zone members that do not belong to the switch WWNs of the two switches.

In this example, “no member interface” in the output shows that the following members are not associated with the switch WWN of either of the switches in the fabric and must be removed:

- zone name FC-VI_Zone_1_10 vsan 10
 - interface fc1/1 swwn 20:00:54:7f:ee:e3:86:50
 - interface fc1/2 swwn 20:00:54:7f:ee:e3:86:50
 - zone name STOR_Zone_1_20_25A vsan 20
 - interface fc1/5 swwn 20:00:54:7f:ee:e3:86:50
 - interface fc1/8 swwn 20:00:54:7f:ee:e3:86:50
 - interface fc1/9 swwn 20:00:54:7f:ee:e3:86:50
 - interface fc1/10 swwn 20:00:54:7f:ee:e3:86:50
 - interface fc1/11 swwn 20:00:54:7f:ee:e3:86:50
 - zone name STOR_Zone_1_20_25B vsan 20
 - interface fc1/8 swwn 20:00:54:7f:ee:e3:86:50
 - interface fc1/9 swwn 20:00:54:7f:ee:e3:86:50
 - interface fc1/10 swwn 20:00:54:7f:ee:e3:86:50
 - interface fc1/11 swwn 20:00:54:7f:ee:e3:86:50
- The following example shows the removal of these interfaces:

```

FC_switch_B_1# conf t
FC_switch_B_1(config)# zone name FC-VI_Zone_1_10 vsan 10
FC_switch_B_1(config-zone)# no member interface fc1/1 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# no member interface fc1/2 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# zone name STOR_Zone_1_20_25A vsan
20
FC_switch_B_1(config-zone)# no member interface fc1/5 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# no member interface fc1/8 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# no member interface fc1/9 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# no member interface fc1/10 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# no member interface fc1/11 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# zone name STOR_Zone_1_20_25B vsan
20
FC_switch_B_1(config-zone)# no member interface fc1/8 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# no member interface fc1/9 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# no member interface fc1/10 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# no member interface fc1/11 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# save running-config startup-config
FC_switch_B_1(config-zone)# zoneset distribute full 10
FC_switch_B_1(config-zone)# zoneset distribute full 20
FC_switch_B_1(config-zone)# end
FC_switch_B_1# copy running-config startup-config

```

- b. Add the ports of the replacement switch to the zones.

All the cabling on the replacement switch must be the same as on the impaired switch:

```

FC_switch_B_1# conf t
FC_switch_B_1(config)# zone name FC-VI_Zone_1_10 vsan 10
FC_switch_B_1(config-zone)# member interface fc1/1 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# member interface fc1/2 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# zone name STOR_Zone_1_20_25A vsan 20
FC_switch_B_1(config-zone)# member interface fc1/5 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# member interface fc1/8 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# member interface fc1/9 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# member interface fc1/10 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# member interface fc1/11 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# zone name STOR_Zone_1_20_25B vsan 20
FC_switch_B_1(config-zone)# member interface fc1/8 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# member interface fc1/9 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# member interface fc1/10 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# member interface fc1/11 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# save running-config startup-config
FC_switch_B_1(config-zone)# zoneset distribute full 10
FC_switch_B_1(config-zone)# zoneset distribute full 20
FC_switch_B_1(config-zone)# end
FC_switch_B_1# copy running-config startup-config

```

c. Verify that the zoning is properly configured:

```
show zone
```

The following example output shows the three zones:

```
FC_switch_B_1# show zone
zone name FC-VI_Zone_1_10 vsan 10
  interface fc1/1 swwn 20:00:54:7f:ee:c6:80:78
  interface fc1/2 swwn 20:00:54:7f:ee:c6:80:78
  interface fc1/1 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/2 swwn 20:00:54:7f:ee:b8:24:c0

zone name STOR_Zone_1_20_25A vsan 20
  interface fc1/5 swwn 20:00:54:7f:ee:c6:80:78
  interface fc1/8 swwn 20:00:54:7f:ee:c6:80:78
  interface fc1/9 swwn 20:00:54:7f:ee:c6:80:78
  interface fc1/10 swwn 20:00:54:7f:ee:c6:80:78
  interface fc1/11 swwn 20:00:54:7f:ee:c6:80:78
  interface fc1/8 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/9 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/10 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/11 swwn 20:00:54:7f:ee:b8:24:c0

zone name STOR_Zone_1_20_25B vsan 20
  interface fc1/8 swwn 20:00:54:7f:ee:c6:80:78
  interface fc1/9 swwn 20:00:54:7f:ee:c6:80:78
  interface fc1/10 swwn 20:00:54:7f:ee:c6:80:78
  interface fc1/11 swwn 20:00:54:7f:ee:c6:80:78
  interface fc1/5 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/8 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/9 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/10 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/11 swwn 20:00:54:7f:ee:b8:24:c0
FC_switch_B_1#
```

- d. Enable the connectivity to storage and the controllers.

The following example shows the port usage:

```

FC_switch_A_1# conf t
FC_switch_A_1(config)# interface fc1/1
FC_switch_A_1(config)# no shut
FC_switch_A_1(config)# interface fc1/4
FC_switch_A_1(config)# shut
FC_switch_A_1(config)# interface fc1/2-3
FC_switch_A_1(config)# shut
FC_switch_A_1(config)# interface fc1/5-6
FC_switch_A_1(config)# shut
FC_switch_A_1(config)# interface fc1/7-8
FC_switch_A_1(config)# shut
FC_switch_A_1# copy running-config startup-config
FC_switch_A_1#

```

9. 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.

Change ISL properties and the IOD/OOD configuration on a Cisco FC switch

You can add Inter-Switch Links (ISLs), change ISL speed, and reconfigure in-order delivery (IOD) or out of-order delivery (OOD) settings on a Cisco FC switch.

Add ISLs to a Cisco FC switch

You might need add ISLs to a switch if you are adding or upgrading hardware, for example, adding or upgrading to faster controllers or faster switches.

About this task

Perform these steps on both switches in the fabric to verify ISL connectivity.

Steps

1. Disable the ISL ports of the ISLs to be added on both switches in the fabric:

```
FC_switch_A_1#config t
```

Enter the following configuration commands, one per line. Enter CTRL-Z after you have entered all of the configuration commands.

```
FC_switch_A_1(config)# interface fc1/36
FC_switch_A_1(config-if)# shut
FC_switch_A_1(config)# end
```

2. Insert SFPs into the ports you are adding as ISL ports, and cable them according to [Cable a fabric-attached MetroCluster configuration](#).

Verify that these ports are listed in the cabling documentation for the switch model that you are adding them to.

3. Configure the ISL ports by following the steps in [Cabling the ISLs between MetroCluster sites](#).
4. Enable all ISL ports (if not enabled) on both switches in the fabric:

```
FC_switch_A_1# config t
```

Enter the following configuration commands, one per line. End with CTRL-Z after you have entered all of the configuration commands.

```
FC_switch_A_1# interface fc1/36
FC_switch_A_1(config-if)# no shut
FC_switch_A_1(config)# end
```

5. Verify that the ISLs are established between both switches:

```
show topology isl
```

6. Repeat the procedure on the second fabric:

```

-----
-----

```

	Local			Remote			VSAN	Cost	I/F	PC
I/F	Band	PC Domain	SwName	Port	Port	SwName	Domain	PC	Stat	Stat
Speed	width									
16g	64g	1	0x11 cisco9	fc1/36	fc1/36	cisco9	0xbc	1	1	15 up up
16g	64g	1	0x11 cisco9	fc1/40	fc1/40	cisco9	0xbc	1	1	15 up up
16g	64g	1	0x11 cisco9	fc1/44	fc1/44	cisco9	0xbc	1	1	15 up up
16g	64g	1	0x11 cisco9	fc1/48	fc1/48	cisco9	0xbc	1	1	15 up up

```

-----
-----

```

Change ISL port speeds on a Cisco FC switch

You can change the speed of ISL ports on a switch to improve the quality of the ISL, for example, lowering the speed on ISLs traveling a greater distance.

About this task

Perform these steps on both switches in the fabric to verify ISL connectivity.

Steps

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

```
FC_switch_A_1# config t
```

Enter the following configuration commands, one per line. End with CTRL-Z after you have entered all of the configuration commands.

```
FC_switch_A_1(config)# interface fc1/36
FC_switch_A_1(config-if)# shut
FC_switch_A_1(config)# end
```

2. Change the speed of the ISL ports on both switches in the fabric:

```
FC_switch_A_1# config t
```

Enter the following configuration commands, one per line. End with CTRL-Z after you have entered all of the configuration commands.

```
FC_switch_A_1(config)# interface fc1/36
FC_switch_A_1(config-if)# switchport speed 16000
```



Speeds for the ports are 16 = 16,000 Gbps, 8 = 8,000 Gbps, and 4 = 4,000 Gbps.

Verify that the ISL ports for your switch are listed in [Install a fabric-attached MetroCluster configuration](#).

3. Enable all ISL ports (if not enabled) on both switches in the fabric:

```
FC_switch_A_1# config t
```

Enter the following configuration commands, one per line. End with CTRL-Z after you have entered all of the configuration commands.

```
FC_switch_A_1(config)# interface fc1/36
FC_switch_A_1(config-if)# no shut
FC_switch_A_1(config)# end
```

4. Verify that the ISLs are established between both switches:

```
show topology isl
```

```
-----
-----
          _____Local_____          _____Remote_____  VSAN Cost I/F  PC
I/F  Band
      PC Domain SwName   Port   Port   SwName Domain PC           Stat Stat
Speed width
-----
-----
      1   0x11  cisco9  fc1/36  fc1/36  cisco9  0xbc   1    1   15 up   up
16g   64g
      1   0x11  cisco9  fc1/40  fc1/40  cisco9  0xbc   1    1   15 up   up
16g   64g
      1   0x11  cisco9  fc1/44  fc1/44  cisco9  0xbc   1    1   15 up   up
16g   64g
      1   0x11  cisco9  fc1/48  fc1/48  cisco9  0xbc   1    1   15 up   up
16g   64g
```

5. Repeat the procedure for the second switch fabric.

Reconfigure the VSAN to guarantee IOD or OOD of frames

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

Reconfigure IOD

Perform the following step to reconfigure IOD of frames.

Steps

1. Enter configuration mode:

```
conf t
```

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

```
in-order-guarantee vsan <vsan-ID>
```



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

- a. Enable load balancing for the VSAN:

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

- b. Exit configuration mode:

```
end
```

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

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

The commands to configure IOD of frames on FC_switch_A_1:

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

The commands to configure IOD of frames on FC_switch_B_1:

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

Reconfigure OOD

Perform the following steps to reconfigure OOD of frames.

Steps

1. Enter configuration mode:

```
conf t
```

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

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

3. Enable load balancing for the VSAN:

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

4. Exit configuration mode:

```
end
```

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

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

The commands to configure OOD of frames on FC_switch_A_1:

```
FC_switch_A_1# config t
FC_switch_A_1(config)# no in-order-guarantee vsan 10
FC_switch_A_1(config)# vsan database
FC_switch_A_1(config-vsan-db)# vsan 10 loadbalancing src-dst-id
FC_switch_A_1(config-vsan-db)# end
FC_switch_A_1# copy running-config startup-config
```

The commands to configure OOD of frames on FC_switch_B_1:

```
FC_switch_B_1# config t
FC_switch_B_1(config)# no in-order-guarantee vsan 10
FC_switch_B_1(config)# vsan database
FC_switch_B_1(config-vsan-db)# vsan 10 loadbalancing src-dst-id
FC_switch_B_1(config-vsan-db)# end
FC_switch_B_1# copy running-config startup-config
```



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

[Learn about configuring IOD or OOD of frames on ONTAP software.](#)

Change the vendor or model of FC switches

You might need to change the vendor for FC switches from Cisco to Brocade or vice versa, change the switch model, or change both.

About this task

- This procedure applies when you are using NetApp validated switches.
- [Enable console logging](#) before performing this task.
- You must perform the steps in this procedure on one Fabric at a time, for both Fabrics in the configuration.

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
```

```
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 switches before installation.

Follow the steps in [Configure the FC switches](#).
3. Disconnect the connections from the old switches by removing the connections in the following order:
 - a. Disconnect the MetroCluster FC and FCVI interfaces.
 - b. Disconnect the ATTO FibreBridge bridges.
 - c. Disconnect the MetroCluster ISLs.
4. Power off the old switches, remove the cables, and physically replace the old switches with the new switch.
5. Cable the switches in the following order:

You must follow the steps in [Cabling a fabric-attached MetroCluster configuration](#).

- a. Cable the ISLs to the remote site.
 - b. Cable the ATTO FibreBridge bridges.
 - c. Cable the MetroCluster FC and FCVI interfaces.
6. Power up the switches.
 7. Verify that the MetroCluster configuration is healthy by repeating [Step 1](#).
 8. Repeat Step 1 to Step 7 for the second Fabric in the configuration.

Replacing a shelf nondisruptively in a fabric-attached MetroCluster configuration

You might need to know how to replace a shelf nondisruptively in a fabric-attached MetroCluster configuration.



This procedure is only for use in a fabric-attached MetroCluster configuration.

Disabling access to the shelf

You must disable access to the shelf before you replace the shelf modules.

Check the overall health of the configuration. If the system does not appear healthy, address the issue first before proceeding.

Steps

1. From both clusters, offline all plexes with disks on the affected shelf stack:

```
aggr offline plex_name
```

The example shows the commands for offlining plexes for a controller running ONTAP.

```
cluster_A_1::> storage aggregate plex offline -aggr aggrA_1_0 -plex
plex0
cluster_A_1::> storage aggregate plex offline -aggr dataA_1_data -plex
plex0
cluster_A_2::> storage aggregate plex offline -aggr aggrA_2_0 -plex
plex0
cluster_A_2::> storage aggregate plex offline -aggr dataA_2_data -plex
plex0
```

2. Verify that the plexes are offline:


```
aggr status -raggr_name
```

The example shows the commands for verifying that the aggregates are offline for a controller running cMode.

```
Cluster_A_1::> storage aggregate show -aggr aggrA_1_0
Cluster_A_1::> storage aggregate show -aggr dataA_1_data
Cluster_A_2::> storage aggregate show -aggr aggrA_2_0
Cluster_A_2::> storage aggregate show -aggr dataA_2_data
```

3. Disable the SAS ports or switch ports depending on whether the bridges connecting the target shelf are connecting a single SAS stack or two or more SAS stacks:
 - If the bridges are connecting a single SAS stack, disable the switch ports that the bridges are connected to using the appropriate command for your switch.

The following example shows a pair of bridges that connect a single SAS stack, which contains the target shelf:



Switch ports 8 and 9 on each switch connect the bridges to the network.

The following example shows ports 8 and 9 being disabled on a Brocade switch.

```
FC_switch_A_1:admin> portDisable 8
FC_switch_A_1:admin> portDisable 9

FC_switch_A_2:admin> portDisable 8
FC_switch_A_2:admin> portDisable 9
```


The following example shows port 8 and 9 being disabled on a Cisco switch.

```
FC_switch_A_1# conf t
FC_switch_A_1(config)# int fc1/8
FC_switch_A_1(config)# shut
FC_switch_A_1(config)# int fc1/9
FC_switch_A_1(config)# shut
FC_switch_A_1(config)# end

FC_switch_A_2# conf t
FC_switch_A_2(config)# int fc1/8
FC_switch_A_2(config)# shut
FC_switch_A_2(config)# int fc1/9
FC_switch_A_2(config)# shut
FC_switch_A_2(config)# end
```

- If the bridges are connecting two or more SAS stacks, disable the SAS ports connecting the bridges to the target shelf:
SASportDisable port number

The following example shows a pair of bridges that connect four SAS stacks. SAS stack 2 contains the target shelf:



SAS port B connects the bridges to the target shelf. By disabling only SAS port B on both shelves, the other SAS stacks can continue to serve data during the replacement procedure.

In this case, disable the SAS port connecting the bridge to the target shelf:

```
SASportDisable port number
```

The following example shows SAS port B being disabled from the bridge and also verifies that it is disabled. You must repeat the command on both bridges.

```
Ready. *
SASPortDisable B

SAS Port B has been disabled.
```

4. If you previously disabled the switch ports, verify that they are disabled:

```
switchShow
```

The example shows that the switch ports are disabled on a Brocade switch.

```
FC_switch_A_1:admin> switchShow
FC_switch_A_2:admin> switchShow
```

The example shows that the switch ports are disabled on a Cisco switch.

```
FC_switch_A_1# show interface fc1/6
FC_switch_A_2# show interface fc1/6
```

5. Wait for ONTAP to realize that the disk is missing.
6. Power off the shelf that you want to replace.

Replacing the shelf

You must physically remove all of the cables and the shelf before inserting and cabling the new shelf and shelf modules.

Steps

1. Remove all disks and disconnect all cables from the shelf that is being replaced.
2. Remove the shelf modules.
3. Insert the new shelf.
4. Insert the new disks into the new shelf.
5. Insert the shelf modules.
6. Cable the shelf (SAS or Power).
7. Power on the shelf.

Reenabling access and verifying the operation

After the shelf has been replaced, you need to reenabling access and verify that the new shelf is operating correctly.

Steps

1. Verify that the shelf powers properly and the links on the IOM modules are present.
2. Enable the switch ports or SAS port according to the following scenarios:

Option	Step
--------	------

If you previously disabled switch ports

- a. Enable the switch ports:

```
portEnable port number
```

The example shows the switch port being enabled on a Brocade switch.

```
Switch_A_1:admin> portEnable 6  
Switch_A_2:admin> portEnable 6
```

The example shows the switch port being enabled on a Cisco switch.

```
Switch_A_1# conf t  
Switch_A_1(config)# int fc1/6  
Switch_A_1(config)# no shut  
Switch_A_1(config)# end  
  
Switch_A_2# conf t  
Switch_A_2(config)# int fc1/6  
Switch_A_2(config)# no shut  
Switch_A_2(config)# end
```

If you previously disabled a SAS port

- a. Enable the SAS port connecting the stack to the shelf location:

```
SASportEnable port number
```

The example shows SAS port A being enabled from the bridge and also verifies that it is enabled.

```
Ready. *  
SASPortEnable A  
  
SAS Port A has been enabled.
```

3. If you previously disabled the switch ports, verify that they are enabled and online and that all devices are logged in correctly:

```
switchShow
```

The example shows the `switchShow` command for verifying that a Brocade switch is online.

```
Switch_A_1:admin> SwitchShow  
Switch_A_2:admin> SwitchShow
```

The example shows the `switchShow` command for verifying that a Cisco switch is online.

```
Switch_A_1# show interface fc1/6
Switch_A_2# show interface fc1/6
```



After several minutes, ONTAP detects that new disks have been inserted and displays a message for each new disk.

4. Verify that the disks have been detected by ONTAP:

```
sysconfig -a
```

5. Online the plexes that were offline earlier:

```
aggr onlineplex_name
```

The example shows the commands for placing plexes on a controller running cMode back online.

```
Cluster_A_1::> storage aggregate plex online -aggr aggr1 -plex plex2
Cluster_A_1::> storage aggregate plex online -aggr aggr2 -plex plex6
Cluster_A_1::> storage aggregate plex online -aggr aggr3 -plex plex1
```

The plexes begin to resynchronize.



You can monitor the progress of resynchronization using the `aggr status -raggr_name` command.

Hot add storage to a MetroCluster FC configuration

Hot-adding a SAS disk shelf in a direct-attached MetroCluster FC configuration using SAS optical cables

You can use SAS optical cables to hot-add a SAS disk shelf to an existing stack of SAS disk shelves in a direct-attached MetroCluster FC configuration, or as a new stack to a SAS HBA or an onboard SAS port on the controller.

- This procedure is nondisruptive and takes approximately two hours to complete.
- You need the admin password and access to an FTP or SCP server.

This task applies to a MetroCluster FC configuration in which the storage is connected directly to the storage controllers with SAS cables. It does not apply to MetroCluster FC configurations using FC-to-SAS bridges or FC switch fabrics.

Steps

1. Follow the instructions for hot-adding a SAS disk shelf in the *Installation Guide* for your disk shelf model to perform the following tasks to hot-add a disk shelf:

- a. Install a disk shelf for a hot-add.
- b. Turn on the power supplies and set the shelf ID for a hot-add.
- c. Cable the hot-added disk shelf.
- d. Verify SAS connectivity.

Hot add SAS storage to a bridge-attached MetroCluster FC configuration

Hot-adding a stack of SAS disk shelves to an existing pair of FibreBridge 7600N or 7500N bridges

You can hot-add a stack of SAS disk shelves to an existing pair of FibreBridge 7600N or 7500N bridges that have available ports.

Before you begin

- You must have downloaded the latest disk and disk shelf firmware.
- All of the disk shelves in the MetroCluster configuration (existing shelves) must be running the same firmware version. If one or more of the disks or shelves are not running the latest firmware version, update the firmware before attaching the new disks or shelves.

[NetApp Downloads: Disk Drive Firmware](#)

[NetApp Downloads: Disk Shelf Firmware](#)

- The FibreBridge 7600N or 7500N bridges must be connected and have available SAS ports.

About this task

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

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

You can use other management interfaces, if required. These options include using a serial port or Telnet to configure and manage a bridge and to configure the Ethernet management 1 port, and using FTP to update the bridge firmware. If you choose any of these management interfaces, you must meet the applicable requirements in [Other bridge management interfaces](#).



If you insert a SAS cable into the wrong port, when you remove the cable from a SAS port, you must wait at least 120 seconds before plugging the cable into a different SAS port. If you fail to do so, the system will not recognize that the cable has been moved to another port.

Steps

1. Properly ground yourself.
2. From the console of either controller, verify that your system has disk autoassignment enabled:

```
storage disk option show
```

The Auto Assign column indicates whether disk autoassignment is enabled.

Node	BKg. FW. Upd.	Auto Copy	Auto Assign	Auto Assign Policy
node_A_1	on	on	on	default
node_A_2	on	on	on	default
2 entries were displayed.				

3. On each bridge in the pair, enable the SAS port that will connect to the new stack:

```
SASPortEnable port-letter
```

The same SAS port (B, C, or D) must be used on both bridges.

4. Save the configuration and reboot each bridge:

```
SaveConfiguration Restart
```

5. Cable the disk shelves to the bridges:

- a. Daisy-chain the disk shelves in each stack.

The *Installation and Service Guide* for your disk shelf model provides detailed information about daisy-chaining disk shelves.

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

[Fabric-attached MetroCluster installation and configuration](#)

[Stretch MetroCluster installation and configuration](#)

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



The bridge SAS port B is disabled.

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

If you are using the...	Then...
-------------------------	---------

ATTO ExpressNAV GUI	<p>a. In a supported web browser, enter the IP address of a bridge in the browser box.</p> <p>You are brought to the ATTO FibreBridge home page, which has a link.</p> <p>b. Click the link, and then enter your user name and the password that you designated when you configured the bridge.</p> <p>The ATTO FibreBridge status page appears with a menu to the left.</p> <p>c. Click Advanced in the menu.</p> <p>d. View the connected devices:</p> <pre>sastargets</pre> <p>e. Click Submit.</p>
Serial port connection	<p>View the connected devices:</p> <pre>sastargets</pre>

The output shows the devices (disks and disk shelves) to which the bridge is connected. The output lines are sequentially numbered so that you can quickly count the devices.



If the text “response truncated” appears at the beginning of the output, you can use Telnet to connect to the bridge, and then view all of the output by using the `sastargets` command.

The following output shows that 10 disks are connected:

Tgt	VendorID	ProductID	Type	SerialNumber
0	NETAPP	X410_S15K6288A15	DISK	3QP1CLE300009940UHJV
1	NETAPP	X410_S15K6288A15	DISK	3QP1ELF600009940V1BV
2	NETAPP	X410_S15K6288A15	DISK	3QP1G3EW00009940U2M0
3	NETAPP	X410_S15K6288A15	DISK	3QP1EWMP00009940U1X5
4	NETAPP	X410_S15K6288A15	DISK	3QP1FZLE00009940G8YU
5	NETAPP	X410_S15K6288A15	DISK	3QP1FZLF00009940TZKZ
6	NETAPP	X410_S15K6288A15	DISK	3QP1CEB400009939MGXL
7	NETAPP	X410_S15K6288A15	DISK	3QP1G7A900009939FNNT
8	NETAPP	X410_S15K6288A15	DISK	3QP1FY0T00009940G8PA
9	NETAPP	X410_S15K6288A15	DISK	3QP1FXW600009940VERQ

7. Verify that the command output shows that the bridge is connected to all of the appropriate disks and disk shelves in the stack.

If the output is...	Then...
----------------------------	----------------

Correct	Repeat the previous step for each remaining bridge.
Not correct	<ol style="list-style-type: none"> a. Check for loose SAS cables or correct the SAS cabling by repeating the step to cable the disk shelves to the bridges. b. Repeat the previous step for each remaining bridge.

8. Update the disk drive firmware to the most current version from the system console:

```
disk_fw_update
```

You must run this command on both controllers.

[NetApp Downloads: Disk Drive Firmware](#)

9. Update the disk shelf firmware to the most current version by using the instructions for the downloaded firmware.

You can run the commands in the procedure from the system console of either controller.

[NetApp Downloads: Disk Shelf Firmware](#)

10. If your system does not have disk autoassignment enabled, assign disk drive ownership.

[Disk and aggregate management](#)



If you are splitting the ownership of a single stack of disk shelves among multiple controllers, you must disable disk autoassignment (`storage disk option modify -autoassign off *` from both nodes in the cluster) before assigning disk ownership; otherwise, when you assign any single disk drive, the remaining disk drives might be automatically assigned to the same controller and pool.



You must not add disk drives to aggregates or volumes until after the disk drive firmware and disk shelf firmware have been updated and the verification steps in this task have been completed.

11. 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 bridges after adding the new stacks:

```
storage bridge 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.

12. If applicable, repeat this procedure for the partner site.

Hot-adding a stack of SAS disk shelves and bridges to a MetroCluster system

You can hot-add (nondisruptively add) an entire stack, including the bridges, to the MetroCluster system. There must be available ports on the FC switches and you must update switch zoning to reflect the changes.

About this task

- This procedure can be used to add a stack using FibreBridge 7600N or 7500N bridges.
- This procedure is written with the assumption that you are using the recommended bridge management interfaces: the ATTO ExpressNAV GUI and the ATTO QuickNAV utility.
 - You use the ATTO ExpressNAV GUI to configure and manage a bridge, and to update the bridge firmware. You use the ATTO QuickNAV utility to configure the bridge Ethernet management 1 port.
 - You can use other management interfaces, if needed. These options include using a serial port or Telnet to configure and manage a bridge, and to configure the Ethernet management 1 port, and using FTP to update the bridge firmware. If you choose any of these management interfaces, your system must meet the applicable requirements in [Other bridge management interfaces](#)

Preparing to hot-add a stack of SAS disk shelves and bridges

Preparing to hot-add a stack of SAS disk shelves and a pair of bridges involves downloading documents as well as the disk drive and disk shelf firmware.

Before you begin

- Your system must be a supported configuration and must be running a supported version of ONTAP.

[NetApp Interoperability Matrix Tool](#)

- All disk drives and disk shelves in the system must be running the latest firmware version.

You might want to update the disk and shelf firmware throughout the MetroCluster configuration prior to adding shelves.

[Upgrade, revert, or downgrade](#)

- Each FC switch must have one FC port available for one bridge to connect to it.



You might need to upgrade the FC switch depending on the FC switch compatibility.

- The computer you are using to set up the bridges must be running an ATTO supported web browser to use the ATTO ExpressNAV GUI: Internet Explorer 8 or 9, or Mozilla Firefox 3.

The *ATTO Product Release Notes* have an up-to-date list of supported web browsers. You can access this document using the information in the steps.

Steps

1. Download or view the following documents from the NetApp Support Site:
 - [NetApp Interoperability Matrix Tool](#)
 - The *Installation and Service Guide* for your disk shelf model.
2. Download content from the ATTO website and from the NetApp website:
 - a. Go to the ATTO FibreBridge Description page.
 - b. Using the link on the ATTO FibreBridge Description page, access the ATTO web site and download the following:
 - *ATTO FibreBridge Installation and Operation Manual* for your bridge model.
 - ATTO QuickNAV utility (to the computer you are using for setup).
 - c. Go to the ATTO FibreBridge Firmware Download page by clicking **Continue** at the end of the ATTO FibreBridge Description page, and then do the following:
 - Download the bridge firmware file as directed on the download page.

In this step, you are only completing the download portion of the instructions provided in the links. You update the firmware on each bridge later, when instructed to do so in the [Hot-adding the stack of shelves](#) section.
 - Make a copy of the ATTO FibreBridge Firmware Download page and release notes for reference later.
3. Download the latest disk and disk shelf firmware, and make a copy of the installation portion of the instructions for reference later.



In this step, you are only completing the download portion of the instructions provided in the links and making a copy of the installation instructions. You update the firmware on each disk and disk shelf later, when instructed to do so in the [Hot-adding the stack of shelves](#) section.

- a. Download the disk firmware and make a copy of the disk firmware instructions for reference later.

[NetApp Downloads: Disk Drive Firmware](#)

- b. Download the disk shelf firmware and make a copy of the disk shelf firmware instructions for reference later.

[NetApp Downloads: Disk Shelf Firmware](#)

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

- a. Acquire a standard Ethernet cable to connect from the bridge Ethernet management 1 port to your network.
- b. Determine a non-default user name and password for accessing the bridges.

It is recommended that you change the default user name and password.

- c. Obtain an IP address, subnet mask, and gateway information for the Ethernet management 1 port on each bridge.
- d. Disable VPN clients on the computer you are using for setup.

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

5. Acquire four screws for each bridge to flush-mount the bridge “L” brackets securely to the front of the rack.

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

6. If necessary, update the FC switch zoning to accommodate the new bridges that are being added to the configuration.

If you are using the Reference Configuration Files provided by NetApp, the zones have been created for all ports, so you do not need to make any zoning updates. There must be a storage zone for each switch port that connects to the FC ports of the bridge.

Hot-adding a stack of SAS disk shelves and bridges

You can hot-add a stack of SAS disk shelves and bridges to increase the capacity of the bridges.

The system must meet all of the requirements to hot-add a stack of SAS disk shelves and bridges.

Preparing to hot-add a stack of SAS disk shelves and bridges

- Hot-adding a stack of SAS disk shelves and bridges is a nondisruptive procedure if all of the interoperability requirements are met.

[NetApp Interoperability Matrix Tool](#)

[Using the Interoperability Matrix Tool to find MetroCluster information](#)

- Multipath HA is the only supported configuration for MetroCluster systems that are using bridges.

Both controller modules must have access through the bridges to the disk shelves in each stack.

- You should hot-add an equal number of disk shelves at each site.
- If you will be using in-band management of the bridge rather than IP management, the steps for configuring the Ethernet port and IP settings can be skipped, as noted in the relevant steps.



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



If you insert a SAS cable into the wrong port, when you remove the cable from a SAS port, you must wait at least 120 seconds before plugging the cable into a different SAS port. If you fail to do so, the system will not recognize that the cable has been moved to another port.

Steps

1. Properly ground yourself.
2. From the console of either controller module, check whether your system has disk autoassignment enabled:

```
storage disk option show
```

The Auto Assign column indicates whether disk autoassignment is enabled.

Node	BKg. FW. Upd.	Auto Copy	Auto Assign	Auto Assign Policy
node_A_1	on	on	on	default
node_A_2	on	on	on	default
2 entries were displayed.				

3. Disable the switch ports for the new stack.
4. If configuring for in-band management, connect a cable from FibreBridge RS-232 serial port to the serial (COM) port on a personal computer.

The serial connection will be used for initial configuration, and then in-band management via ONTAP and the FC ports can be used to monitor and manage the bridge.

5. If configuring for IP management, configure the Ethernet management 1 port for each bridge by following the procedure in section 2.0 of the *ATTO FibreBridge Installation and Operation Manual* for your bridge model.

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

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

6. Configure the bridge.

If you retrieved the configuration information from the old bridge, use the information to configure the new bridge.

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

The *ATTO FibreBridge Installation and Operation Manual* for your bridge model has the most current information on available commands and how to use them.



Do not configure time synchronization on ATTO FibreBridge 7600N or 7500N. The time synchronization for ATTO FibreBridge 7600N or 7500N is set to the cluster time after the bridge is discovered by ONTAP. It is also synchronized periodically once a day. The time zone used is GMT and is not changeable.

- a. If configuring for IP management, configure the IP settings of the bridge.

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

If using the CLI, you must run the following commands:

```
set ipaddress mp1 ip-address
```

```
set ipsubnetmask mp1 subnet-mask
```

```
set ipgateway mp1 x.x.x.x
```

```
set ipdhcp mp1 disabled
```

```
set ethernetspeed mp1 1000
```

- b. Configure the bridge name.

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

Example bridge names for one stack group on each site:

- *bridge_A_1a*
 - *bridge_A_1b*
 - *bridge_B_1a*
 - *bridge_B_1b*
- If using the CLI, you must run the following command:

```
set bridgename bridgename
```

- c. If running ONTAP 9.4 or earlier, enable SNMP on the bridge:

```
set SNMP enabled
```

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

7. Configure the bridge FC ports.

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

The supported FC data rate depends on your model bridge.

- The FibreBridge 7600N bridge supports up to 32, 16, or 8 Gbps.
- The FibreBridge 7500N bridge supports up to 16, 8, or 4 Gbps.



The FCDataRate speed you select is limited to the maximum speed supported by both the bridge and the switch to which the bridge port connects. Cabling distances must not exceed the limitations of the SFPs and other hardware.

If using the CLI, you must run the following command:

```
set FCDataRate port-number port-speed
```

- b. If you are configuring a FibreBridge 7500N bridge, configure the connection mode that the port uses to "ptp".



The FCConnMode setting is not required when configuring a FibreBridge 7600N bridge.

If using the CLI, you must run the following command:

```
set FCConnMode port-number ptp
```

- c. If you are configuring a FibreBridge 7600N or 7500N bridge, you must configure or disable the FC2 port.

- If you are using the second port, you must repeat the previous substeps for the FC2 port.
- If you are not using the second port, then you must disable the port:

```
FCPortDisable port-number
```

- d. If you are configuring a FibreBridge 7600N or 7500N bridge, disable the unused SAS ports:

```
SASPortDisable sas-port
```



SAS ports A through D are enabled by default. You must disable the SAS ports that are not being used. If only SAS port A is used, then SAS ports B, C, and D must be disabled.

8. Secure access to the bridge and save the bridge's configuration.

- a. From the controller prompt check the status of the bridges:

```
storage bridge show
```

The output shows which bridge is not secured.

- b. Check the status of the unsecured bridge's ports:

```
info
```

The output shows the status of Ethernet ports MP1 and MP2.

- c. If Ethernet port MP1 is enabled, run the following command:

```
set EthernetPort mp1 disabled
```



If Ethernet port MP2 is also enabled, repeat the previous substep for port MP2.

- d. Save the bridge's configuration.

You must run the following commands:

SaveConfiguration

FirmwareRestart

You are prompted to restart the bridge.

9. Update the FibreBridge firmware on each bridge.

If the new bridge is the same type as the partner bridge upgrade to the same firmware as the partner bridge. If the new bridge is a different type to the partner bridge, upgrade to the latest firmware supported by the bridge and version of ONTAP. See the section "Updating firmware on a FibreBridge bridge" in *MetroCluster Maintenance*.

10. Cable the disk shelves to the bridges:

- a. Daisy-chain the disk shelves in each stack.

The *Installation Guide* for your disk shelf model provides detailed information about daisy-chaining disk shelves.

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

[Fabric-attached MetroCluster installation and configuration](#)

[Stretch MetroCluster installation and configuration](#)

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



The bridge SAS port B is disabled.

11. Verify that each bridge can detect all of the disk drives and disk shelves to which the bridge is connected.

If you are using the...	Then...
ATTO ExpressNAV GUI	<p>a. In a supported web browser, enter the IP address of a bridge in the browser box.</p> <p>You are brought to the ATTO FibreBridge home page, which has a link.</p> <p>b. Click the link, and then enter your user name and the password that you designated when you configured the bridge.</p> <p>The ATTO FibreBridge status page appears with a menu to the left.</p> <p>c. Click Advanced in the menu.</p> <p>d. View the connected devices: sastargets</p> <p>e. Click Submit.</p>

Serial port connection	View the connected devices: sastargets
------------------------	---

The output shows the devices (disks and disk shelves) to which the bridge is connected. The output lines are sequentially numbered so that you can quickly count the devices.



If the text response truncated appears at the beginning of the output, you can use Telnet to connect to the bridge, and then view all of the output by using the `sastargets` command.

The following output shows that 10 disks are connected:

Tgt	VendorID	ProductID	Type	SerialNumber
0	NETAPP	X410_S15K6288A15	DISK	3QP1CLE300009940UHJV
1	NETAPP	X410_S15K6288A15	DISK	3QP1ELF600009940V1BV
2	NETAPP	X410_S15K6288A15	DISK	3QP1G3EW00009940U2M0
3	NETAPP	X410_S15K6288A15	DISK	3QP1EWMP00009940U1X5
4	NETAPP	X410_S15K6288A15	DISK	3QP1FZLE00009940G8YU
5	NETAPP	X410_S15K6288A15	DISK	3QP1FZLF00009940TZKZ
6	NETAPP	X410_S15K6288A15	DISK	3QP1CEB400009939MGXL
7	NETAPP	X410_S15K6288A15	DISK	3QP1G7A900009939FNNT
8	NETAPP	X410_S15K6288A15	DISK	3QP1FY0T00009940G8PA
9	NETAPP	X410_S15K6288A15	DISK	3QP1FXW600009940VERQ

12. Verify that the command output shows that the bridge is connected to all of the appropriate disks and disk shelves in the stack.

If the output is...	Then...
Correct	Repeat Step 11 for each remaining bridge.
Not correct	<ol style="list-style-type: none"> Check for loose SAS cables or correct the SAS cabling by repeating Step 10. Repeat Step 11.

13. If you are configuring a fabric-attached MetroCluster configuration, cable each bridge to the local FC switches, using the cabling shown in the table for your configuration, switch model, and FC-to-SAS bridge model:



Brocade and Cisco switches use different port numbering, as shown in the following tables.

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

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

DR GROUP 1												
			Brocade 6505		Brocade 6510, Brocade DCX 8510-8		Brocade 6520		Brocade G620, Brocade G620-1, Brocade G630, Brocade G630-1		Brocade G720	
Component	Port		Switc h 1	Switc h 2	Switc h 1	Switc h 2	Switc h 1	Switc h 2	Switc h 1	Switc h 2	Switc h 1	Switc h 2
Stack 1	bridge_x_1a	FC1	8		8		8		8		10	
		FC2	-	8	-	8	-	8	-	8	-	10
	bridge_x_1B	FC1	9	-	9	-	9	-	9	-	11	-
		FC2	-	9	-	9	-	9	-	9	-	11
Stack 2	bridge_x_2a	FC1	10	-	10	-	10	-	10	-	14	-
		FC2	-	10	-	10	-	10	-	10	-	14
	bridge_x_2B	FC1	11	-	11	-	11	-	11	-	17	-
		FC2	-	11	-	11	-	11	-	11	-	17
Stack 3	bridge_x_3a	FC1	12	-	12	-	12	-	12	-	18	-
		FC2	-	12	-	12	-	12	-	12	-	18
	bridge_x_3B	FC1	13	-	13	-	13	-	13	-	19	-
		FC2	-	13	-	13	-	13	-	13	-	19
Stack y	bridge_x_ya	FC1	14	-	14	-	14	-	14	-	20	-
		FC2	-	14	-	14	-	14	-	14	-	20
	bridge_x_yb	FC1	15	-	15	-	15	-	15	-	21	-
		FC2		15		15		15	-	15	-	21



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

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

DR GROUP 2

			Brocade G620, Brocade G620- 1, Brocade G630, Brocade G630-1		Brocade 6510, Brocade DCX 8510-8		Brocade 6520		Brocade G720	
Component	Port	Port	Switch 1	Switch 2	Switch 1	Switch 2	Switch 1	Switch 2	Switch 1	Switch 2
Stack 1	bridge_x_51a	FC1	26	-	32	-	56	-	32	-
		FC2	-	26	-	32	-	56	-	32
	bridge_x_51b	FC1	27	-	33	-	57	-	33	-
		FC2	-	27	-	33	-	57	-	33
Stack 2	bridge_x_52a	FC1	30	-	34	-	58	-	34	-
		FC2	-	30	-	34	-	58	-	34
	bridge_x_52b	FC1	31	-	35	-	59	-	35	-
		FC2	-	31	-	35	-	59	-	35
Stack 3	bridge_x_53a	FC1	32	-	36	-	60	-	36	-
		FC2	-	32	-	36	-	60	-	36
	bridge_x_53b	FC1	33	-	37	-	61	-	37	-
		FC2	-	33	-	37	-	61	-	37
Stack y	bridge_x_5ya	FC1	34	-	38	-	62	-	38	-
		FC2	-	34	-	38	-	62	-	38
	bridge_x_5yb	FC1	35	-	39	-	63	-	39	-
		FC2	-	35	-	39	-	63	-	39



Additional bridges can be cabled to ports 36 - 39 in G620, G630, G620-1, and G-630-1 switches.

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

DR GROUP 1

Comp onent	Port	Brocade 6505		Brocade 6510, Brocade DCX 8510-8		Brocade 6520		Brocade G620, brocade G620- 1, Brocade G630, Brocade G630-1		Brocade G720	
		Switch 1	Switch 2	Switch 1	Switch 2	Switch 1	Switch 2	Switch 1	Switch 2	Switch 1	Switch 2
Stack 1	bridge _x_1a	8		8		8		8		10	
	bridge _x_1b	-	8	-	8	-	8	-	8	-	10
Stack 2	bridge _x_2a	9	-	9	-	9	-	9	-	11	-
	bridge _x_2b	-	9	-	9	-	9	-	9	-	11
Stack 3	bridge _x_3a	10	-	10	-	10	-	10	-	14	-
	bridge _x_4b	-	10	-	10	-	10	-	10	-	14
Stack y	bridge _x_ya	11	-	11	-	11	-	11	-	15	-
	bridge _x_yb	-	11	-	11	-	11	-	11	-	15



Additional bridges can be cabled to ports 12 - 17, 20 and 21 in G620, G630, G620-1, and G630-1 switches. Additional bridges can be cabled to ports 16 - 17, 20 and 21 G720 switches.

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

DR GROUP 2

		Brocade G720		Brocade G620, Brocade G620-1, Brocade G630, Brocade G630-1		Brocade 6510, Brocade DCX 8510-8		Brocade 6520	
Stack 1	bridge_x_51a	32	-	26	-	32	-	56	-
	bridge_x_51b	-	32	-	26	-	32	-	56
Stack 2	bridge_x_52a	33	-	27	-	33	-	57	-
	bridge_x_52b	-	33	-	27	-	33	-	57
Stack 3	bridge_x_53a	34	-	30	-	34	-	58	-
	bridge_x_54b	-	34	-	30	-	34	-	58
Stack y	bridge_x_ya	35	-	31	-	35	-	59	-
	bridge_x_yb	-	35	-	31	-	35	-	59
 Additional bridges can be cabled to ports 32 through 39 in G620, G630, G620-1, and G630-1 switches. Additional bridges can be cabled to ports 36 through 39 in G720 switches.									

14. If you are configuring a bridge-attached MetroCluster system, cable each bridge to the controller modules:
 - a. Cable FC port 1 of the bridge to a 16 Gb or 8 Gb FC port on the controller module in cluster_A.
 - b. Cable FC port 2 of the bridge to the same speed FC port of the controller module in cluster_A.
 - c. Repeat these substeps on other subsequent bridges until all of the bridges have been cabled.

15. Update the disk drive firmware to the most current version from the system console:

```
disk_fw_update
```

You must run this command on both controller modules.

[NetApp Downloads: Disk Drive Firmware](#)

16. Update the disk shelf firmware to the most current version by using the instructions for the downloaded firmware.

You can run the commands in the procedure from the system console of either controller module.

[NetApp Downloads: Disk Shelf Firmware](#)

17. If your system does not have disk autoassignment enabled, assign disk drive ownership.

Disk and aggregate management



If you are splitting the ownership of a single stack of disk shelves among multiple controller modules, you must disable disk autoassignment on both nodes in the cluster (`storage disk option modify -autoassign off *`) before assigning disk ownership; otherwise, when you assign any single disk drive, the remaining disk drives might be automatically assigned to the same controller module and pool.



You must not add disk drives to aggregates or volumes until after the disk drive firmware and disk shelf firmware have been updated and the verification steps in this task have been completed.

18. Enable the switch ports for the new stack.
19. 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.
20. If applicable, repeat this procedure for the partner site.

Related information

[In-band management of the FC-to-SAS bridges](#)

Hot add a SAS disk shelf to a stack of SAS disk shelves

You can hot-add a disk shelf when you want to increase storage without any reduction in performance.

Step 1: Prepare to hot-add a SAS disk shelf

To prepare for hot-adding a SAS disk shelf, you need to download documents along with the disk drive and disk shelf firmware.

Before you begin

- Verify that your system is a supported configuration and is running a supported version of ONTAP.
- Verify that all disk drives and disk shelves in the system are running the latest firmware version.

You might want to update the disk and shelf firmware throughout the MetroCluster configuration before you add shelves.

[Upgrade, revert, or downgrade](#)

Steps

1. Download or view the following documents from the NetApp Support Site:
 - [Interoperability Matrix Tool](#)
 - The *Installation Guide* for your disk shelf model.
2. Verify that the disk shelf you are hot-adding is supported.

[Interoperability Matrix Tool](#)

3. Download the latest disk and disk shelf firmware:



In this step, you only complete the download portion of the instructions. You need to follow the steps in [hot-add a disk shelf](#) to install the disk shelf.

- a. Download the disk firmware and make a copy of the disk firmware instructions for reference later.

[NetApp Downloads: Disk Drive Firmware](#)

- b. Download the disk shelf firmware and make a copy of the disk shelf firmware instructions for reference later.

[NetApp Downloads: Disk Shelf Firmware](#)

Step 2: Hot-add a disk shelf

Use the following procedure to hot-add a disk shelf to a stack.

Before you begin

- Verify that the system meets all of the requirements in [Prepare to hot-add SAS disk shelves](#).
- Verify that your environment meets one of the following scenarios before you hot-add a shelf:
 - You have two FibreBridge 7500N bridges connected to a stack of SAS disk shelves.
 - You have two FibreBridge 7600N bridges connected to a stack of SAS disk shelves.
 - You have one FibreBridge 7500N bridge and one FibreBridge 7600N bridge connected to a stack of SAS disk shelves.

About this task

- This procedure is for hot-adding a disk shelf to the last disk shelf in a stack.

This procedure is written with the assumption that the last disk shelf in a stack is connected from IOM A to bridge A and from IOM B to bridge B.

- This is a nondisruptive procedure.
- You should hot-add an equal number of disk shelves at each site.
- If you are hot-adding more than one disk shelf, you must hot-add one disk shelf at a time.

Each pair of FibreBridge 7500N or 7600N bridges can support up to four stacks.



Hot-adding a disk shelf requires you to update the disk drive firmware on the hot-added disk shelf by running the `storage disk firmware update` command in advanced mode. Running this command can be disruptive if the firmware on existing disk drives in your system is an older version.

If you insert a SAS cable into the wrong port, after you remove the cable from a SAS port, you must wait at least 120 seconds before plugging the cable into a different SAS port. If you fail to do so, the system will not recognize that you have moved the cable to a different port.

Steps

1. Properly ground yourself.
2. Verify disk shelf connectivity from the system console of either controller:

```
sysconfig -v
```

The output is similar to the following:

- Each bridge on a separate line and under each FC port to which it is visible; for example, hot-adding a disk shelf to a set of FibreBridge 7500N bridges results in the following output:

```
FC-to-SAS Bridge:
cisco_A_1-1:9.126L0: ATTO  FibreBridge7500N 2.10  FB7500N100189
cisco_A_1-2:1.126L0: ATTO  FibreBridge7500N 2.10  FB7500N100162
```

- Each disk shelf on a separate line under each FC port to which it is visible:

```
Shelf    0: IOM6  Firmware rev. IOM6 A: 0173 IOM6 B: 0173
Shelf    1: IOM6  Firmware rev. IOM6 A: 0173 IOM6 B: 0173
```

- Each disk drive on a separate line under each FC port to which it is visible:

```
cisco_A_1-1:9.126L1    : NETAPP    X421_HCOBD450A10 NA01 418.0GB
(879097968 520B/sect)
cisco_A_1-1:9.126L2    : NETAPP    X421_HCOBD450A10 NA01 418.0GB
(879097968 520B/sect)
```

3. Check whether your system has disk auto-assignment enabled from the console of either controller:

```
storage disk option show
```

The auto-assignment policy is shown in the Auto Assign column.

Node	BKg. FW. Upd.	Auto Copy	Auto Assign	Auto Assign Policy
node_A_1	on	on	on	default
node_A_2	on	on	on	default
2 entries were displayed.				

4. If your system does not have disk auto-assignment enabled, or if disk drives in the same stack are owned by both controllers, assign disk drives to the appropriate pools.

Disk and aggregate management



- If you are splitting a single stack of disk shelves between two controllers, disk auto-assignment must be disabled before you assign disk ownership; otherwise, when you assign any single disk drive, the remaining disk drives might be automatically assigned to the same controller and pool.

The `storage disk option modify -node <node-name> -autoassign off` command disables disk autoassignment.

- You cannot add drives to aggregates or volumes until after you have updated the disk drive and disk shelf firmware.

5. Update the disk shelf firmware to the most current version by using the instructions for the downloaded firmware.

You can run the commands in the procedure from the system console of either controller.

NetApp Downloads: Disk Shelf Firmware

6. Install and cable the disk shelf:



Do not force a connector into a port. The mini-SAS cables are keyed; when oriented correctly into a SAS port, the SAS cable clicks into place and the disk shelf SAS port LNK LED illuminates green. For disk shelves, you insert a SAS cable connector with the pull tab oriented up (on the topside of the connector).

- a. Install the disk shelf, power it on, and set the shelf ID.

The *Installation Guide* for your disk shelf model provides detailed information about installing disk shelves.



You must power-cycle the disk shelf and keep the shelf IDs unique for each SAS disk shelf within the entire storage system.

- b. Disconnect the SAS cable from the IOM B port of the last shelf in the stack, and then reconnect it to the

same port in the new shelf.

The other end of this cable remains connected to bridge B.

- c. Daisy-chain the new disk shelf by cabling the new shelf IOM ports (of IOM A and IOM B) to the last shelf IOM ports (of IOM A and IOM B).

The *Installation Guide* for your disk shelf model provides detailed information about daisy-chaining disk shelves.

7. Update the disk drive firmware to the most current version from the system console.

NetApp Downloads: Disk Drive Firmware

- a. Change to the advanced privilege level:

```
set -privilege advanced
```

You need to respond with **y** when prompted to continue into advanced mode and see the advanced mode prompt (*>).

- b. Update the disk drive firmware to the most current version from the system console:

```
storage disk firmware update
```

- c. Return to the admin privilege level:

```
set -privilege admin
```

- d. Repeat the previous substeps on the other controller.

8. 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 Active IQ 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.

9. If you are hot-adding more than one disk shelf, repeat the previous steps for each disk shelf that you are hot-adding.

Hot-removing storage from a MetroCluster FC configuration

You can hot-remove drive shelves—physically remove shelves that have had the aggregates removed from the drives—from a MetroCluster FC configuration that is up and serving data. You can hot-remove one or more shelves from anywhere within a stack of shelves or remove a stack of shelves.

- Your system must be a multipath HA, multipath, quad-path HA, or quad-path configuration.
- In a four-node MetroCluster FC configuration, the local HA pair cannot be in a takeover state.
- You must have already removed all aggregates from the drives in the shelves that you are removing.



If you attempt this procedure on non-MetroCluster FC configurations with aggregates on the shelf you are removing, you could cause the system to fail with a multidrive panic.

Removing aggregates involves splitting the mirrored aggregates on the shelves you are removing, and then re-creating the mirrored aggregates with another set of drives.

Disk and aggregate management

- You must have removed drive ownership after removing the aggregates from the drives in the shelves that you are removing.

Disk and aggregate management

- If you are removing one or more shelves from within a stack, you must have factored the distance to bypass the shelves that you are removing.

If the current cables are not long enough, you need to have longer cables available.

This task applies to the following MetroCluster FC configurations:

- Direct-attached MetroCluster FC configurations, in which the storage shelves are directly connected to the storage controllers with SAS cables
- Fabric-attached or bridge-attached MetroCluster FC configurations, in which the storage shelves are connected using FC-to-SAS bridges

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. 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.

2. Set the privilege level to advanced:

```
set -privilege advanced
```

3. Verify that no mailbox drive is on the shelves: **storage failover mailbox-disk show**

4. Remove the shelf according to the steps for the relevant scenario.

Scenario	Steps
To remove an aggregate when the shelf contains either unmirrored, mirrored, or both types of aggregate...	<p>a. Use the <code>storage aggregate delete -aggregate aggregate name</code> command to remove the aggregate.</p> <p>b. Use the standard procedure to remove ownership of all drives in that shelf, and then physically remove the shelf.</p> <p>Follow the instructions in the <i>SAS Disk Shelves Service Guide</i> for your shelf model to hot-remove shelves.</p>

To remove a plex from a mirrored aggregate, you need to unmirror the aggregate.

- a. Identify the plex that you want to remove by using the `run -node local sysconfig -r` command.

In the following example, you can identify the plex from the line `Plex`

`/dpg_mcc_8020_13_a1_aggr1/plex0`. In this case, the plex to specify is “plex0”.


```
dpgmcc_8020_13_a1a2::storage
aggregate> run -node local
sysconfig -r
*** This system has taken over
dpg-mcc-8020-13-a1
Aggregate
dpg_mcc_8020_13_a1_aggr1
(online, raid_dp, mirrored)
(block checksums)
  Plex
/dpg_mcc_8020_13_a1_aggr1/plex
0 (online, normal, active,
pool0)
  RAID group
/dpg_mcc_8020_13_a1_aggr1/plex
0/rg0 (normal, block
checksums)
  RAID Disk Device
HA  SHELF BAY CHAN Pool Type
RPM  Used (MB/blks)  Phys
(MB/blks)
-----
-----
-----
-----
      dparity  mcc-cisco-8Gb-
fab-2:1-1.126L16 0c 32 15
FC:B 0 SAS 15000
272000/557056000
274845/562884296
      parity  mcc-cisco-8Gb-
fab-2:1-1.126L18 0c 32 17
FC:B 0 SAS 15000
272000/557056000
274845/562884296
      data  mcc-cisco-8Gb-
fab-2:1-1.126L19 0c 32 18
FC:B 0 SAS 15000
272000/557056000
274845/562884296
      data  mcc-cisco-8Gb-
```

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

If you need to perform site maintenance or relocate a single site in a MetroCluster FC 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

- Before you begin, check that any non-mirrored aggregates at the site are offline.

- Verify the operation of the MetroCluster configuration in ONTAP:
 - Check whether the system is multipathed:

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

- Check for any health alerts on both clusters:

```
system health alert show
```

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

```

FC:B 0 SAS 15000
272000/557056000
274845/562884296
data mcc-cisco-8Gb-
fab-3:1-1.126L22 0d 32 21
272000/557056000
274845/562884296

```

```

fab-3:1-1.126L37 0d 34 10
FC:A 1 SAS 15000
272000/557056000
280104/573653840

```

```

parity mcc-cisco-8Gb-
fab-3:1-1.126L14 0d 33 13
FC:A 1 SAS 15000
272000/557056000
280104/573653840

```

```

data mcc-cisco-8Gb-
fab-3:1-1.126L41 0d 34 14
FC:A 1 SAS 15000
272000/557056000
280104/573653840

```

```

data mcc-cisco-8Gb-
fab-3:1-1.126L15 0d 33 14
FC:A 1 SAS 15000
272000/557056000
280104/573653840

```

```

data mcc-cisco-8Gb-
fab-3:1-1.126L45 0d 34 18

```

```
metrocluster show
```

```
FC:A 1 SAS 15000  
272000/557056000  
280104/573653840
```

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 address any issues discovered.

b. Use the storage aggregate plex delete -aggregate *aggr_name* -plex *plex_name* command to remove the plex.

plex defines the plex name, such as "plex3" or "plex6".

c. Use the standard procedure to remove ownership of all drives in that shelf, and then physically remove the shelf.

Follow the instructions in the [SAS Disk Shelves Service Guide](#) for your shelf model to hot-remove shelves.

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.

The unmirrored aggregates will only be online after a switchover if the remote disks in the aggregate are accessible. If the ISLs fail, the local node might be unable to access the data in the unmirrored remote disks. The failure of an aggregate can lead to a reboot of the local node.

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. 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>
```

6. Move unmirrored aggregates offline: storage aggregate offline

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

7. 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.

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

Switch type	Action
-------------	--------

For Brocade FC switches...

- a. Use the `portcfgpersistentdisable <port>` command to persistently disable the ports as shown in the following example. This must be done on both switches at the surviving site.

```
Switch_A_1:admin> portcfgpersistentdisable 14
Switch_A_1:admin> portcfgpersistentdisable 15
Switch_A_1:admin>
```

- b. Verify that the ports are disabled using the `switchshow` command shown in the following example:

```
Switch_A_1:admin> switchshow
switchName:      Switch_A_1
switchType:      109.1
switchState:     Online
switchMode:      Native
switchRole:      Principal
switchDomain:    2
switchId:        fffc02
switchWwn:       10:00:00:05:33:88:9c:68
zoning:          ON (T5_T6)
switchBeacon:    OFF
FC Router:       OFF
FC Router BB Fabric ID: 128
Address Mode:    0

  Index Port Address Media Speed State      Proto
  =====
  ...
  14  14   020e00  id   16G   No_Light  FC
Disabled (Persistent)
  15  15   020f00  id   16G   No_Light  FC
Disabled (Persistent)
  ...
Switch_A_1:admin>
```

For Cisco FC switches...

- a. Use the `interface` command to persistently disable the ports. The following example shows ports 14 and 15 being disabled:

```
Switch_A_1# conf t
Switch_A_1(config)# interface fc1/14-15
Switch_A_1(config)# shut

Switch_A_1(config-if)# end
Switch_A_1# copy running-config startup-config
```

- b. Verify that the switch port is disabled using the `show interface brief` command as shown in the following example:

```
Switch_A_1# show interface brief
Switch_A_1
```

9. Power off the equipment at the disaster site.

The following equipment must be turned off in the order shown:

- Storage controllers - the storage controllers should currently be at the `LOADER` prompt, you must power them off completely.
- MetroCluster FC switches
- ATTO FibreBridges (if present)
- 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, FC switches, FibreBridges, and storage shelves.
3. Configure the ISL ports and verify the intersite connectivity.
 - a. Power on the FC switches.



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

- b. Enable the ports.

Enable the ports according to the correct switch types in the following table:

Switch type	Command
-------------	---------

For Brocade FC switches...

- a. Use the `portcfgpersistentenable <port number>` command to persistently enable the port. This must be done on both switches at the surviving site.

The following example shows ports 14 and 15 being enabled on Switch_A_1.

```
switch_A_1:admin>
portcfgpersistentenable 14
switch_A_1:admin>
portcfgpersistentenable 15
switch_A_1:admin>
```

- b. Verify that the switch port is enabled: `switchshow`

The following example shows that ports 14 and 15 are enabled:

```
switch_A_1:admin> switchshow
switchName: Switch_A_1
switchType: 109.1

switchState:    Online
switchMode: Native
switchRole: Principal
switchDomain:    2
switchId:    fffc02
switchWwn:    10:00:00:05:33:88:9c:68
zoning:    ON (T5_T6)
switchBeacon:    OFF
FC Router:    OFF
FC Router BB Fabric ID: 128
Address Mode:    0

Index Port Address Media Speed State
Proto
=====
=====
...
 14 14 020e00 id 16G Online
FC E-Port 10:00:00:05:33:86:89:cb
"Switch_A_1"
 15 15 020f00 id 16G Online
FC E-Port 10:00:00:05:33:86:89:cb
"Switch_A_1" (downstream)
...
switch_A_1:admin>
```

For Cisco FC switches...

a. Enter the `interface` command to enable the port.

The following example shows ports 14 and 15 being enabled on `Switch_A_1`.

```
switch_A_1# conf t
switch_A_1(config)# interface fc1/14-15
switch_A_1(config)# no shut
switch_A_1(config-if)# end
switch_A_1# copy running-config
startup-config
```

b. Verify that the switch port is enabled: `show interface brief`

```
switch_A_1# show interface brief
switch_A_1#
```

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.

Disable the ports based on whether you are using Brocade or Cisco switches as shown in the following table:

Switch type	Command
-------------	---------

For Brocade FC switches...

- a. Enter the `portcfgpersistentdisable <port_number>` command to persistently disable the port.

This must be done on both switches at the surviving site. The following example shows ports 14 and 15 being disabled on Switch_A_1:

```
switch_A_1:admin> portpersistentdisable
14
switch_A_1:admin> portpersistentdisable
15
switch_A_1:admin>
```

- b. Verify that the switch port is disabled: `switchshow`

The following example shows that ports 14 and 15 are disabled:

```
switch_A_1:admin> switchshow
switchName: Switch_A_1
switchType: 109.1
switchState: Online
switchMode: Native
switchRole: Principal
switchDomain: 2
switchId: fffc02
switchWwn: 10:00:00:05:33:88:9c:68
zoning: ON (T5_T6)
switchBeacon: OFF
FC Router: OFF
FC Router BB Fabric ID: 128
Address Mode: 0

Index Port Address Media Speed State
Proto
=====
=====
...
14 14 020e00 id 16G No_Light
FC Disabled (Persistent)
15 15 020f00 id 16G No_Light
FC Disabled (Persistent)
...
switch_A_1:admin>
```

For Cisco FC switches...

a. Disable the port using the `interface` command.

The following example shows ports fc1/14 and fc1/15 being disabled on Switch A_1:

```
switch_A_1# conf t

switch_A_1(config)# interface fc1/14-15
switch_A_1(config)# shut
switch_A_1(config-if)# end
switch_A_1# copy running-config startup-
config
```

b. Verify that the switch port is disabled using the `show interface brief` command.

```
switch_A_1# show interface brief
switch_A_1#
```

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. Disable the ISLs on the FC switches.
 3. Power on the shelves and allow enough time for them to power on completely.
 4. Power on the FibreBridge bridges.
 - a. On the FC switches, verify that the ports connecting the bridges are coming online.

You can use a command such as `switchshow` for Brocade switches, and `show interface brief` for Cisco switches.

b. Verify that the shelves and disks on the bridges are clearly visible.

You can use a command such as `sastargets` on the ATTO CLI.

5. Enable the ISLs on the FC switches.

Enable the ports based on whether you are using Brocade or Cisco switches as shown in the following table:

Switch type	Command
-------------	---------

For Brocade FC switches...

- a. Enter the `portcfgpersistentenable <port>` command to persistently enable the ports. This must be done on both switches at the surviving site.

The following example shows ports 14 and 15 being enabled on Switch_A_1:

```
Switch_A_1:admin> portcfgpersistentenable 14
Switch_A_1:admin> portcfgpersistentenable 15
Switch_A_1:admin>
```

- b. Verify that the switch port is enabled using the `switchshow` command:

```
switch_A_1:admin> switchshow
switchName:      Switch_A_1
switchType:      109.1
switchState:     Online
switchMode:      Native
switchRole:      Principal
switchDomain:    2
switchId:        fffc02
switchWwn:       10:00:00:05:33:88:9c:68
zoning:          ON (T5_T6)
switchBeacon:    OFF
FC Router:       OFF
FC Router BB Fabric ID: 128
Address Mode:    0

  Index Port Address Media Speed State      Proto
  =====
  ...
  14  14   020e00   id   16G   Online   FC
E-Port 10:00:00:05:33:86:89:cb "Switch_A_1"
  15  15   020f00   id   16G   Online   FC
E-Port 10:00:00:05:33:86:89:cb "Switch_A_1"
(downstream)
  ...
switch_A_1:admin>
```

For Cisco FC switches...

- a. Use the `interface` command to enable the ports.

The following example shows port fc1/14 and fc1/15 being enabled on Switch A_1:

```
switch_A_1# conf t
switch_A_1(config)# interface fc1/14-15
switch_A_1(config)# no shut
switch_A_1(config-if)# end
switch_A_1# copy running-config startup-config
```

- b. Verify that the switch port is disabled:

```
switch_A_1# show interface brief
switch_A_1#
```

6. 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
```

7. 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 FC configuration

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

About this task

You must perform the steps in this procedure from both sites, at the same time.



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.

Steps

1. Verify the MetroCluster configuration from both sites in the MetroCluster configuration.
 - a. Confirm the MetroCluster configuration and that the operational mode is normal.
metrocluster show

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

```
run local sysconfig -v
```

- c. Run the following command:

```
storage bridge show
```

- d. Run the following command:

```
storage port show
```

- e. Run the following command:

```
storage switch show
```

- f. Run the following command:

```
network port show
```

- g. Perform a MetroCluster check:

```
metrocluster check run
```

- h. Display the results of the MetroCluster check:

```
metrocluster check show
```

2. Disable AUSO by modifying the AUSO Failure Domain to

```
auso-disabled
```

```
cluster_A_site_A::*>metrocluster modify -auto-switchover-failure-domain
auso-disabled
```

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 by using the following command: **halt**

- For a four-node or eight-node MetroCluster configuration, use the **inhibit-takeover** and **skip-lif-migration-before-shutdown** parameters:

```
system node halt -node node1_SiteA -inhibit-takeover true -ignore
-quorum-warnings true -skip-lif-migration-before-shutdown true
```

- For a two-node MetroCluster configuration, use the command:

```
system node halt -node nodel_SiteA -ignore-quorum-warnings true
```

5. Power off the following equipment at the site:
 - Storage controllers
 - MetroCluster FC switches (if in use and the configuration is not a two-node stretch configuration)
 - ATTO FibreBridges
 - Storage shelves
6. Wait for thirty minutes and then power on the following equipment at the site:
 - Storage shelves
 - ATTO FibreBridges
 - MetroCluster FC switches
 - 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

Maintenance procedures for MetroCluster IP configurations

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. Revert all cluster LIFs that are not currently on their home ports to their home ports:

```
network interface revert -vserver * -lif *
```

13. 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	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           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.

```

Upgrade firmware on MetroCluster IP switches

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

Verify that the RCF is supported

When you change ONTAP version or the switch firmware version, you should verify that you have a reference configuration file (RCF) that is supported for that version. If you use the [RcfFileGenerator](#) tool, the correct RCF is generated for your configuration.

Steps

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

From this switch...	Issue this command...
Broadcom switch	(IP_switch_A_1) # show clibanner
Cisco switch	IP_switch_A_1# show banner motd
NVIDIA SN2100 switch	cumulus@mcc1:mgmt:~\$ nv config find message

Locate the line in the command output that indicates the RCF version. For example, the following output from a Cisco switch indicates that the RCF 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 for MetroCluster IP](#). If you can generate the RCF 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 the switch firmware

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
```

The operation runs in the background.

d. 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        ok
volumes            ok
7 entries were displayed.

```

e. 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 a reference configuration file (RCF) file on a MetroCluster IP switch. For example, if the RCF version that you are running on the switches is not supported by the ONTAP version, the switch firmware version, or both.

Before you begin

- If you are installing new switch firmware, you must install the switch firmware before upgrading the RCF file.
- Before you upgrade the RCF, [verify that the RCF is supported](#).
- [Enable console logging](#) before performing this task.

About this task

- This procedure disrupts traffic on the switch where the RCF file is upgraded. Traffic resumes when 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.

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       ok
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 network, 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/	Local	Discovered	Interface	
Protocol	Port	Device		
Platform				

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	Status	Network	Current	
Vserver	Interface	Admin/Oper	Address/Mask	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.

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 on supported systems 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 9.17.1	<ul style="list-style-type: none">• AFF A800, AFF C800• AFF A20, AFF A30, AFF C30, AFF A50, AFF C60• AFF A70, AFF A90, AFF A1K, AFF C80• FAS50, FAS70, FAS90
ONTAP 9.15.1	<ul style="list-style-type: none">• AFF A400• AFF C400• 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        ok
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        ok
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, 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.



Don't power up the shelves before you power up the controllers. This prevents the controllers from an unintended boot into ONTAP.

3. Power on the shelves, allowing enough time for them to power on completely.
4. Verify that the storage is visible from maintenance mode.

- a. Boot into maintenance mode:

```
boot_ontap maint
```

- b. Verify that the storage is visible from the surviving site.
- c. Verify that the local and remote storage is visible from the node in maintenance mode:

```
disk show -v
```

5. Halt the nodes:

```
halt
```

6. 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.

Power off an entire MetroCluster IP configuration

Power off the entire MetroCluster IP configuration and all of the equipment before starting maintenance or relocation.



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.
 - 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

```
cluster_A_site_A::*>metrocluster modify -auto-switchover-failure-domain  
auso-disabled
```



In a MetroCluster IP configuration, the AUSO Failure Domain is already set to 'auso-disabled' unless the configuration uses 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 at both sites:

```
cluster_A::*>system node halt -node * -inhibit-takeover true -ignore
-quorum-warnings true
cluster_B::*>system node halt -node * -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

Maintenance procedures for all MetroCluster configurations

Replacing a shelf nondisruptively in a stretch MetroCluster configuration

You can replace disk shelves without disruption in a stretch MetroCluster configuration with a fully populated disk shelf or a disk shelf chassis and transfer components from the shelf you are removing.

The disk shelf model you are installing must meet the storage system requirements specified in the [Hardware Universe](#), which includes supported shelf models, supported disk drive types, the maximum number of disk shelves in a stack, and supported ONTAP versions.

Steps

1. Properly ground yourself.
2. Identify all aggregates and volumes that have disks from the loop that contains the shelf you are replacing and make note of the affected plex name.

Either node might contain disks from the loop of the affected shelf and host aggregates or host volumes.

3. Choose one of the following two options based on the replacement scenario you are planning.
 - If you are replacing a complete disk shelf, including the shelf chassis, disks, and I/O modules (IOM), take the corresponding action as described in the table below:

Scenario	Action
The affected plex contains fewer disks from the affected shelf.	Replace the disks one-by-one on the affected shelf with spares from another shelf. <div style="display: flex; align-items: center; margin-top: 10px;">  <p>You can take the plex offline after completing the disk replacement.</p> </div>
The affected plex contains more disks than are in the affected shelf.	Move the plex offline and then delete the plex.
The affected plex has any disk from the affected shelf.	Move the plex offline but do not delete it.

- If you are replacing only the disk shelf chassis and no other components, perform the following steps:
 - a. Offline the affected plexes from the controller where they are hosted:

```
aggregate offline
```

- b. Verify that the plexes are offline:

```
aggregate status -r
```

4. Identify the controller SAS ports to which the affected shelf loop is connected and disable the SAS ports on both site controllers:

```
storage port disable -node node_name -port SAS_port
```

The affected shelf loop is connected to both sites.

5. Wait for ONTAP to recognize that the disk is missing.

- a. Verify that the disk is missing:

```
sysconfig -a or sysconfig -r
```

6. Turn off the power switch on the disk shelf.
7. Unplug all power cords from the disk shelf.
8. Make a record of the ports from which you unplug the cables so that you can cable the new disk shelf in the same way.
9. Unplug and remove the cables connecting the disk shelf to the other disk shelves or the storage system.
10. Remove the disk shelf from the rack.

To make the disk shelf lighter and easier to maneuver, remove the power supplies and IOM. If you will be installing a disk shelf chassis, also remove the disk drives or carriers. Otherwise, avoid removing disk drives or carriers if possible because excessive handling can cause internal drive damage.

11. Install and secure the replacement disk shelf onto the support brackets and rack.
12. If you installed a disk shelf chassis, reinstall power supplies and IOM.
13. Reconfigure the stack of disk shelves by connecting all cables to the replacement disk shelf ports exactly as they were configured on the disk shelf that you removed.
14. Turn on the power to the replacement disk shelf and wait for the disk drives to spin up.
15. Change the disk shelf ID to a unique ID from 0 through 98.
16. Enable any SAS ports that you previously disabled .

- a. Wait for ONTAP to recognize that the disks are inserted.
- b. Verify that the disks are inserted:

```
sysconfig -a or sysconfig -r
```

17. If you are replacing the complete disk shelf (disk shelf chassis, disks, IOM), perform the following steps:



If you are replacing only the disk shelf chassis and no other components, go to Step 19.

- a. Determine whether disk auto assignment is enabled (on).

```
storage disk option modify -autoassign
```

Disk assignment will occur automatically.

- b. If disk auto assignment is not enabled, assign disk ownership manually.

18. Move the plexes back online:

```
aggregate online plex name
```

19. Recreate any plexes that were deleted by mirroring the aggregate.

20. Monitor the plexes as they begin resynchronizing:

```
aggregate status -r <aggregate name>
```

21. Verify that the storage system is functioning as expected:

```
system health alert show
```

When to migrate root volumes to a new destination

You might need to move root volumes to another root aggregate within a two-node or four-node MetroCluster configuration.

Migrating root volumes within a two-node MetroCluster configuration

To migrate root volumes to a new root aggregate within a two-node MetroCluster configuration, you should refer to [How to move mroot to a new root aggregate in a 2-node Clustered MetroCluster with Switchover](#). This procedure shows you how to non-disruptively migrate the root volumes during a MetroCluster switchover operation. This procedure is slightly different than the procedure used on a four-node configuration.

Migrating root volumes within a four-node MetroCluster configuration

To migrate root volumes to a new root aggregate within a four-node MetroCluster configuration, you can use the [system node migrate-root](#) command while meeting the following requirements.

- You can use system node migrate-root to move root aggregates within a four-node MetroCluster configuration.
- All root aggregates must be mirrored.
- You can add new shelves on both sites with smaller drives to host the root aggregate.
- You must check the drive limits that the platform supports before attaching new drives.

[NetApp Hardware Universe](#)

- If you move the root aggregate to smaller drives, you need to accommodate the minimum root volume size of the platform to ensure all core files are saved.



The four-node procedure can also be applied to an eight-node configuration.

Moving a metadata volume in MetroCluster configurations

You can move a metadata volume from one aggregate to another aggregate in a MetroCluster configuration. You might want to move a metadata volume when the source aggregate is decommissioned or unmirrored, or for other reasons that make the aggregate ineligible.

- You must have cluster administrator privileges to perform this task.
- The target aggregate must be mirrored and should not be in the degraded state.
- The available space in the target aggregate must be larger than the metadata volume that you are moving.

Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

2. Identify the metadata volume that should be moved:

```
volume show MDV_CRS*
```

```

Cluster_A::*> volume show MDV_CRS*
Vserver   Volume                Aggregate             State      Type      Size
Available Used%
-----
Cluster_A
          MDV_CRS_14c00d4ac9f311e7922800a0984395f1_A
                Node_A_1_aggr1
                        online      RW        10GB
9.50GB    5%
Cluster_A
          MDV_CRS_14c00d4ac9f311e7922800a0984395f1_B
                Node_A_2_aggr1
                        online      RW        10GB
9.50GB    5%
Cluster_A
          MDV_CRS_15035e66c9f311e7902700a098439625_A
                Node_B_1_aggr1
                        -          RW        -
-         -
Cluster_A
          MDV_CRS_15035e66c9f311e7902700a098439625_B
                Node_B_2_aggr1
                        -          RW        -
-         -
4 entries were displayed.

Cluster_A::>

```

3. Identify an eligible target aggregate:

metrocluster check config-replication show-aggregate-eligibility

The following command identifies the aggregates in cluster_A that are eligible to host metadata volumes:

```
Cluster_A::*> metrocluster check config-replication show-aggregate-
eligibility
```

```
Aggregate Hosted Config Replication Vols Host Addl Vols Comments
-----
-----
Node_A_1_aggr0 - false Root Aggregate
Node_A_2_aggr0 - false Root Aggregate
Node_A_1_aggr1 MDV_CRS_1bc7134a5ddf11e3b63f123478563412_A true -
Node_A_2_aggr1 MDV_CRS_1bc7134a5ddf11e3b63f123478563412_B true -
Node_A_1_aggr2 - true
Node_A_2_aggr2 - true
Node_A_1_Aggr3 - false Unable to determine available space of aggregate
Node_A_1_aggr5 - false Unable to determine mirror configuration
Node_A_2_aggr6 - false Mirror configuration does not match requirement
Node_B_1_aggr4 - false NonLocal Aggregate
```



In the previous example, Node_A_1_aggr2 and Node_A_2_aggr2 are eligible.

4. Start the volume move operation:

```
volume move start -vserver svm_name -volume metadata_volume_name -destination
-aggregate destination_aggregate_name
```

The following command moves metadata volume MDV_CRS_14c00d4ac9f311e7922800a0984395f1 from aggregate Node_A_1_aggr1 to aggregate Node_A_1_aggr2:

```
Cluster_A::*> volume move start -vserver svm_cluster_A -volume
MDV_CRS_14c00d4ac9f311e7922800a0984395f1
-destination-aggregate aggr_cluster_A_02_01

Warning: You are about to modify the system volume
"MDV_CRS_9da04864ca6011e7b82e0050568be9fe_A". This may cause
severe
performance or stability problems. Do not proceed unless
directed to
do so by support. Do you want to proceed? {y|n}: y
[Job 109] Job is queued: Move
"MDV_CRS_9da04864ca6011e7b82e0050568be9fe_A" in Vserver
"svm_cluster_A" to aggregate "aggr_cluster_A_02_01".
Use the "volume move show -vserver svm_cluster_A -volume
MDV_CRS_9da04864ca6011e7b82e0050568be9fe_A" command to view the status
of this operation.
```

5. Verify the state of the volume move operation:

```
volume move show -volume vol_constituent_name
```

6. Return to the admin privilege level:

```
set -privilege admin
```

Renaming a cluster in MetroCluster configurations

Renaming a cluster in a MetroCluster configuration involves making the changes, and then verifying on both the local and remote clusters that the change took effect correctly.

Steps

1. View the cluster names using the

```
metrocluster node show
```

command:

```
cluster_1::*> metrocluster node show
DR
Group Cluster Node          Configuration  DR
-----
State          Mirroring Mode
-----
1      cluster_1
      node_A_1      configured    enabled      normal
      node_A_2      configured    enabled      normal
      cluster_2
      node_B_1      configured    enabled      normal
      node_B_2      configured    enabled      normal
4 entries were displayed.
```

2. Rename the cluster:

```
cluster identity modify -name new_name
```

In the following example, the `cluster_1` cluster is renamed `cluster_A`:

```
cluster_1::*> cluster identity modify -name cluster_A
```

3. Verify on the local cluster that the renamed cluster is running normally:

```
metrocluster node show
```

In the following example, the newly renamed `cluster_A` is running normally:

```

cluster_A::*> metrocluster node show
DR
Group Cluster Node          Configuration  DR
State          Mirroring Mode
-----
-----
1      cluster_A
      node_A_1      configured   enabled   normal
      node_A_2      configured   enabled   normal
      cluster_2
      node_B_1      configured   enabled   normal
      node_B_2      configured   enabled   normal
4 entries were displayed.

```

4. Rename the remote cluster:

```
cluster peer modify-local-name -name cluster_2 -new-name cluster_B
```

In the following example, cluster_2 is renamed cluster_B:

```

cluster_A:::> cluster peer modify-local-name -name cluster_2 -new-name
cluster_B

```

5. Verify on the remote cluster that the local cluster was renamed and is running normally:

```
metrocluster node show
```

In the following example, the newly renamed cluster_B is running normally:

```

cluster_B::*> metrocluster node show
DR
Group Cluster Node          Configuration  DR
State          Mirroring Mode
-----
-----
1      cluster_B
      node_B_1      configured   enabled   normal
      node_B_2      configured   enabled   normal
      cluster_A
      node_A_1      configured   enabled   normal
      node_A_2      configured   enabled   normal
4 entries were displayed.

```

6. Repeat these steps for each cluster that you want to rename.

Verify the health of a MetroCluster configuration

Learn how to verify that the MetroCluster components are healthy.

About this task

- In MetroCluster IP and FC configurations, you can use the CLI to run health check commands and verify the state of the MetroCluster components.
- In MetroCluster IP configurations running ONTAP 9.8 or later, you can also use ONTAP System Manager to monitor and troubleshoot health check alerts.

Steps

Verify the health of the MetroCluster configuration depending on whether you are using the CLI or System Manager.

CLI

Use the follow steps to check the health of a MetroCluster configuration using the CLI.

Steps

1. Verify that the MetroCluster components are healthy:

```
metrocluster check run
```

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

The operation runs in the background.

2. After the `metrocluster check run` operation completes, display the results:

```
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	ok
volumes	ok

7 entries were displayed.

3. Check the status of the running MetroCluster check operation:

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

4. Verify that there are no health alerts:

```
system health alert show
```

ONTAP System Manager (MetroCluster IP only)

Beginning with ONTAP 9.8, System Manager monitors the health of MetroCluster IP configurations and helps you identify and correct problems that might occur.

System Manager periodically checks the health of your MetroCluster IP configuration. When you view the MetroCluster section in the Dashboard, usually the message is "MetroCluster systems are healthy."

However, when a problem occurs, the message will show the number of events. You can click on this message and view the results of the health check for the following components:

- Node
- Network Interface
- Tier (Storage)
- Cluster
- Connection
- Volume
- Configuration Replication

The **Status** column identifies which components have problems, and the **Details** column suggests how to correct the problem.

Steps

1. In System Manager, select **Dashboard**.
2. View the message in the **MetroCluster** section:
 - a. If the message indicates that your MetroCluster configuration is healthy, and the connections between the clusters and the ONTAP Mediator are healthy (shown with check marks), then you have no problems to correct.
 - b. If the message lists the number of events, or the connections have gone down (shown with an "X"), then continue to the next step.
3. Click the message that shows the number of events.

The MetroCluster Health Report displays.

4. Troubleshoot the problems that appear in the report using the suggestions in the **Details** column.
5. When all the problems have been corrected, click **Check MetroCluster Health**.



You should perform all your troubleshooting tasks before running the check because the MetroCluster Health Check uses an intensive amount of resources.

The MetroCluster Health Check runs in the background. You can work on other tasks while you wait for it to finish.

Where to find additional information

You can learn more about configuring, operating, and monitoring a MetroCluster configuration in NetApp's extensive documentation.

Information	Subject
-------------	---------

MetroCluster documentation	<ul style="list-style-type: none"> • All MetroCluster information
NetApp MetroCluster Solution Architecture and Design	<ul style="list-style-type: none"> • A technical overview of the MetroCluster configuration and operation. • Best practices for MetroCluster configuration.
Fabric-attached MetroCluster installation and configuration	<ul style="list-style-type: none"> • Fabric-attached MetroCluster architecture • Cabling the configuration • Configuring the FC-to-SAS bridges • Configuring the FC switches • Configuring the MetroCluster in ONTAP
Stretch MetroCluster installation and configuration	<ul style="list-style-type: none"> • Stretch MetroCluster architecture • Cabling the configuration • Configuring the FC-to-SAS bridges • Configuring the MetroCluster in ONTAP
MetroCluster IP installation and configuration	<ul style="list-style-type: none"> • MetroCluster IP architecture • Cabling the MetroCluster IP configuration • Configuring the MetroCluster in ONTAP
NetApp Documentation: Product Guides and Resources	<ul style="list-style-type: none"> • Monitoring the MetroCluster configuration and performance
MetroCluster Tiebreaker Software installation and configuration	<ul style="list-style-type: none"> • Monitoring the MetroCluster configuration with the MetroCluster Tiebreaker software
Copy-based transition	<ul style="list-style-type: none"> • Transitioning data from 7-Mode storage systems to clustered storage systems

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.