

Maintenance procedures for MetroCluster FC configurations

ONTAP MetroCluster

NetApp August 29, 2025

This PDF was generated from https://docs.netapp.com/us-en/ontapmetrocluster/maintain/task_modify_switch_or_bridge_ip_address_for_health_monitoring.html on August 29, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Maintenance procedures for MetroCluster FC configurations	1
Modify a switch or ATTO bridge IP address for health monitoring	1
Modify a switch IP address	1
Modify an ATTO bridge IP address	1
FC-to-SAS bridge maintenance	2
Support for FibreBridge 7600N bridges in MetroCluster configurations	2
Support for FibreBridge 7500N bridges in MetroCluster configurations	3
Enabling IP port access on the FibreBridge 7600N bridge if necessary	4
Updating firmware on a FibreBridge bridge	5
Replacing a single FC-to-SAS bridge.	11
Replacing a pair of FibreBridge 6500N bridges with 7600N or 7500N bridges	37
Requirements for using other interfaces to configure and manage FibreBridge bridges	58
Hot-replacing a failed power supply module	59
In-band management of the FC-to-SAS bridges	59
Securing or unsecuring the FibreBridge bridge	60
FC switch maintenance and replacement	61
Upgrade or downgrade the firmware on a Brocade FC switch	61
Upgrading or downgrading the firmware on a Cisco FC switch	65
Upgrading to new Brocade FC switches	72
Replacing a Brocade FC switch	76
Renaming a Brocade FC switch	81
Disabling encryption on Brocade FC switches	84
Change ISL properties, ISL ports, or the IOD/OOD configuration on a Brocade switch	88
Replacing a Cisco FC switch	94
Change ISL properties and the IOD/OOD configuration on a Cisco FC switch	103
Change the vendor or model of FC switches	109
Replacing a shelf nondisruptively in a fabric-attached MetroCluster configuration	110
Disabling access to the shelf	111
Replacing the shelf	114
Reenabling access and verifying the operation	114
Hot add storage to a MetroCluster FC configuration	116
Hot-adding a SAS disk shelf in a direct-attached MetroCluster FC configuration using SAS optical	110
	116
Hot add SAS storage to a bridge-attached MetroCluster FC configuration	117
Hot-removing storage from a MetroCluster FC configuration	138
Power off and power on a single site in a MetroCluster FC configuration	142
	142
	148
Powering on the MetroCluster configuration and returning to normal operation	153
Powering on an entire MetroCluster FC configuration	156

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	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
---	--	---	---

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	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.	Hot-swapping a FibreBridge 6500N bridge with a FibreBridge 7600N or 7500N bridge
Consolidating multiple stacks by replacing multiple pairs of FibreBridge 6500N bridges with a single pair of FibreBridge 7500N bridges	Yes	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. 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.	Replacing a pair of FibreBridge 6500N bridges with 7600N or 7500N bridges

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_la -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 ' Owning	"system brid	dge firmware upc	late"		
Job ID Name	Vserver	Node	State		
2246 job-name					
		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"				
	Owning			
Job ID Name	Vserver	Node	State	
	·			
2246 System bridge firmwa	re 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.

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.

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:

cluster_A> system bridge show -fields fw-version,symbolic-name name fw-version symbolic-name ATTO_20000010affeaffe 3.10 A06X bridge_A_1a ATTO_20000010affeffae 3.10 A06X bridge_A_1b ATTO_20000010affeafff 3.10 A06X bridge_A_2a ATTO_20000010affeaffa 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.

- 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.

15. 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

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 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.

lf	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:		
	storage bridge run-cli -name <i>bridge-name</i> -command "info"		
	storage bridge run-cli -name <i>bridge-name</i> -command "sasportlist"		

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
MP2 IP Gateway	=	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					
;Connec	tor	РНҮ	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	В	1	Disabled	12Gb	5001086000a1c704
Device	В	2	Disabled	12Gb	5001086000a1c704
Device	В	3	Disabled	12Gb	5001086000a1c704
Device	В	4	Disabled	12Gb	5001086000a1c704
Device	С	1	Disabled	12Gb	5001086000a1c708
Device	С	2	Disabled	12Gb	5001086000a1c708
Device	С	3	Disabled	12Gb	5001086000a1c708
Device	С	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	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. View the connected devices:
	sastargets
	e. Click Submit .
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	Туре	SerialNumber
0	NETAPP	X410_S15K6288A15	DISK	3QP1CLE300009940UHJV
1	NETAPP	X410_S15K6288A15	DISK	3QP1ELF600009940V1BV
2	NETAPP	X410_S15K6288A15	DISK	3QP1G3EW00009940U2M0
3	NETAPP	X410_S15K6288A15	DISK	3QP1EWMP00009940U1X5
4	NETAPP	X410_S15K6288A15	DISK	3QP1FZLE00009940G8YU
5	NETAPP	X410_S15K6288A15	DISK	3QP1FZLF00009940TZKZ
6	NETAPP	X410_S15K6288A15	DISK	3QP1CEB400009939MGXL
7	NETAPP	X410_S15K6288A15	DISK	3QP1G7A900009939FNTT
8	NETAPP	X410_S15K6288A15	DISK	3QP1FY0T00009940G8PA
9	NETAPP	X410_S15K6288A15	DISK	3QP1FXW600009940VERQ

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	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: BO		
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	storage bridge add -address 0.0.0.0 -managed-by in-band -name bridge-name
9.4 and earlier	storage bridge add -address bridge-ip-address -name bridge-name

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

storage bridge show

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

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

controller A 1::> storage bridge show Symbolic Name Is Monitored Monitor Status Bridge Vendor Model Bridge WWN _____ ____ ATTO 10.10.20.10 atto01 true ok Atto FibreBridge 7500N 2000010867038c0 ATTO 10.10.20.11 atto02 true ok Atto FibreBridge 7500N 2000010867033c0 ATTO_10.10.20.12 atto03 true ok Atto FibreBridge 7500N 2000010867030c0 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 mpl ip-address
set ipsubnetmask mpl subnet-mask
set ipgateway mpl x.x.x.x
set ipdhcp mpl disabled
set ethernetspeed mpl 1000
```

b. Configure the bridge name.

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

Example bridge names for one stack group on each site:

- bridge_A_1a
- bridge_A_1b
- bridge_B_1a
- bridge_B_1b

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

set bridgename bridgename

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

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

- 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	Туре	SerialNumber
0	NETAPP	X410_S15K6288A15	DISK	3QP1CLE300009940UHJV
1	NETAPP	X410_S15K6288A15	DISK	3QP1ELF600009940V1BV
2	NETAPP	X410_S15K6288A15	DISK	3QP1G3EW00009940U2M0
3	NETAPP	X410_S15K6288A15	DISK	3QP1EWMP00009940U1X5
4	NETAPP	X410_S15K6288A15	DISK	3QP1FZLE00009940G8YU
5	NETAPP	X410_S15K6288A15	DISK	3QP1FZLF00009940TZKZ
6	NETAPP	X410_S15K6288A15	DISK	3QP1CEB400009939MGXL
7	NETAPP	X410_S15K6288A15	DISK	3QP1G7A900009939FNTT
8	NETAPP	X410_S15K6288A15	DISK	3QP1FY0T00009940G8PA
9	NETAPP	X410_S15K6288A15	DISK	3QP1FXW600009940VERQ

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 the previous step for each remaining bridge.
Not correct	a. Check for loose SAS cables or correct the SAS cabling by repeating Step 18.b. Repeat the previous step.

21. If the bridge is in a fabric-attached MetroCluster configuration, reenable 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:

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

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

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 Symbolic Name Is Monitored Monitor Status Bridge Vendor Model Bridge WWN _____ ____ _____ _____ ATTO 10.10.20.10 atto01 true ok Atto FibreBridge 7500N 2000010867038c0 ATTO 10.10.20.11 atto02 true ok Atto FibreBridge 7500N 2000010867033c0 ATTO 10.10.20.12 atto03 true ok Atto FibreBridge 7500N 2000010867030c0 ATTO 10.10.20.13 atto04 true ok Atto FibreBridge 7500N 200001086703b80 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	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:
	sastargets
Serial port connection	Display the list of disks visible to the bridge:
	sastargets

The output shows the devices (disks and disk shelves) that the bridge is connected to. Output lines are sequentially numbered so that you can quickly count the devices. For example, the following output shows that 10 disks are connected:
Tgt	VendorID	ProductID	Туре	SerialNumber
0	NETAPP	X410_S15K6288A15	DISK	3QP1CLE300009940UHJV
1	NETAPP	X410_S15K6288A15	DISK	3QP1ELF600009940V1BV
2	NETAPP	X410_S15K6288A15	DISK	3QP1G3EW00009940U2M0
3	NETAPP	X410_S15K6288A15	DISK	3QP1EWMP00009940U1X5
4	NETAPP	X410_S15K6288A15	DISK	3QP1FZLE00009940G8YU
5	NETAPP	X410_S15K6288A15	DISK	3QP1FZLF00009940TZKZ
6	NETAPP	X410_S15K6288A15	DISK	3QP1CEB400009939MGXL
7	NETAPP	X410_S15K6288A15	DISK	3QP1G7A900009939FNTT
8	NETAPP	X410_S15K6288A15	DISK	3QP1FY0T00009940G8PA
9	NETAPP	X410_S15K6288A15	DISK	3QP1FXW600009940VERQ



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

20. 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	 a. Check for loose SAS cables or correct the SAS cabling by repeating Step 18. b. Repeat the previous step for each remaining bridge. 	

- 21. Reenable the FC switch port that connects to the bridge.
- 22. 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	storage bridge add -address 0.0.0.0 -managed-by in-band -name bridge-name
9.4 and earlier	storage bridge add -address <i>bridge-ip-address</i> -name <i>bridge-name</i>

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
        Symbolic Name Is Monitored Monitor Status
Bridge
Vendor Model
                   Bridge WWN
_____
               _____ ____
_____ ____
                    _____
ATTO_10.10.20.10 atto01 true
                                 ok
                                               Atto
FibreBridge 7500N 2000010867038c0
ATTO 10.10.20.11 atto02 true
                                  ok
                                               Atto
FibreBridge 7500N 2000010867033c0
ATTO 10.10.20.12 atto03 true
                                  ok
                                              Atto
FibreBridge 7500N 2000010867030c0
ATTO 10.10.20.13 atto04 true
                                   ok
                                             Atto
FibreBridge 7500N 200001086703b80
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: UR0009R
       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:
 - $\circ\,$ 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>)**<===</pre> 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 FINISAR CORP. SFP Vendor: SFP Part Number: FTLF8529P3BCVAN1 SFP Serial Number: URQOR1R 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 \ensuremath{-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: URQOR1R
       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:



3. 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:



4. 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: URQOR1R
        SFP Capabilities: 4, 8 or 16 Gbit
        Link Data Rate: 16 Gbit
        Switch Port:
                          brcd6505-fcs40:1
```

```
**<List of disks visible to port\>**
        ТD
               Vendor Model
                                               Size
                                         FW
       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
                                                 FibreBridge7500N A30H
                                : ATTO
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	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
the four FibreBridge 6500N bridges.	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	STOR_A_1a-FC1	7,1; 7,4; 5,1; 5,4; 5,6	Purple + blue
zone for each FC port on the two FibreBridge	STOR_A_1b-FC1	7,2; 7,5; 5,2; 5,5; 5,7	Dashed purple + red
7500N blidges.	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</pre>
```

```
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: URQOR1R
       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.

 (\mathbf{i})

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.



(i)

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 reenable 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 reenable 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 reenable the switch.

About this task

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

- 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 Ciscospecific 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:

```
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 $-\boldsymbol{v}$

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

e 0
e 1
.e 2
.e 3
e 4
e 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: 0FF
    switchBeacon: 0FF
```

- 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.

<pre>FC_switch_B_1:admin></pre>	portcfgpersistentdisable	0
<pre>FC_switch_B_1:admin></pre>	portcfgpersistentdisable	1
<pre>FC_switch_B_1:admin></pre>	portcfgpersistentdisable	2
<pre>FC_switch_B_1:admin></pre>	portcfgpersistentdisable	3
<pre>FC_switch_B_1:admin></pre>	portcfgpersistentdisable	4
<pre>FC_switch_B_1:admin></pre>	portcfgpersistentdisable	5
<pre>FC_switch_B_1:admin></pre>	portcfgpersistentdisable	6
<pre>FC_switch_B_1:admin></pre>	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"
         030B00 id 16G Online FC E-Port
11
    11
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,
DPO Configuration replay has completed.
2018/03/09-07:50:48, [ESM-1011], 146105, FID 128, INFO, FC switch 1 B,
DPO 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
10:00:50:eb:1a:9a:a5:79
switchWwn:
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
                    Name Vendor Model Switch WWN
                                                         Monitored
Switch
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…	Use the portcfgspeed port number port speed command on both switches on the fabric.
	In the following example, you change the ISL port speed from 40 Gbps to 16 Gbps:
	<pre>brocade_switch_A_1:admin> portcfgspeed 40 16</pre>
	You can verify that the speed has changed using the switchshow command:
	brocade_switch_A_1:admin> switchshow
	You should see the following output:
	 40 40 062800 id 16G No_Sync FC Disabled
To change the distance of an ISL port…	Use the portcfglongdistance <i>port number port distance</i> 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
```

a. Configure the trunking mode for each ISL <code>portcfgtrunkport</code> *port number* as shown in the following table:

Set the portcfgtrunkport port number to 0:	

b. Configure the fabric for IOD or OOD as required.

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

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

4. 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	If the ISLs are trunked, only a single ISL appears in the output for the islshow command. Either port 40 or 41 can appear depending on which is the trunk master. The output of trunkshow 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:
	<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) # 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(config) # shut
FC_switch_B_1(config) # shut
FC_switch_B_1(config) # shut
```

- 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) # interface fc1/5-6
FC_switch_A_1(config) # shut
FC_switch_A_1(config) # interface fc1/7-8
FC_switch_A_1(config) # interface fc1/7-8
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
```

Insert SFPs into the ports you are adding as ISL ports, and cable them according to Cable a fabricattached 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 _____ 1 0x11 cisco9 fc1/36 fc1/36 cisco9 0xbc 1 1 15 up up 64g 16q 1 0x11 cisco9 fc1/40 fc1/40 cisco9 0xbc 1 1 15 up up 16g 64g 0x11 cisco9 fc1/44 fc1/44 cisco9 0xbc 1 1 1 15 up up 64q 16g 0x11 cisco9 fc1/48 fc1/48 cisco9 0xbc 1 1 1 15 up up 64g 16g

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
```

 (\mathbf{i})

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
                    _____
_____
        0x11 cisco9 fc1/36 fc1/36 cisco9 0xbc 1
    1
                                            1 15 up
                                                       up
16g
    64g
    1 0x11 cisco9 fc1/40 fc1/40 cisco9 0xbc 1 1 15 up
                                                       up
16g
    64a
    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
    64a
```

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 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

show	State
Entry Name	State
Configuration state	configured
Mode	normal
AUSO Failure Domain	auso-on-cluster-
Configuration state	configured
Mode	normal
AUSO Failure Domain	auso-on-cluster-
	show Entry Name Configuration state Mode AUSO Failure Domain Configuration state Mode AUSO Failure Domain

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

cluste DR	er_A::> metrocluster no	de show 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 ent	ries 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) # int fc1/9
FC_switch_A_2(config) # shut
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 reenable 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	<pre>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) # no shut Switch_A_2# conf t Switch_A_2 (config) # int fc1/6 Switch_A_2 (config) # no shut Switch_A_2 (config) # end</pre>
If you previously disabled a SAS port	<pre>Switch_A_2(config) # end 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.</pre>

3. If you previously disabled the switch ports, verify that they are enabled and online and that and 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.
- If you are adding an IOM12 shelf to a stack of IOM6 shelves, see Hot-adding IOM12 shelves to a stack of IOM6 shelves.

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 w	ere 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 daisychaining 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	 a. In a supported web browser, enter the IP address of a bridge in the browser box. You are brought to the ATTO FibreBridge home page, 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. View the connected devices: sastargets e. Click Submit.
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.

(i)

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:

```
Tqt VendorID ProductID
                                          SerialNumber
                              Type
 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
                                          30P1G7A900009939FNTT
 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	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.

All disk shelves in the MetroCluster configuration (both the new shelves and existing shelves) must be running the same firmware version.



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

- 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 w	ere 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 mpl 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	a. In a supported web browser, enter the IP address of a bridge in the browser box.
	You are brought to the ATTO FibreBridge home page, 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. View the connected devices: sastargets
	e. Click Submit.

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:

VendorID	ProductID	Туре	SerialNumber
NETAPP	X410_S15K6288A15	DISK	3QP1CLE300009940UHJV
NETAPP	X410_S15K6288A15	DISK	3QP1ELF600009940V1BV
NETAPP	X410_S15K6288A15	DISK	3QP1G3EW00009940U2M0
NETAPP	X410_S15K6288A15	DISK	3QP1EWMP00009940U1X5
NETAPP	X410_S15K6288A15	DISK	3QP1FZLE00009940G8YU
NETAPP	X410_S15K6288A15	DISK	3QP1FZLF00009940TZKZ
NETAPP	X410_S15K6288A15	DISK	3QP1CEB400009939MGXL
NETAPP	X410_S15K6288A15	DISK	3QP1G7A900009939FNTT
NETAPP	X410_S15K6288A15	DISK	3QP1FY0T00009940G8PA
NETAPP	X410_S15K6288A15	DISK	3QP1FXW600009940VERQ
	<pre>VendorID NETAPP NETAPP NETAPP NETAPP NETAPP NETAPP NETAPP NETAPP</pre>	VendorIDProductIDNETAPPX410_S15K6288A15NETAPPX410_S15K6288A15NETAPPX410_S15K6288A15NETAPPX410_S15K6288A15NETAPPX410_S15K6288A15NETAPPX410_S15K6288A15NETAPPX410_S15K6288A15NETAPPX410_S15K6288A15NETAPPX410_S15K6288A15NETAPPX410_S15K6288A15NETAPPX410_S15K6288A15NETAPPX410_S15K6288A15NETAPPX410_S15K6288A15	VendorIDProductIDTypeNETAPPX410_S15K6288A15DISKNETAPPX410_S15K6288A15DISKNETAPPX410_S15K6288A15DISKNETAPPX410_S15K6288A15DISKNETAPPX410_S15K6288A15DISKNETAPPX410_S15K6288A15DISKNETAPPX410_S15K6288A15DISKNETAPPX410_S15K6288A15DISKNETAPPX410_S15K6288A15DISKNETAPPX410_S15K6288A15DISKNETAPPX410_S15K6288A15DISK

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	a. Check for loose SAS cables or correct the SAS cabling by repeating Step 10.
	b. 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".
- $\circ\,$ 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				
Compo	onent	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			
Compor	nent	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											
		Brocad	e 6505	Brocade 6510, Brocade DCX 8510-8		Brocade 6520		Brocade G620, brocade G620- 1, Brocade G630, Brocade G630-1		Brocade G720	
Comp onent	Port	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
(i)	Add and	itional bri G630-1 s	dges car switches.	n be cabl Addition	ed to por al bridge	rts 12 - 1 es can be	7, 20 and cabled t	d 21 in G to ports 1	620, G63 6 - 17, 2	30, G620 0 and 21	-1,

g , G720 switches.

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

DR GROUP 2

		Brocade	G720	Brocade Brocade Brocade Brocade	G620, G620-1, G630, G630-1	Brocade Brocade 8510-8	6510, DCX	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
i	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



A mix of IOM12 modules and IOM6 modules is supported within the same stack if your system is running a supported ONTAP version. To establish whether your ONTAP version supports shelf mixing, refer to the Interoperability Matrix Tool (IMT). If your version of ONTAP is not supported and you cannot upgrade or downgrade the IOM modules on the existing stack or the new shelf that is to be added to a supported combination of IOM modules, you need to do one of the following:

- Start a new stack on a new SAS port (if supported by the bridge-pair).
- Start a new stack on an additional bridge-pair.

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:
 - $\circ\,$ You have two FibreBridge 7500N bridges connected to a stack of SAS disk shelves.
 - $\circ\,$ 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:

Shelf0: IOM6Firmware rev. IOM6 A: 0173IOM6 B: 0173Shelf1: IOM6Firmware rev. IOM6 A: 0173IOM6 B: 0173

 $\circ\,$ 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 w	ere 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 autoassignment 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 \mathbf{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-adding an IOM12 disk shelf to a stack of IOM6 disk shelves in a bridge-attached MetroCluster configuration

Depending on your version of ONTAP, you can hot-add an IOM12 disk shelf to a stack of IOM6 disk shelves in a bridge-attached MetroCluster configuration.

To perform this procedure, see Hot-adding shelves with IOM12 modules to a stack of shelves with IOM6 modules.

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	 a. Use the storage aggregate delete aggregate aggregate name command to remove the aggregate. b. 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.			

```
To remove a plex from a mirrored aggregate, you
                                       a. Identify the plex that you want to remove by
need to unmirror the aggregate.
                                         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 al 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 Oc 32 17
                                           FC:B 0 SAS 15000
                                           272000/557056000
                                           274845/562884296
                                                data mcc-cisco-8Gb-
                                           fab-2:1-1.126L19 Oc 32 18
                                           FC:B 0 SAS 15000
                                           272000/557056000
                                           274845/562884296
                                                                         141
                                                 data mcc-cisco-8Gb-
```

Power off and power on a single site in: a MetroCluster²FC⁰ 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 of the site. mcc-cisco-8Gbfab-2:1-1.126L22 0c 32 21

If you need to relocate and reconfigure a site (for example, if your cell to expand from 15 1000 node to an eightnode cluster), you cannot complete these tasks at the same time 7 his processings 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.



fab-3:1-1.126L45 Od 34 18
c. Confirm the MetroCluster configuration and that the operational mode is parmal 5000

 metrocluster show
 272000/557056000

 d. Perform a MetroCluster check:
metrocluster check run
 280104/573653840

 e. Display the results of the MetroCluster check:
metrocluster check show
 b. Use the storage aggregate plex delete
-aggregate aggr_name -plex
plex_name command to remove the plex.

f. Check for any health alerts on the switches (if present) ":plex6".

storage switch show

g. Run Config Advisor.

NetApp Downloads: Config Advisor

c. Use the standard procedure to remove ownership of all drives in that shelf, and then physically remove the shelf.

plex defines the plex name, such as "plex3" or

Follow the instructions in the SAS Disk Shelves Service Guide for your shelf model to hot-

- h. After running Config Advisor, review the tool's output are 1000 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.

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>

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:

 $\circ\,$ 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
                                            0
Node B 2 aggr0 plex0 normal, active true
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_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 portofgpersistent disable <port> command to persistently disable the ports as shown in the following example. This must be done on both switches at the surviving site.</port>
	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:
	<pre>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: OFF FC Router BB Fabric ID: 128 Address Mode: 0</pre>
	Index Port Address Media Speed State Proto
	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# 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.</port>
	The following example shows ports 14 and 15 being enabled on Switch_A_1.
	<pre>switch_A_1:admin></pre>
	switch & lighting
	portofapersistentenable 15
	<pre>switch_A_1:admin></pre>
	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: U
	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.
	<pre>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</pre>
	startup-configb. Verify that the switch port is enabled: show interface
	switch_A_1# show interface brief

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

(i)

Command

```
For Brocade FC switches...
                            a. Enter the portcfgpersistent disable <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
	<pre>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</pre>
	b. Verify that the switch port is disabled using the show interface brief command.
	<pre>switch_A_1# show interface brief switch_A_1#</pre>

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
                   a. Enter the portcfgpersistentenable <port> command to
switches...
                     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# copy running-config startup-config

      b. Verify that the switch port is disabled:

      switch_A_1#

      switch_A_1#

      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 node1_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 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**

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at http://www.netapp.com/TM are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.