Lenovo

Metrocluster Service Guide



ONTAP® 9

Fourth edition (March 2023)

© Copyright Lenovo 2019, 2023.
LIMITED AND RESTRICTED RIGHTS NOTICE: If data or software is delivered pursuant to a General Services Administration (GSA) contract, use, reproduction, or disclosure is subject to restrictions set forth in Contract No. GS-35F-05925

Contents

Contents i

Chapter 1. Deciding whether to use the MetroCluster Service Guide 1

Chapter 2. Where to find procedures for MetroCluster maintenance tasks . . 3

Chapter 3. MetroCluster failure and recovery scenarios. 5

Chapter 4. Performing FC-to-SAS bridge maintenance 7

Support for FibreBridge 7600N bridges in MetroCluster configurations 7

Enabling IP port access on the FibreBridge 7600N bridge if necessary 7

Updating firmware on a FibreBridge bridge 8

Updating firmware on FibreBridge 7600N bridge on configurations running ONTAP 9.4 and later 8

Replacing a single FC-to-SAS bridge 11

Verifying storage connectivity 12

Hot-swapping a bridge with a replacement bridge of the same model 13

Hot-swapping a FibreBridge 7500N with a 7600N bridge 19

Hot-swapping a FibreBridge 6500N bridge with a FibreBridge 7600N bridge 24

Replacing a pair of FibreBridge 6500N bridges with 7600N bridge 30

Verifying storage connectivity 30

Hot-swapping FibreBridge 6500N bridges to create a pair of FibreBridge 7600N bridge . . . 31

Cabling the bridge SAS ports when consolidating storage behind FibreBridge 7600N bridge 35

Updating zoning when adding FibreBridge 7600N bridge to a configuration. 38

Cabling the second bridge FC port when adding FibreBridge 7600N bridge to a configuration 39

Disabling unused SAS ports on the FC-to-SAS bridges 41

Requirements for using other interfaces to configure and manage FibreBridge bridges 42

Hot-replacing a failed power supply module. . . . 42

In-band management of the FC-to-SAS bridges . . 43

Managing a FibreBridge bridge from ONTAP 43

Securing or unsecuring the FibreBridge bridge. . . 44

Chapter 5. Performing FC switch maintenance and replacement 45

Upgrading or downgrading the firmware on a Brocade FC switch 45

Upgrading or downgrading the firmware on a Cisco FC switch 47

Upgrading to new Brocade FC switches 48

Replacing a Brocade FC switch 52

Renaming a Brocade FC switch 55

Changing ISL properties, ISL ports, or the IOD/OOD configuration on a Brocade switch 57

Replacing a Cisco FC switch. 61

Changing speed of ISL ports on a Cisco FC switch 67

Adding ISLs to a Cisco switch 68

Chapter 6. Performing IP switch maintenance and replacement 71

Replacing an IP switch 71

Upgrading firmware on MetroCluster IP switches 73

Upgrading RCF files on MetroCluster IP switches 74

Renaming a Cisco IP switch 75

Chapter 7. Adding, removing, or changing ISL ports nondisruptively . . 79

Chapter 8. Identifying storage in a MetroCluster IP configuration 81

Chapter 9. Hot-adding storage to a MetroCluster FC configuration 85

Hot-adding a SAS disk shelf in a direct-attached MetroCluster FC configuration using SAS optical cables 85

Hot-adding SAS storage to a bridge-attached MetroCluster FC configuration 85

Hot-adding a stack of SAS disk shelves to an existing pair of FibreBridge 7600N bridges . . 85

Hot-adding a stack of SAS disk shelves and bridges to a MetroCluster system 91

Hot-adding a SAS disk shelf to a stack of SAS disk shelves 97

Chapter 10. Hot-removing storage from a MetroCluster FC configuration. 101

© Copyright Lenovo 2019, 2023

Chapter 11. Replacing a shelf nondisruptively in a fabric-attached MetroCluster configuration 103
Disabling access to the shelf 103
Replacing the shelf 105
Reenabling access and verifying the operation . . . 106

Chapter 12. When to migrate root volumes to a new destination 109

Chapter 13. Moving a metadata volume in MetroCluster configurations 111

Chapter 14. Renaming a cluster in MetroCluster configurations. 113

Chapter 15. Powering off and powering on a data center 115
Powering off a MetroCluster site 115
Relocating the powered-off site of the MetroCluster 119
Powering on the MetroCluster configuration and returning to normal operation 121

Chapter 16. Powering off an entire MetroCluster IP configuration 125

Chapter 17. Powering off an entire MetroCluster FC configuration 127

Chapter 18. Reconfiguring the FC switch layout 129
Sending a custom AutoSupport message prior to reconfiguring switches 129
Verifying the health of the MetroCluster configuration 129
Persistently disabling the switches 130
Determining the new cabling layout 130
Applying RCF files and recabling the switches . . . 131
Persistently enable the switches 131
Verifying switchover, healing, and switchback . . . 132

Chapter 19. Port assignments for FC switches 133
Port assignments for systems using two initiator ports 133
Port assignments for FC switches. 134

Chapter 20. Where to find additional information 147

Appendix A. Contacting Support . . .149

Appendix B. Notices 151
Trademarks 152

Chapter 1. Deciding whether to use the MetroCluster Service Guide

This guide describes how to perform maintenance on the MetroCluster hardware components.

You should use this guide for maintaining the MetroCluster-specific hardware components in your MetroCluster configuration.

You can find other MetroCluster documentation in the following location:

ONTAP 9 Documentation Center

© Copyright Lenovo 2019, 2023

Chapter 2. Where to find procedures for MetroCluster maintenance tasks

You should be sure you select the correct procedure when you perform MetroCluster hardware maintenance tasks.

Note: With the release of ONTAP 9.8, the MetroCluster upgrade and expansion procedures have moved to the *MetroCluster Transition, Upgrade, and Expansion Guide.*MetroCluster Upgrade, Transition, and Expansion Guide

Component	MetroCluster type (FC or IP)	Task	Procedure
ONTAP software	Both	ONTAP software upgrade	Upgrade, Revert or Downgrade Guide
Controller module	Both	FRU replacement (including controller modules, PCIe cards, FC-VI card, and so on) **Note:** Moving a storage controller module or NVRAM card among the MetroCluster storage systems is not supported.	AFA and Hybrid Documentation Center
		Upgrade, transition, expansion	MetroCluster Upgrade, Transition, and Expansion Guide
Drive shelf	FC	Shelf addition (stack with bridges or individual shelf)	"Hot-adding a stack of SAS disk shelves to an existing pair of FibreBridge 7600N bridges" on page 85 "Hot-adding a stack of SAS disk shelves and bridges to a MetroCluster system " on page 91 "Hot-adding a SAS disk shelf to a stack of SAS disk shelves" on page 97
	FC	Shelf removal	Chapter 10 "Hot-removing storage from a MetroCluster FC configuration" on page 101
	FC	All other shelf maintenance procedures. The standard procedures can be used.	Lenovo Data Center Support
	IP	All shelf maintenance procedures. The standard procedures can be used. If adding shelves for an unmirrored aggregate, see Considerations when using unmirrored aggregates	Lenovo Data Center Support

© Copyright Lenovo 2019, 2023

Component	MetroCluster type (FC or IP)	Task	Procedure
FC-to-SAS bridge	FC	Bridge replacement	"Replacing a single FC-to-SAS bridge" on page 11 "Replacing a pair of FibreBridge 6500N bridges with 7600N bridge" on page 30 "Hot-swapping a FibreBridge 7500N with a 7600N bridge" on page 19
	FC	Firmware upgrade	"Updating firmware on FibreBridge 7600N bridge on configurations running ONTAP 9.4 and later" on page 8
	FC	Replacing a failed power supply module	"Hot-replacing a failed power supply module" on page 42
FC switch	FC	Switch upgrade	"Upgrading to new Brocade FC switches" on page 48
	FC	Switch replacement	"Replacing a Brocade FC switch (MetroCluster)" on page 52 "Replacing a Cisco FC switch" on page 61
	FC	Firmware upgrade	"Updating firmware on a Brocade FC switch" on page 45 "Updating firmware on a Cisco FC switch" on page 47
SAS optical cables	FC	Cable replacement	Lenovo Data Center Support
Cluster interconnect switches	FC	Switch replacement	Lenovo Data Center Support
	Both	Firmware upgrade	Lenovo Data Center Support

Chapter 3. MetroCluster failure and recovery scenarios

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

Note: For additional information about recovery from node failures, see the section "Choosing the correct recovery procedure" in the *MetroCluster Management and Disaster Recovery Guide*.

MetroCluster management and disaster recovery

Event	Impact	Recovery
Single node failure	A failover is triggered.	The configuration recovers through a local takeover. RAID is not impacted. Review system messages and replace failed FRUs as necessary.
Two nodes fail at one site	Two nodes will fail only if automated switchover is enabled in the MetroCluster Tiebreaker software.	Manual USO if automated switchover in MetroCluster Tiebreaker software is not enabled.
MetroCluster IP interface—failure of one port	The system is degraded. Additional port failure impacts HA mirroring.	The second port is used. Health Monitor generates an alert if the physical link to the port is broken. Review system messages and replace failed FRUs as necessary.
MetroCluster IP interface—failure of both ports	HA capability is impacted. RAID SyncMirror of the node stops syncing.	Immediate manual recovery is required as there is no HA takeover. Review system messages and replace failed FRUs as necessary.
Failure of one MetroCluster IP switch	No impact. Redundancy is provided through the second network.	Replace the failed switch as necessary. “Replacing an IP switch” on page 71
Failure of two MetroCluster IP switches that are in the same network	No impact. Redundancy is provided through the second network.	Replace the failed switch as necessary. “Replacing an IP switch” on page 71

© Copyright Lenovo 2019, 2023

Event	Impact	Recovery
Failure of two MetroCluster IP switches that are at one site	RAID SyncMirror of the node stops syncing. HA capability is impacted and the cluster goes out of quorum.	Replace the failed switch as necessary. “Replacing an IP switch” on page 71
Failure of two MetroCluster IP switches that are at different sites and not on the same network (diagonal failure)	RAID SyncMirror of the node stops syncing.	RAID SyncMirror of the node stops syncing. Cluster and HA capability are not impacted. Replace the failed switch as necessary. “Replacing an IP switch” on page 71

Chapter 4. Performing FC-to-SAS bridge maintenance

If necessary, you can nondisruptively replace the FC-to-SAS bridges or upgrade their firmware in the MetroCluster configuration.

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

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 7600N 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 7600N bridge.

Note: 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 7600N bridge with a single FibreBridge 7600N bridge	No	The FibreBridge 7600N bridge must be configured exactly the same as the FibreBridge 7600N bridge.	“Hot-swapping a FibreBridge 7500N with a 7600N bridge” on page 19
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 bridge” on page 24
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 7600N bridges can support up to four stacks.	“Hot-adding a stack of SAS disk shelves and bridges to a MetroCluster system ” on page 91

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.

About this task

Unlike the ATTO FibreBridge 7600N and 6500N bridges, the FibreBridge 7600N bridge is shipped with all IP port protocols and services disabled.

Starting 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.

Starting with ONTAP 9.8, *in-band monitoring* of the bridges is supported by default.

© Copyright Lenovo 2019, 2023

This task is required if you have FibreBridge 7600N bridges and you are running ONTAP 9.5 or later and are **not** planning on using in-band management to manage the bridges. Therefore, you need to configure the bridge via the Ethernet management port.

Step 1. Access the bridge's console interface by connecting a serial cable to the serial port on the FibreBridge 7600N bridge.

Step 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.

Step 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.

Step 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.

Step 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.

Updating firmware on FibreBridge 7600N bridge 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 bridges on configurations running ONTAP 9.4 and later.

Before you begin

- 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 server.
- You must be using a supported firmware version.

About this task

- You can use this task only on FibreBridge 7600N bridge 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.

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

Note: Starting 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.

Step 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`

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

Step 3. 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.

Step 4. 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 (*>).

Step 5. Update the bridge firmware by using the following command:
`storage bridge firmware update -bridge` *name* `-uri` *URL-of-firmware-package*

Example

```
cluster_A> storage bridge firmware update -bridge bridge_A_1a -uri http://192.168.132.97/firmware.spf
```

Step 6. Return to the admin privilege level:
`set -privilege admin`

Step 7. Verify that the firmware upgrade is complete:
`job show -name` *Storage bridge firmware update*

Example

The following example shows that the job is still running:

```
cluster_A> job show -name storage bridge firmware update
Owning

Job ID Name                 Vserver    Node           State
------ -------------------- ---------- -------------- ----------
2246   job-name
                            cluster_A
                                       node_A_1
                                                      Running
Description: Storage bridge firmware update job
```

Result

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

```
cluster_A> job show -name Storage bridge firmware update

                            Owning
```

```
Job ID Name                  Vserver    Node           State
------ -------------------- ---------- -------------- ----------
2246   Storage bridge firmware update job name
                            cluster_A
                                       node_A_1
                                                      Success
Description: Storage bridge firmware update job
```

Step 8. 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:

1. Run the `flashimages` command on the console of the bridge and confirm that the correct firmware versions are displayed.

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

   ```
   flashimages

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

2. 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:

1. Run the `storage bridge run-cli -name` *bridge-name* `-command FlashImages` command.

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

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

   [Job 2257]

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


   [Job 2257] Job succeeded.
   ```

2. If necessary, restart the bridge:
   `storage bridge run-cli -name ATTO_7600N_IB_1 -command FirmwareRestart`

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

Step 9. 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).

Example

```
cluster_A> node run -node cluster_A-01 -command sysconfig
Lenovo 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
```

Step 10. Verify that the FibreBridge firmware was updated:
`storage bridge show -fields fw-version,symbolic-name`

Example

```
cluster_A> storage 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.
```

Step 11. Verify that both 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.

Example

```
Ready.
flashimages

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

 Ready.
```

Step 12. Repeat steps 5 to 10 to ensure that both flash images are updated to the same version.

Step 13. Verify that both flash images are updated to the same version.`flashimages`The output should show the same version for both partitions.

Example

```
Ready.
flashimages

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

 Ready.
```

Step 14. 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" on page 42

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.

Step 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:

Example

```
node_A_1> run local sysconfig -v
Lenovo Release 9.6P5: Fri Dec 13 13:23:47 EST 2019
System ID: 4068741258 (node_A_1); partner ID: 4068741260 (node_B_1)
System Serial Number: 940001025471 (node_A_1)
System Rev: A4
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.

Note: Starting 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.

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

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

Example

Enter the command: `storage bridge run-cli -name bridge_A1 -command "info"`

```
info

Device Status          = Good
Unsaved Changes        = None
Device                 = "FibreBridge 7500N"
Serial Number          = FB7500N100000
Device Version         = 3.10
Board Revision         = 7
Build Number           = 007A
Build Type             = Release
Build Date             = "Aug 20 2019" 11:01:24
Flash Revision         = 0.02
Firmware Version       = 3.10
BCE Version (FPGA 1)   = 15
BAU Version (FPGA 2)   = 33
User-defined name      = "bridgeA1"
World Wide Name        = 20 00 00 10 86 A1 C7 00
MB of RAM Installed    = 512
FC1 Node Name          = 20 00 00 10 86 A1 C7 00
```

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

Ready.
```

Enter the command: `storage bridge run-cli -name bridge_A1 -command "sasportlist"`

```
SASPortList

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

Step 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.

Step 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`

Step 4. Properly ground yourself.

Step 5. Power down the ATTO bridge.

If you are using a...	Then...
FibreBridge 7600N bridge	Remove the power cables connected to the bridge.

Step 6. Disconnect the cables that are connected to the old bridge.
You should make note of the port to which each cable was connected.

Step 7. Remove the old bridge from the rack.

Step 8. Install the new bridge into the rack.

Step 9. Reconnect the power cord and, if configuring for IP access to the bridge, a shielded Ethernet cable.

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

Step 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.

Step 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.

Step 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.

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.

Step 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.

Note: 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, enable SNMP on the bridge:
`set SNMP enabled`

If running ONTAP 9.5 or later, ONTAP will use in-band monitoring via the FC connection to the bridge, so you should leave SNMP disabled.

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

Note: 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.

Note: 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*

Note: 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.

Step 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`

   Note: 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:
   `SaveConfigurationFirmwareRestart`You are prompted to restart the bridge.

Step 16. 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" on page 8.

Step 17. Reconnect the SAS and FC cables to the same ports on the new bridge.

If the new bridge is a FibreBridge 7600N, you must replace the cables connecting the bridge to the top or bottom of the shelf stack. The FibreBridge 6500N bridge used SAS cables; the FibreBridge 7600N bridge require mini-SAS cables for these connections.

Note: 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.

Step 18. 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	1. 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. 2. 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. 3. Click **Advanced** in the menu. 4. View the connected devices: `sastargets` 5. 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.

Note: 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.

Example

The following output shows that 10 disks are connected:

```
Tgt VendorID ProductID        Type        SerialNumber
  0 NETAPP   X410_S15K6288A15 DISK        3QP1CLE300009940UHJV
  1 NETAPP   X410_S15K6288A15 DISK        3QP1ELF600009940V1BV
  2 NETAPP   X410_S15K6288A15 DISK        3QP1G3EW00009940U2M0
  3 NETAPP   X410_S15K6288A15 DISK        3QP1EWMP00009940U1X5
  4 NETAPP   X410_S15K6288A15 DISK        3QP1FZLE00009940G8YU
  5 NETAPP   X410_S15K6288A15 DISK        3QP1FZLF00009940TZKZ
  6 NETAPP   X410_S15K6288A15 DISK        3QP1CEB400009939MGXL
  7 NETAPP   X410_S15K6288A15 DISK        3QP1G7A900009939FNTT
  8 NETAPP   X410_S15K6288A15 DISK        3QP1FYOT00009940G8PA
  9 NETAPP   X410_S15K6288A15 DISK        3QP1FXW600009940VERQ
```

Step 19. 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 Step 18 on page 17 for each remaining bridge.
Not correct	1. Check for loose SAS cables or correct the SAS cabling by repeating Step Step 17 on page 17. 2. Repeat Step Step 18 on page 17.

Step 20. 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.

Step 21. 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`

Note: 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.

Example

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

```
Lenovo Release 9.6P5: Fri Dec 13 13:23:47 EST 2019
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: A4
System Storage Configuration: Multi-Path HA
System ACP Connectivity: Inband Active
```

Attention: 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.

Step 22. 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`

Step 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	`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.

Example

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

```
controller_A_1::> storage bridge show

Bridge               Symbolic Name Is Monitored  Monitor Status  Vendor Model                Bridge WWN
------------------  ------------- -----------  --------------  ------ -----------------    ----------
ATTO_10.10.20.10  atto01        true           ok              Atto   FibreBridge 7600N   20000010867038c0
ATTO_10.10.20.11  atto02        true           ok              Atto   FibreBridge 7600N   20000010867033c0
ATTO_10.10.20.12  atto03        true           ok              Atto   FibreBridge 7600N   20000010867030c0
ATTO_10.10.20.13  atto04        true           ok              Atto   FibreBridge 7600N   2000001086703b80

4 entries were displayed

 controller_A_1::>
```

Step 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`

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.

Note: Starting 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.

Step 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.

Step 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`

Step 3. Properly ground yourself.

Step 4. Remove the power cables connected to the bridge to power down the bridge.

Step 5. Disconnect the cables that are connected to the old bridge.
You should make note of the port to which each cable was connected.

Step 6. Remove the old bridge from the rack.

Step 7. Install the new bridge into the rack.

Step 8. Reconnect the power cord and shielded Ethernet cable.

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

Step 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.

Step 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.

Step 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.

Step 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.

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.

Step 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.

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

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

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

   If using the CLI, you must run the following commands:
   `set ipaddress mp1` *ip-address*
   `set ipsubnetmask mp1` *subnet-mask*
   `set ipgateway mp1 x.x.x.x`
   `set ipdhcp mp1 disabled`
   `set ethernetspeed mp1 1000`

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

   Example bridge names for one stack group on each site:

   - bridge_A_1a
   - bridge_A_1b
   - bridge_B_1a
   - bridge_B_1b

   If using the CLI, you must run the following command:
   `set bridgename"` *bridgename*

c. If running ONTAP 9.4, enable SNMP on the bridge:
   `set SNMP enabled`
   If running ONTAP 9.5 or later, ONTAP will use in-band monitoring via the FC connection to the bridge, so you should leave SNMP disabled.

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

   Note: 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*

Note: 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.
```

Step 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`

Note: 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:
`SaveConfigurationFirmwareRestart`You are prompted to restart the bridge.

Step 16. Update the FibreBridge firmware on each bridge.

"Updating firmware on FibreBridge 7600N bridge on configurations running ONTAP 9.4 and later" on page 8

Step 17. Reconnect the SAS and FC cables to the same ports on the new bridge.

Note: 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.

Step 18. 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.

Example

The following output shows that 10 disks are connected:

```
Tgt VendorID ProductID        Type        SerialNumber
  0 NETAPP   X410_S15K6288A15 DISK        3QP1CLE300009940UHJV
  1 NETAPP   X410_S15K6288A15 DISK        3QP1ELF600009940V1BV
  2 NETAPP   X410_S15K6288A15 DISK        3QP1G3EW00009940U2M0
  3 NETAPP   X410_S15K6288A15 DISK        3QP1EWMP00009940U1X5
  4 NETAPP   X410_S15K6288A15 DISK        3QP1FZLE00009940G8YU
  5 NETAPP   X410_S15K6288A15 DISK        3QP1FZLF00009940TZKZ
  6 NETAPP   X410_S15K6288A15 DISK        3QP1CEB400009939MGXL
```

```
 7 NETAPP   X410_S15K6288A15 DISK        3QP1G7A900009939FNTT
 8 NETAPP   X410_S15K6288A15 DISK        3QP1FYOT0000994OG8PA
 9 NETAPP   X410_S15K6288A15 DISK        3QP1FXW600009940VERQ
```

Step 19. 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	1. Check for loose SAS cables or correct the SAS cabling by repeating Step Step 17 on page 17. 2. Repeat the previous step.

Step 20. 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.

Step 21. 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`

Note: 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.

Example

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

```
Lenovo Release 9.6P5: Fri Dec 13 13:23:47 EST 2019
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: A4
System Storage Configuration: Multi-Path HA
System ACP Connectivity: Inband Active
```

Attention: 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.

Step 22. 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`

Step 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	`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.

Example

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

```
controller_A_1::> storage bridge show

Bridge               Symbolic Name Is Monitored  Monitor Status  Vendor Model                Bridge WWN
------------------  ------------- ------------  --------------  ------ -----------------    ----------
ATTO_10.10.20.10  atto01        true          ok              Atto   FibreBridge 7600N   20000010867038c0
ATTO_10.10.20.11  atto02        true          ok              Atto   FibreBridge 7600N   20000010867033c0
ATTO_10.10.20.12  atto03        true          ok              Atto   FibreBridge 7600N   2000001086703Oc0
ATTO_10.10.20.13  atto04        true          ok              Atto   FibreBridge 7600N   2000001086703b80

4 entries were displayed

 controller_A_1::>
```

Step 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`

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

You can hot-swap a FibreBridge 6500N bridge with a FibreBridge 7600N 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 bridge.
- When you hot-swap a FibreBridge 6500N bridge with a FibreBridge 7600N bridge, you must use only one FC port and one SAS port on the FibreBridge 7600N 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.

Attention: If you are hot-swapping both FibreBridge 6500N bridges in a pair, you must use the "Replacing a pair of FibreBridge 6500N bridges with 7600N bridge" on page 30 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 bridge.

Note: Starting 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.

Step 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.

Step 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`

Step 3. Properly ground yourself.

Step 4. Turn off the power switch of the bridge.

Step 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.

Step 6. Remove the FibreBridge 6500N bridge that you need to replace from the rack.

Step 7. Install the new FibreBridge 7600N bridge into the rack.

Step 8. Reconnect the power cord and, if necessary, the shielded Ethernet cable.

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

Step 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.

Step 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.

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.

Step 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.

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.

Step 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.

Note: 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, enable SNMP on the bridge:
   `set SNMP enabled`

   If running ONTAP 9.5 or later, ONTAP will use in-band monitoring via the FC connection to the bridge, so you should leave SNMP disabled.

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

   Note: 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.

Note: 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*

Note: 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.

Step 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`

Note: 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:
`SaveConfigurationFirmwareRestart`You are prompted to restart the bridge.

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

Step 16. 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" on page 8.

Step 17. Reconnect the SAS and FC cables to the SAS A and Fibre Channel 1 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.

The FC port must be cabled to the same switch or controller port that the FibreBridge 6500N bridge had been connected to.

Note: 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.

Step 18. Verify that the bridge can see all of the disk drives and disk shelves it is connected to.

If you are using the...	Then...
ATTO ExpressNAV GUI	1. 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. 2. 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. 3. Click **Advanced** in the menu. 4. 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`

Example

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

```
Tgt VendorID ProductID        Type        SerialNumber
  0 NETAPP   X410_S15K6288A15 DISK        3QP1CLE300009940UHJV
  1 NETAPP   X410_S15K6288A15 DISK        3QP1ELF600009940V1BV
  2 NETAPP   X410_S15K6288A15 DISK        3QP1G3EW00009940U2M0
  3 NETAPP   X410_S15K6288A15 DISK        3QP1EWMP00009940U1X5
  4 NETAPP   X410_S15K6288A15 DISK        3QP1FZLE00009940G8YU
  5 NETAPP   X410_S15K6288A15 DISK        3QP1FZLF00009940TZKZ
  6 NETAPP   X410_S15K6288A15 DISK        3QP1CEB400009939MGXL
  7 NETAPP   X410_S15K6288A15 DISK        3QP1G7A900009939FNTT
  8 NETAPP   X410_S15K6288A15 DISK        3QP1FYOT00009940G8PA
  9 NETAPP   X410_S15K6288A15 DISK        3QP1FXW600009940VERQ
```

Note: 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.

Step 19. 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 Step Step 18 on page 27 for each remaining bridge.
Not correct	1. Check for loose SAS cables or correct the SAS cabling by repeating Step Step 17 on page 27. 2. Repeat Step Step 18 on page 27.

Step 20. Reenable the FC switch port that connects to the bridge.

Step 21. 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`

Note: 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:

```
Lenovo Release 9.6P5: Fri Dec 13 13:23:47 EST 2019
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: A4
System Storage Configuration: Multi-Path HA
System ACP Connectivity: Inband Active
```

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.

Attention: 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.

Step 22. 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	`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.

Example

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

```
controller_A_1::> storage bridge show

Bridge               Symbolic Name Is Monitored  Monitor Status  Vendor Model                 Bridge WWN
------------------  ------------- ------------  --------------  ------ ------------------    ----------
ATTO_10.10.20.10  atto01        true          ok              Atto   FibreBridge 7600N   20000010867038c0
ATTO_10.10.20.11  atto02        true          ok              Atto   FibreBridge 7600N   20000010867033c0
ATTO_10.10.20.12  atto03        true          ok              Atto   FibreBridge 7600N   20000010867030c0
ATTO_10.10.20.13  atto04        true          ok              Atto   FibreBridge 7600N   2000001086703b80

4 entries were displayed

 controller_A_1::>
```

Step 23. 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`

Step 24. Return the failed part to Lenovo as described in the RMA instructions shipped with the kit. Contact technical support at Lenovo Support if you need the RMA number or additional help with the replacement procedure.

Replacing a pair of FibreBridge 6500N bridges with 7600N bridge

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

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 bridge.

  After the replacement, both bridges in the pair must be the same model.
- You previously replaced a single FibreBridge 6500N bridge with a 7600N bridge and are now replacing the second bridge in the pair.
- You have a pair of FibreBridge 7600N bridge 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.

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.

Step 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:

Example

```
node_A_1> run local sysconfig -v
Lenovo Release 9.6P5: Fri Dec 13 13:23:47 EST 2019
```

```
System ID: 4068741258 (node_A_1); partner ID: 4068741260 (node_B_1)
System Serial Number: 940001025471 (node_A_1)
System Rev: A4
System Storage Configuration: Multi-Path HA<=== Configuration should be multi-path HA
.
.
.
slot 0: FC Host Adapter 0g (QLogic 8324 rev. 2, N-port, <UP>)<=== Initiator port
Firmware rev:      7.5.0
Flash rev:         0.0.0
Host Port Id:      0x60130
FC Node Name:       5:00a:098201:bae312
FC Port Name:      5:00a:098201:bae312
SFP Vendor:        UTILITIES CORP.
SFP Part Number:   FTLF8529P3BCVAN1
SFP Serial Number: URQ0Q9R
SFP Capabilities:  4, 8 or 16 Gbit
Link Data Rate:    16 Gbit
Switch Port:       brcd6505-fcs40:1
  <List of disks visible to port>
 ID     Vendor   Model            FW    Size
brcd6505-fcs29:12.126L1527     : NETAPP   X302_HJUPI01TSSM NA04 847.5GB (1953525168 512B/sect)
brcd6505-fcs29:12.126L1528     : NETAPP   X302_HJUPI01TSSA NA02 847.5GB (1953525168 512B/sect)
.
.
.
<List of FC-to-SAS bridges visible to port>
FC-to-SAS Bridge:
brcd6505-fcs40:12.126L0        : ATTO     FibreBridge6500N 1.61  FB6500N102980
brcd6505-fcs42:13.126L0        : ATTO     FibreBridge6500N 1.61  FB6500N102980
brcd6505-fcs42:6.126L0         : ATTO     FibreBridge6500N 1.61  FB6500N101167
brcd6505-fcs42:7.126L0         : ATTO     FibreBridge6500N 1.61  FB6500N102974
.
.
.
  <List of storage shelves visible to port>
brcd6505-fcs40:12.shelf6: DS4243  Firmware rev. IOM3 A: 0200  IOM3 B: 0200
brcd6505-fcs40:12.shelf8: DS4243  Firmware rev. IOM3 A: 0200  IOM3 B: 0200
.
.
.
```

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

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

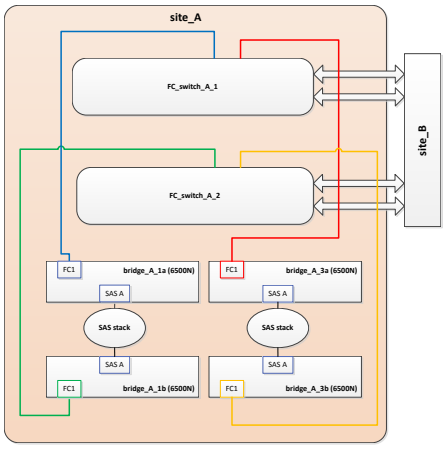
About this task

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 bridge.

  In this case, you should begin with Step 3 Step 3 on page 33.

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

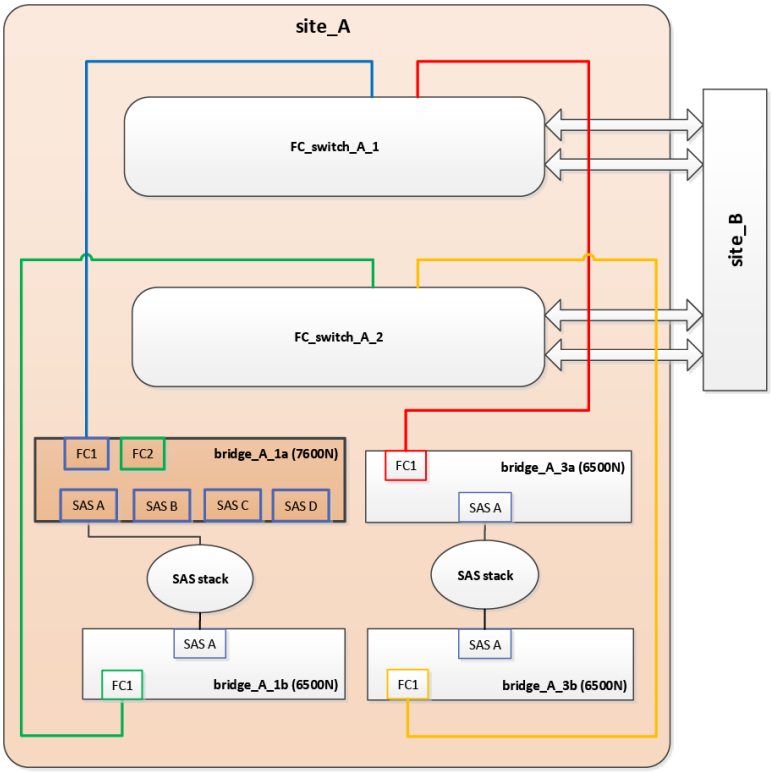


Step 1. Using the following guidelines, hot-swap the “top” FibreBridge 6500N bridge with a FibreBridge 7600N bridge using the procedure in “Hot-swapping a FibreBridge 6500N bridge with a FibreBridge 7600N bridge” on page 24:

- Connect the FibreBridge 7600N 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 bridge FC2 port at this time.

The following diagram shows that bridge_A_1a has been replaced and is now a FibreBridge 7600N bridge:



Step 2. Confirm connectivity to the bridge-connected disks and that the new FibreBridge 7600N is visible in the configuration:

`run local sysconfig -v`

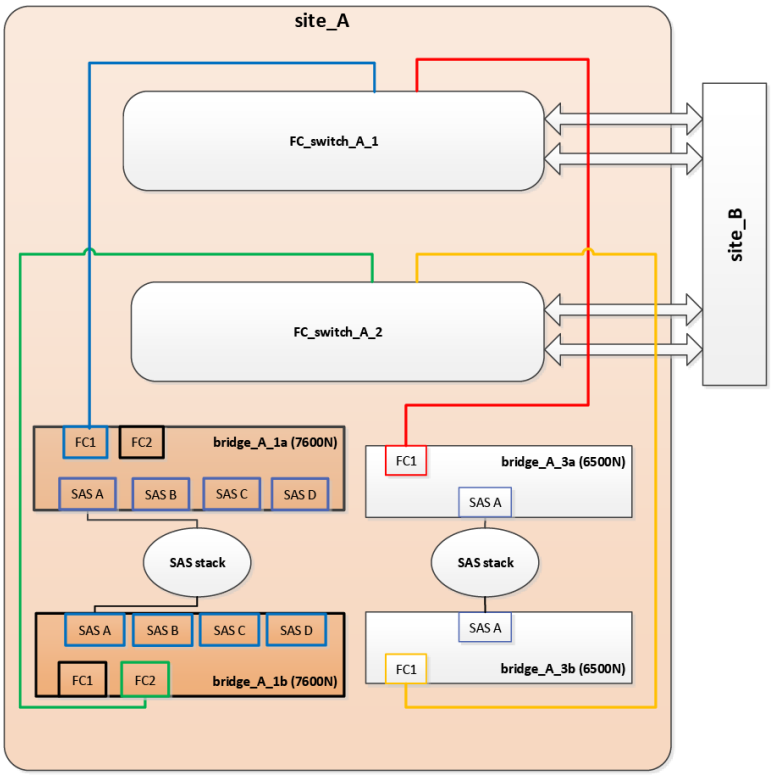
Example

```
node_A_1> run local sysconfig -v
Lenovo Release 9.6P5: Fri Dec 13 13:23:47 EST 2019
System ID: 0536872165 (node_A_1); partner ID: 0536872141 (node_B_1)
System Serial Number: 940001025465 (node_A_1)
System Rev: A4
System Storage Configuration: Multi-Path HA<=== Configuration should be multi-path HA
.
.
.
slot 0: FC Host Adapter 0g (QLogic 8324 rev. 2, N-port, <UP>)<=== Initiator port
Firmware rev:       7.5.0
Flash rev:          0.0.0
Host Port Id:       0x60100
FC Node Name:        5:00a:098201:bae312
FC Port Name:       5:00a:098201:bae312
SFP Vendor:         FINISAR CORP.
SFP Part Number:    FTLF8529P3BCVAN1
SFP Serial Number: URQ0R1R
SFP Capabilities:  4, 8 or 16 Gbit
Link Data Rate:     16 Gbit
Switch Port:        brcd6505-fcs40:1
  <List of disks visible to port>
 ID      Vendor    Model             FW     Size
brcd6505-fcs40:12.126L1527     : NETAPP   X302_HJUPI01TSSM NA04 847.5GB (1953525168 512B/sect)
brcd6505-fcs40:12.126L1528     : NETAPP   X302_HJUPI01TSSA NA02 847.5GB (1953525168 512B/sect)
.
.
.
<List of FC-to-SAS bridges visible to port>
FC-to-SAS Bridge:
brcd6505-fcs40:12.126L0        : ATTO     FibreBridge7600N A30H  FB7600N100104<===
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
.
.
.
```

Step 3. Using the following guidelines, hot-swap the “bottom” FibreBridge 6500N bridge with a FibreBridge 7600N bridge using the procedure in “Hot-swapping a FibreBridge 6500N bridge with a FibreBridge 7600N bridge” on page 24:

- Connect the FibreBridge 7600N 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 bridge FC1 port at this time.



Step 4. Confirm connectivity to the bridge-connected disks:

`run local sysconfig -v`

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

Example

```
node_A_1> run local sysconfig -v
Lenovo Release 9.6P5: Fri Dec 13 13:23:47 EST 2019
System ID: 0536872165 (node_A_1); partner ID: 0536872141 (node_B_1)
System Serial Number: 940001025465 (node_A_1)
System Rev: A4
System Storage Configuration: Multi-Path HA<=== Configuration should be multi-path HA
.
.
.
slot 0: FC Host Adapter 0g (QLogic 8324 rev. 2, N-port, <UP>)<=== Initiator port
Firmware rev:       7.5.0
Flash rev:          0.0.0
Host Port Id:       0x60100
FC Node Name:        5:00a:098201:bae312
FC Port Name:        5:00a:098201:bae312
SFP Vendor:         FINISAR CORP.
SFP Part Number:    FTLF8529P3BCVAN1
SFP Serial Number: URQ0R1R
SFP Capabilities:  4, 8 or 16 Gbit
Link Data Rate:     16 Gbit
Switch Port:        brcd6505-fcs40:1
  <List of disks visible to port>
 ID      Vendor   Model            FW    Size
brcd6505-fcs40:12.126L1527     : NETAPP   X302_HJUPI01TSSM NA04 847.5GB (1953525168 512B/sect)
brcd6505-fcs40:12.126L1528     : NETAPP   X302_HJUPI01TSSA NA02 847.5GB (1953525168 512B/sect)
.
.
.
```

```
<List of FC-to-SAS bridges visible to port>
FC-to-SAS Bridge:
brcd6505-fcs40:12.126LO         : ATTO     FibreBridge7600N A30H  FB7600N100104
brcd6505-fcs42:13.126LO         : ATTO     FibreBridge7600N A30H  FB7600N100104
.
.
.
  <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 bridge

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

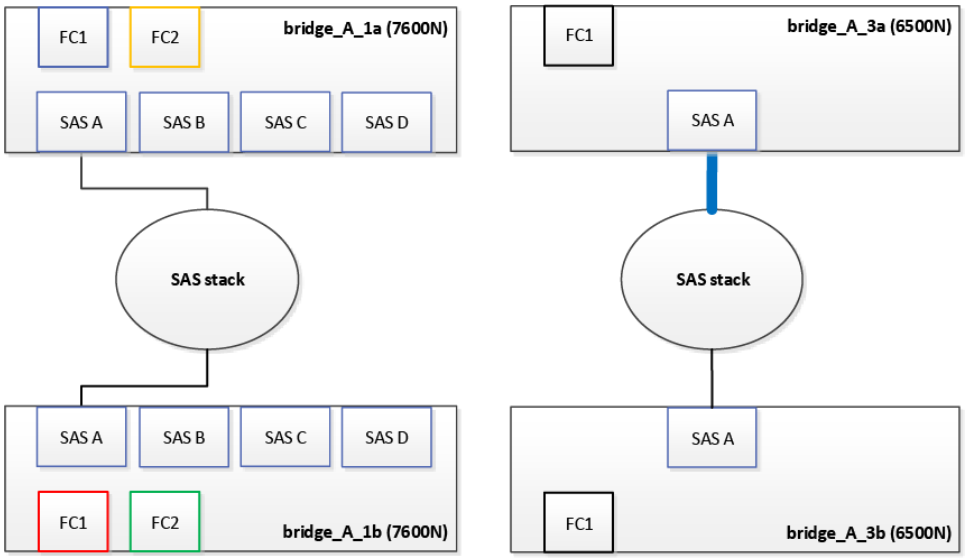
About this task

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

Attention: 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.

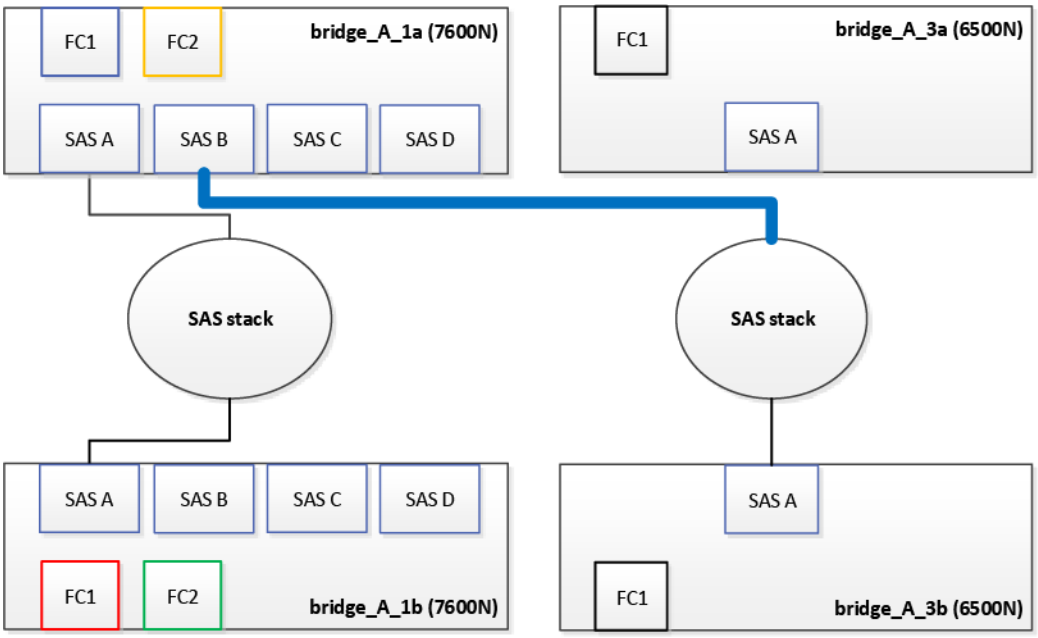
Note: Wait at least 10 seconds before connecting the port. The SAS cable connectors are keyed; when oriented correctly into a SAS port, the connector clicks into place and the disk shelf SAS port LNK LED illuminates green. For disk shelves, you insert a SAS cable connector with the pull tab oriented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

Step 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:

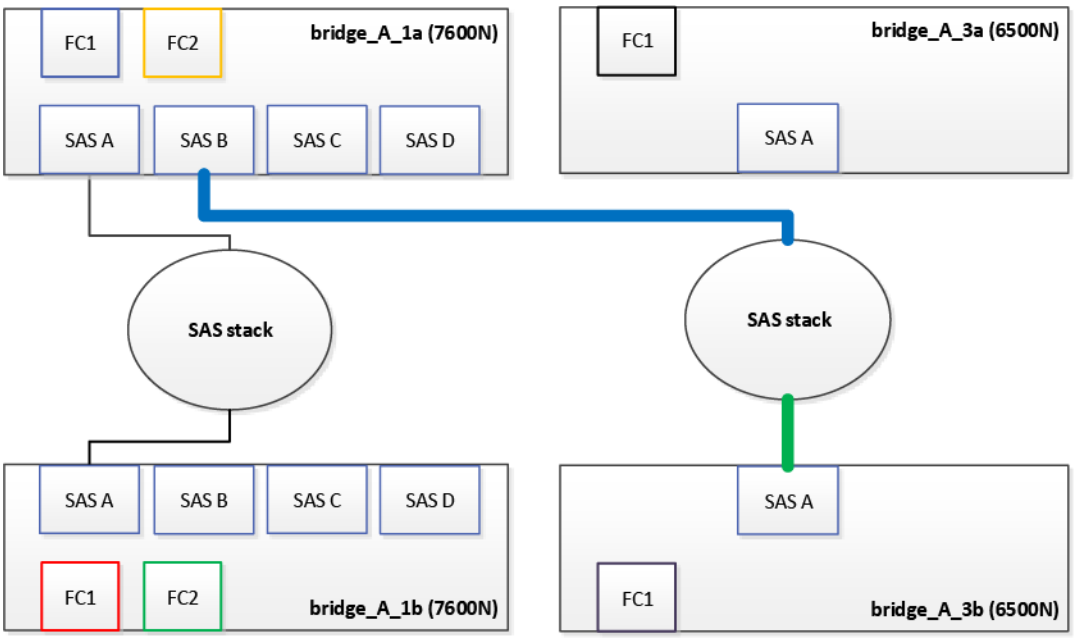


Step 2. Using a cable with a mini-SAS connector, connect the same SAS port on the storage shelf to the SAS B port of the top FibreBridge 7600N bridge.

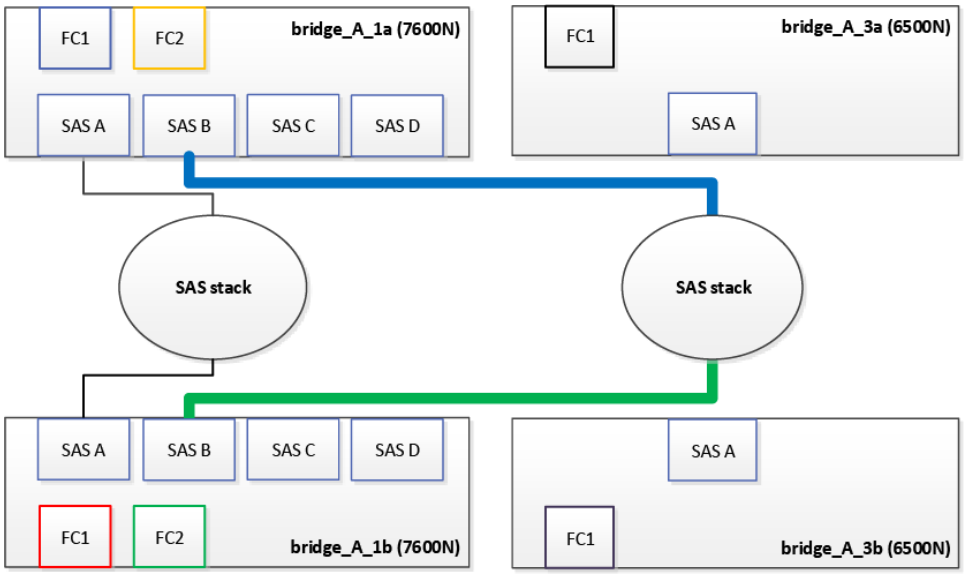
The cable is shown in blue in the following example:



Step 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:



Step 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 bridge.
This cable is shown in green in the following example:



Step 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:

Example

```
node_A_1> run local sysconfig -v
Lenovo Release 9.6P5: Fri Dec 13 13:23:47 EST 2019
System ID: 0536872165 (node_A_1); partner ID: 0536872141 (node_B_1)
System Serial Number: 940001025465 (node_A_1)
System Rev: A4
System Storage Configuration: Multi-Path HA<=== Configuration should be multi-path HA
.
.
.
slot 0: FC Host Adapter 0g (QLogic 8324 rev. 2, N-port, <UP>)<=== Initiator port
Firmware rev:       7.5.0
Flash rev:          0.0.0
Host Port Id:       0x60100
FC Node Name:        5:00a:098201:bae312
FC Port Name:        5:00a:098201:bae312
SFP Vendor:         FINISAR CORP.
SFP Part Number:    FTLF8529P3BCVAN1
SFP Serial Number: URQ0R1R
SFP Capabilities:  4, 8 or 16 Gbit
Link Data Rate:     16 Gbit
Switch Port:        brcd6505-fcs40:1
  <List of disks visible to port>
 ID      Vendor    Model             FW     Size
brcd6505-fcs40:12.126L1527     : NETAPP   X302_HJUPI01TSSM NA04 847.5GB (1953525168 512B/sect)
brcd6505-fcs40:12.126L1528     : NETAPP   X302_HJUPI01TSSA NA02 847.5GB (1953525168 512B/sect)
.
.
.
<List of FC-to-SAS bridges visible to port>
FC-to-SAS Bridge:
brcd6505-fcs40:12.126L0        : ATTO     FibreBridge7600N A30H  FB7600N100104
brcd6505-fcs42:13.126L0        : ATTO     FibreBridge7600N A30H  FB7600N100104
.
.
.
  <List of storage shelves visible to port>
```

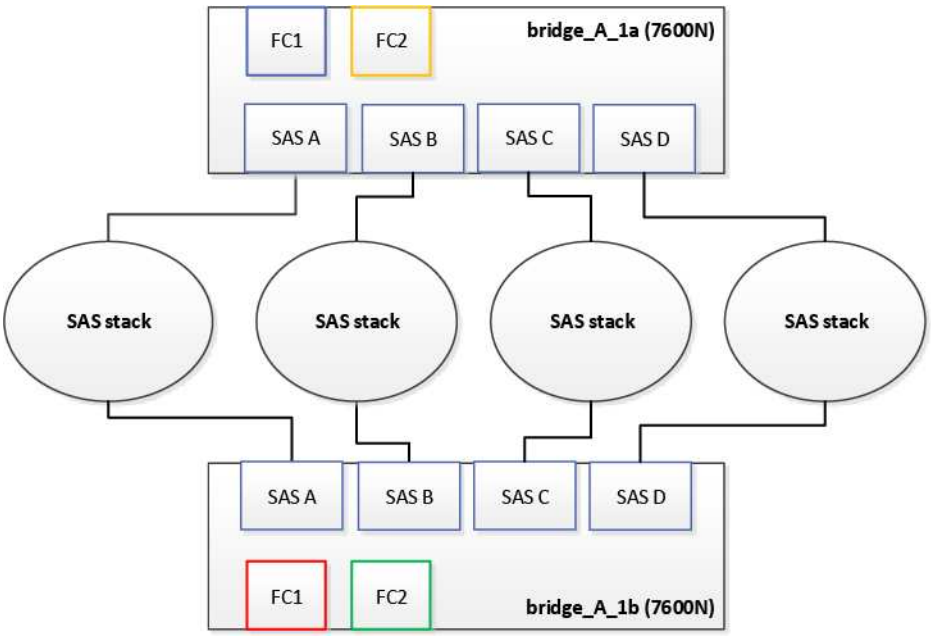
```
brcd6505-fcs40:12.shelf6: DS4243  Firmware rev. IOM3 A: 0200  IOM3 B: 0200
brcd6505-fcs40:12.shelf8: DS4243  Firmware rev. IOM3 A: 0200  IOM3 B: 0200
.
.
.
```

Step 6. Remove the old FibreBridge 6500N bridges that are no longer connected to the SAS storage.

Step 7. Wait two minutes for the system to recognize the changes.

Step 8. If the system was miscabled, remove the cable, correct the cabling, and then reconnect the correct cable.

Step 9. If necessary, repeat the preceding steps to move up to two additional SAS stacks behind the new FibreBridge 7600N bridge, 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 bridge to a configuration

The zoning must be changed when you are replacing FibreBridge 6500N bridges with FibreBridge 7600N bridge and using both FC ports on the FibreBridge 7600N bridge.

Updating zoning when adding FibreBridge 7600N bridge to a configuration

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

About this task

- FibreBridge 7600N bridges are supported in ONTAP 9.6 and later.
- The specific zoning in this task is for ONTAP 9.4 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 FC2 ports.
The zones for the new FC2 ports will contain the initiator ports removed from the existing zones.

The FC switch section of the *Fabric-attached MetroCluster Installation and Configuration Guide* contains details about the zoning commands.

Fabric-attached MetroCluster installation and configuration

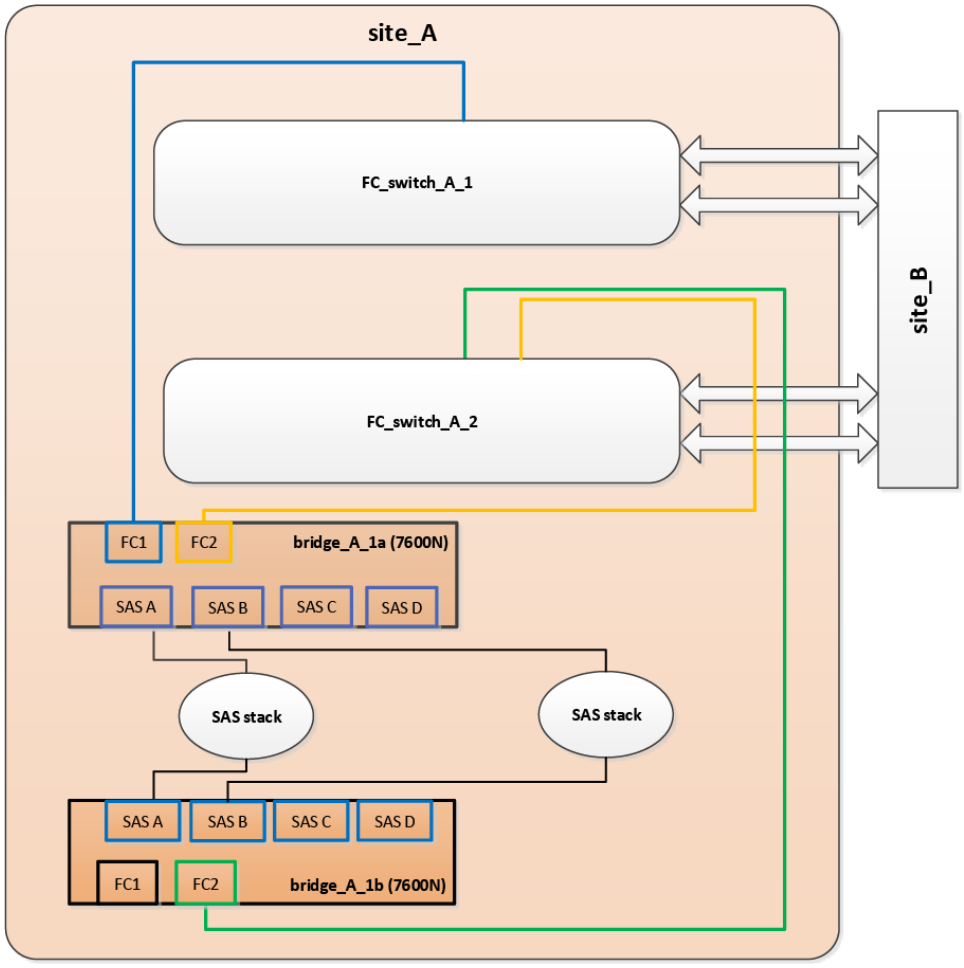
Cabling the second bridge FC port when adding FibreBridge 7600N bridge to a configuration

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

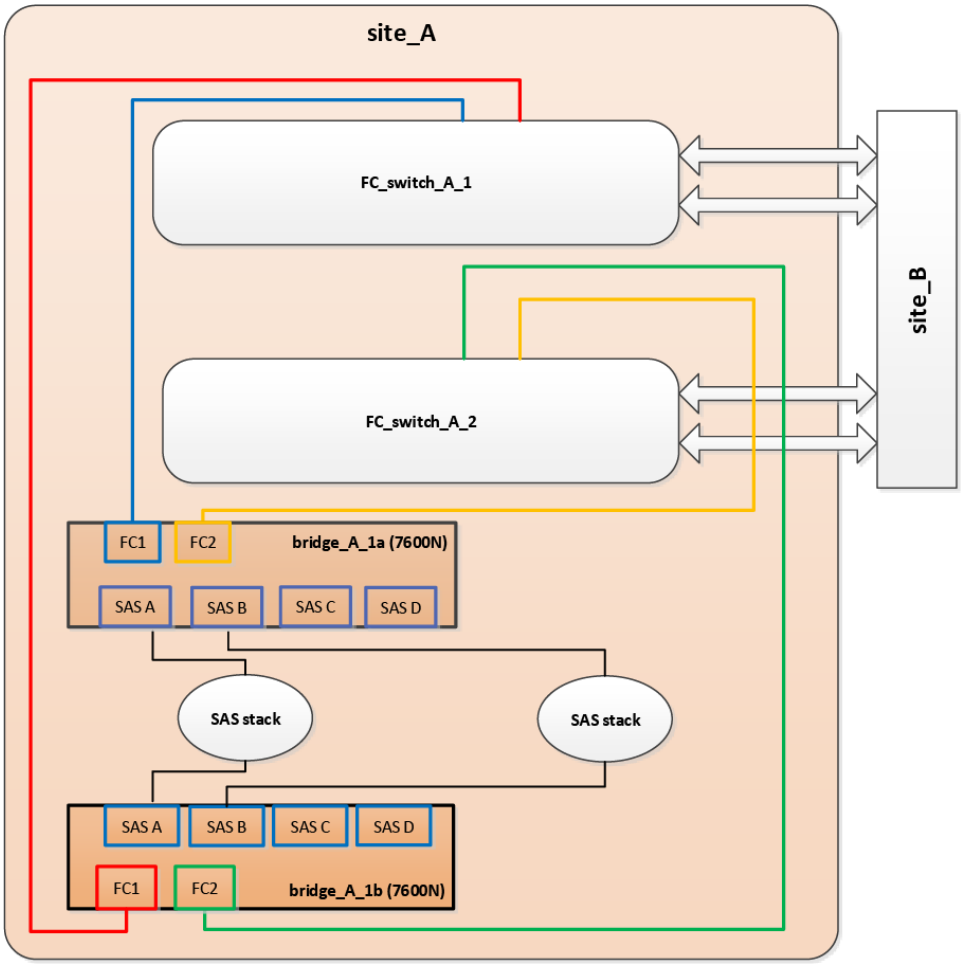
Before you begin

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

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



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



Step 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:

Example

```
node_A_1> run local sysconfig -v
Lenovo Release 9.6P5: Fri Dec 13 13:23:47 EST 2019
System ID: 0536872165 (node_A_1); partner ID: 0536872141 (node_B_1)
System Serial Number: 940001025465 (node_A_1)
System Rev: A4
System Storage Configuration: Multi-Path HA<=== Configuration should be multi-path HA
.
.
.
slot 0: FC Host Adapter 0g (QLogic 8324 rev. 2, N-port, <UP>)<=== Initiator port
Firmware rev:       7.5.0
Flash rev:          0.0.0
Host Port Id:       0x60100
FC Node Name:       5:00a:098201:bae312
FC Port Name:       5:00a:098201:bae312
SFP Vendor:         FINISAR CORP.
SFP Part Number:    FTLF8529P3BCVAN1
SFP Serial Number: URQ0R1R
```

```
SFP Capabilities:  4, 8 or 16 Gbit
Link Data Rate:    16 Gbit
Switch Port:       brcd6505-fcs40:1
  <List of disks visible to port>
 ID     Vendor   Model            FW    Size
brcd6505-fcs40:12.126L1527     : NETAPP   X302_HJUPI01TSSM NA04 847.5GB (1953525168 512B/sect)
brcd6505-fcs40:12.126L1528     : NETAPP   X302_HJUPI01TSSA NA02 847.5GB (1953525168 512B/sect)
.
.
.
<List of FC-to-SAS bridges visible to port>
FC-to-SAS Bridge:
brcd6505-fcs40:12.126L0        : ATTO     FibreBridge7600N A30H  FB7600N100104
brcd6505-fcs42:13.126L0        : ATTO     FibreBridge7600N A30H  FB7600N100104
.
.
.
  <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.

Step 1. Disable unused SAS ports on the top FC-to-SAS bridge:

a. Log in to the bridge CLI.

b. Disable any unused ports.

Note: If you have configured an ATTO 7600N 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*

Example

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`

Example

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.
```

Step 2. Repeat Step 1 on page 41 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.

Note: 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.

Note: 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 7600N bridges, *in-band management* of the bridges is supported as an alternative to IP management of the bridges.

Note: Starting 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.

Starting with ONTAP 9.8, bridges are monitored via in-band connections by default.

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.

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

Managing a FibreBridge bridge from ONTAP

Starting 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

Note: Starting 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.

Step 1. Run the applicable FibreBridge command within the `storage bridge run-cli` command: `storage bridge run-cli -name` *bridge-name* `-command "`*command-text*`"`

Example

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 9.6 you can secure the bridge. This disables the bridge's Ethernet ports. You can also reenable Ethernet access.

About this task

- 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.

Note: Starting 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.

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

Example

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
```

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

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

Example

The following command secures bridge_A_1:

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

Step 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`

Example

The following command secures bridge_A_1:

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

Chapter 5. Performing FC switch maintenance and replacement

If necessary, you can non-disruptively replace the FC switches or upgrade their firmware in the MetroCluster configuration.

Upgrading or downgrading 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.

Before you begin

- You must have the firmware files.
- The system must be properly cabled.
- All paths to the storage shelves must be available.
- The disk shelf stacks must be stable.
- The FC switch fabric must be healthy.
- No failed components can be present in the system.
- The system must be operating normally.
- You must have the admin password and access to an FTP or SCP server.

About this task

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

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

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

Step 1. Log in to each of the switches in the fabric.
The examples in the following steps use the switch `FC_switch_A_1`.

Step 2. Disable each of the switches in the fabric:
`switchCfgPersistentDisable`
If this command is not available, then run the `switchDisable` command.

Example

```
FC_switch_A_1:admin> switchCfgPersistentDisable
```

Step 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.

Example

© Copyright Lenovo 2019, 2023

```
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) [1]: 2
Password:
Server IP: 10.64.203.188, Protocol IPv4
Checking system settings for firmwaredownload...
System settings check passed.
```

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

Example

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
```

Step 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.

Example

```
FC_switch_A_1:admin> reboot
```

Step 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.

Step 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.

Example

```
FC_switch_A_1:admin> switchCfgPersistentEnable
```

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

Example

```
FC_switch_A_1:admin> switchShow
```

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

Example

```
FC_switch_A_1:admin> portbuffershow
```

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

Example

```
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.

Step 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`

Step 12. Repeat this procedure for the second switch fabric.

Upgrading or downgrading the firmware on a Cisco FC switch

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

Before you begin

- The system must be properly cabled.
- All paths to the storage shelves must be available.
- The disk shelf stacks must be stable.
- The FC switch fabric must be healthy.
- All components in the system must be healthy.
- The system must be operating normally.
- You need the admin password and access to an FTP or SCP server.

About this task

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.

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

Step 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.

Step 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.

Step 3. Copy the kickstart and system files to the switches:
`copy` *source_file target_file*

Example

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
```

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

Step 5. Install the desired firmware on the switches:
`install all system bootflash:`*systemfile_name* `kickstart bootflash:`*kickstartfile_name*

Example

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.
```

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

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

Step 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`

Step 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.

Before you begin

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

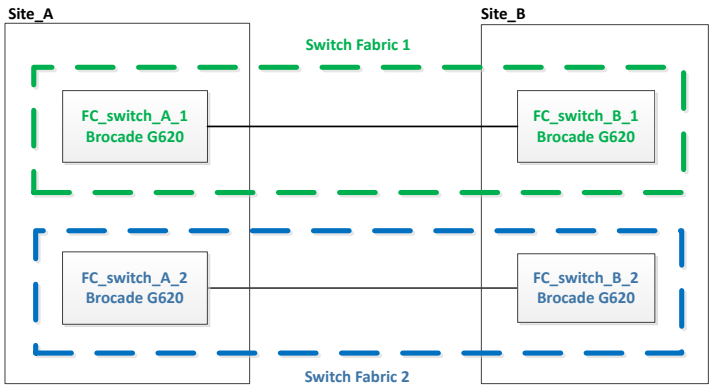
  The illustrations in the following steps show current switches.
- The switches must be running the most recent supported firmware.

About this task

- This procedure is nondisruptive and takes approximately two hours to complete.
- You need the admin password and access to an FTP or SCP server.

The switch fabrics are upgraded one at a time.

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

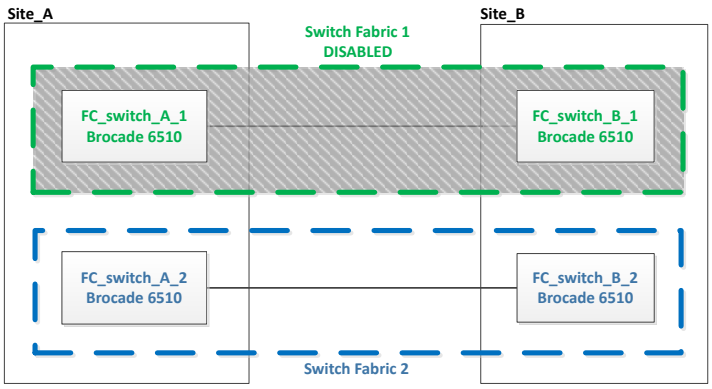


Step 1. Disable the first switch fabric:

```
FC_switch_A_1:admin> switchCfgPersistentDisable
```

Example

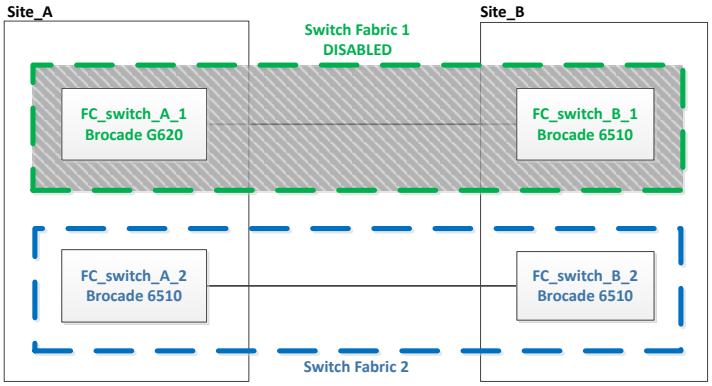
```
FC_switch_A_1:admin> switchCfgPersistentDisable
```



Step 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:
`switchCfgPersistentDisable`
The command disables both switches in the switch fabric.
```
FC_switch_A_1:admin> switchCfgPersistentDisable
```

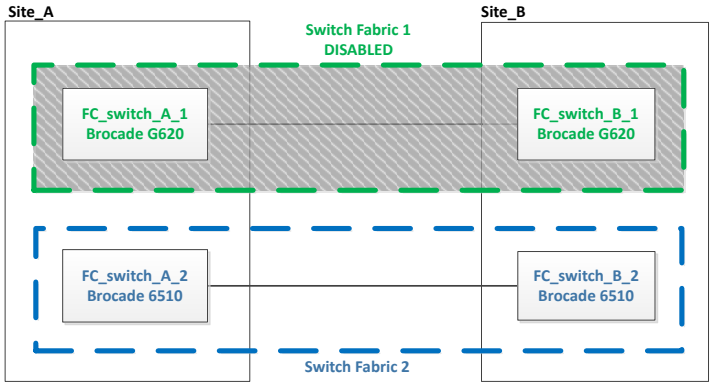
d. Cable the new switch using the recommended port assignments.

"Port assignments for FC switches" on page 134

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

Result

Both switches in fabric 1 have been replaced.



Step 3. Power up the new switches and let them boot up.

Step 4. Download the RCF files for the new switch.

Lenovo Data Center Support Site

Step 5. Apply the RCF files to both new switches in the fabric, following the directions on the download page.

Step 6. Save the switch configuration:
`cfgSave`

Step 7. Wait 10 minutes to allow the configuration to stabilize.

Step 8. 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:

Example

```
node_A_1> run local sysconfig -v
Lenovo Release 9.6P5: Fri Dec 13 13:23:47 EST 2019
System ID: 4068741258 (node_A_1); partner ID: 4068741260 (node_B_1)
System Serial Number: 940001025471 (node_A_1)
System Rev: A4
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
.
.
.
```

Step 9. Returning to the switch prompt, verify the switch firmware version:
`firmwareShow`

The switches must be running the most recent supported firmware.

Step 10. 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 (`*>`).

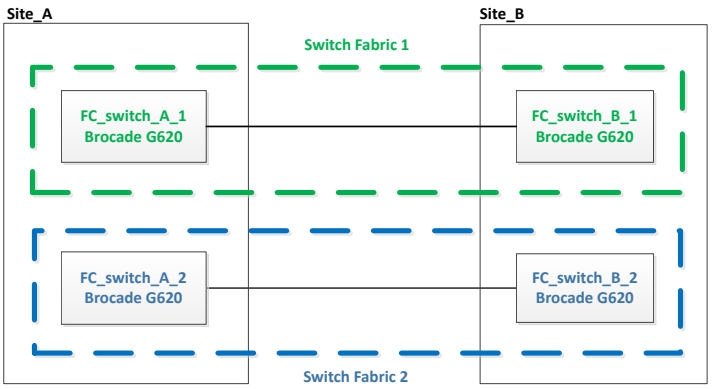
b. Perform the switchover operation with the -simulate parameter:
`metrocluster switchover -simulate`

c. Return to the admin privilege level:
```
set -privilege admin
```

Step 11. Repeat the previous steps on the second switch fabric.

Result

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.

Before you begin

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

About 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.

Note: 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" on page 134

Step 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*

Example

```
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*

Example

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

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

Example

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

Step 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 as described in the *MetroCluster Installation and Configuration Guide*, but skipping the “Configuring zoning on Brocade FC switches” section.

Fabric-attached MetroCluster installation and configuration

You configure zoning later in this procedure.

Note: 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.

Example

```
FC_switch_B_1:admin> portcfgpersistentdisable 0
FC_switch_B_1:admin> portcfgpersistentdisable 1
FC_switch_B_1:admin> portcfgpersistentdisable 2
```

```
FC_switch_B_1:admin> portcfgpersistentdisable 3
FC_switch_B_1:admin> portcfgpersistentdisable 4
FC_switch_B_1:admin> portcfgpersistentdisable 5

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

Step 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.

Step 5. If you want to enable ISL encryption, complete the applicable tasks in the *Fabric-attached MetroCluster Installation and Configuration Guide*.
If you are enabling ISL encryption, you need to complete the following tasks:

- Disabling the virtual fabric
- Setting the payload
- Setting the authentication policy
- Enabling ISL encryption on Brocade switches

Step 6. Complete the configuration of the new switch:

a. Enable the ISLs:
`portcfgpersistentenable` *port-number*

Example

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

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

Example

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

Index Port Address Media Speed State  Proto
==============================================
...
10   10  030A00 id  16G  Online FC E-Port 10:00:00:05:33:86:89:cb "FC_switch_A_1"
11   11  030B00 id  16G  Online FC E-Port 10:00:00:05:33:86:89:cb "FC_switch_A_1" (downstream)
...
```

c. Enable the storage ports that connect to the FC bridges.

Example

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

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

Example

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
```

Step 7. Verify that the ports are online:
`switchshow`

Step 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`

Renaming a Brocade FC switch

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

Step 1. Persistently disable the switch or switches in one fabric: `switchcfgpersistentdisable`
The following example shows the output for the `switchcfgpersistentdisable` command:

Example

```
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>
```

Step 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:

Example

```
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>
```

Step 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:

Example

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

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

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

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


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

Step 4. Persistently enable the switches: `switchcfgpersistentenable`
The following example shows the output for the `switchcfgpersistentenable` command:

Example

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

FC_switch_1_B:admin>


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

```
HIF Mode:       OFF

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

Step 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:

Example

```
cluster_A::*> storage switch show
  (storage switch show)
                       Symbolic                                  Is        Monitor
Switch                 Name     Vendor  Model Switch WWN       Monitored Status
---------------------- -------- ------- ----- ---------------- --------- -------
Brocade_172.20.7.90
                       RTP-FC01-510Q40
                                Brocade Brocade7840
                                        100Oc4f57c904bc8 true      ok
Brocade_172.20.7.91
                       RTP-FC02-510Q40
                                Brocade Brocade7840
                                        100050eb1a9aa579 true      ok
Brocade_172.20.7.92
```

Changing 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.

About 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).

Step 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.

Example

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
```

Step 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:
	`brocade_switch_A_1:admin> portcfgspeed 40 16`
	
	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` *port number port distance* command on both switches in the fabric.
To remove an ISL...	Disconnect the link (cable).
To add an ISL...	Insert SFPs into the ports you are adding as ISL ports. Ensure that these ports are listed in the *Fabric-attached MetroCluster Installation and Configuration Guide* for the switch you are adding them to and cable them according to this guide.
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.

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

Note: 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.

a. Persistently disable the switches using the `switchcfgpersistentdisable` command as shown in the following example:

Example

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

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

Result

Scenario	Steps

Configure the ISL for trunking (IOD)	Set the `portcfgtrunkport` *port number* to 1:
	
	`FC_switch_A_1:admin> portcfgtrunkport 20 1`
	`FC_switch_A_1:admin> portcfgtrunkport 21 1`
	`FC_switch_B_1:admin> portcfgtrunkport 20 1`
	`FC_switch_B_1:admin> portcfgtrunkport 21 1`

Configure the ISL for trunking (OOD)	Set the `portcfgtrunkport` *port number* to 0:
	
	`FC_switch_A_1:admin> portcfgtrunkport 20 0`
	`FC_switch_A_1:admin> portcfgtrunkport 21 0`
	`FC_switch_B_1:admin> portcfgtrunkport 20 0`
	`FC_switch_B_1:admin> portcfgtrunkport 21 0`

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

Result

Scenario	Steps

Configure the fabric for IOD	Set the three settings of IOD, APT, and DLS using the `iodset`, `aptpolicy`*policy* , and `dlsreset` commands as shown in the following example:
	
	`Switch_A_1:admin> iodset`
	`Switch_A_1:admin> aptpolicy 1`
	`Policy updated successfully.`
	`Switch_A_1:admin> dlsreset`
	`FC_switch_A_1:admin>portcfgtrunkport 40 1`
	`FC_switch_A_1:admin>portcfgtrunkport 41 1`
	
	`Switch_B_1:admin> iodset`
	`Switch_B_1:admin> aptpolicy 1`
	`Policy updated successfully.`
	`Switch_B_1:admin> dlsreset`
	`FC_switch_B_1:admin>portcfgtrunkport 20 1`
	`FC_switch_B_1:admin>portcfgtrunkport 21 1`

Configure the fabric for OOD	Set the three settings of IOD, APT, and DLS using the `iodreset`, `aptpolicy`*policy* , and `dlsset` commands as shown in the following example:
	
	`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`
	
	`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`

d. Enable the switches persistently using the `switchcfgpersistentenable` command.

Example

```
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:

Example

```
brocade_switch_A_1:admin>
switchenable
```

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

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
```

Note: You must run these commands on both switches.

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

Example

```
switch_A_1:admin> iodshow
IOD is set

switch_A_1:admin> aptPpolicy
              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
```

Note: You must run these commands on both switches.

Step 4. Verify that the ISLs are online and trunked (if the linking equipment supports trunking) using the `islshow` and `trunkshow` commands.

Note: 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:
	`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`

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`

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

Example

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

Step 6. Enable the ports that were disabled in step 1: `portenable` *port number*

Example

The following example shows ISL ports 0–7 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.

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" on page 134

Step 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.

Example

```
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.

Example

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

Step 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.

Example

```
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.

Example

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

Step 3. If the impaired switch is still operational, determine the WWN for the switch:
`show wwn switch`

Example

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

Step 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.

Example

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

Step 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.

Step 6. Enable the ISL ports on the replacement switch.

Example

```
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#
```

Step 7. Verify that the ISL ports on the replacement switch are up:
`show interface brief`

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

a. Distribute the zoning information from the healthy fabric.

Example

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`

Example

```
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.

Example

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#
```

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

Example

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
```

e. Add the ports of the replacement switch to the zones.

Example

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
```

f. Verify that the zoning is properly configured:
`show zone`

Example

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#
```

g. Enable the connectivity to storage and the controllers.

Example

The following example shows the port usage:

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

Step 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`

Changing speed of ISL ports on a Cisco FC switch

You might need to change the speed of ISL ports on a switch to improve the quality of the ISL. ISLs traveling greater distances might need their speed lowered to improve quality.

Before you begin

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

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

```
FC_switch_A_1# config t
```

Enter 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
```

Step 2. Change the speed of the ISL ports on both switches in the fabric:

```
FC_switch_A_1# config t
```

Enter 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
```

Note: Speeds for ports are 16 = 16,000 Gbps, 8 = 8,000 Gbps, 4 = 4,000 Gbps.

Ensure that these ISL ports for your switch are listed in the *Fabric-attached MetroCluster Installation and Configuration Guide*.

Step 3. Enable all ISL ports (if not enabled) on both switches in the fabric:

```
FC_switch_A_1# config t
```

Enter 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
```

Step 4. Verify the ISLs are established between both switches:

```
show topology isl
-------------------------------------------------------------------------------
_________Local_________ _________Remote_________ VSAN Cost I/F  PC    I/F  Band
 PC Domain SwName   Port  Port   SwName Domain PC          Stat Stat Speed width
-------------------------------------------------------------------------------
  1   0x11 cisco9 fc1/36  fc1/36 cisco9 0xbc    1    1   15 up   up    16g   64g
  1   0x11 cisco9 fc1/40  fc1/40 cisco9 0xbc    1    1   15 up   up    16g   64g
  1   0x11 cisco9 fc1/44  fc1/44 cisco9 0xbc    1    1   15 up   up    16g   64g
  1   0x11 cisco9 fc1/48  fc1/48 cisco9 0xbc    1    1   15 up   up    16g   64g
```

Step 5. Repeat the procedure for the second switch fabric.

Adding ISLs to a Cisco switch

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

Before you begin

Steps that are completed on one switch must also be completed on the other to ensure ISL connectivity.

Step 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 configuration commands, one per line. End with CTRL-Z after all the configuration commands have been entered.

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

Step 2. Insert SFPs into the ports you are adding as ISL ports, and cable them according to the *Installation and Configuration Guide*.
Ensure that these ports are listed in the *Installation and Configuration Guide* for the switch you are adding them to.

Step 3. Configure the ISL ports in accordance with the *Installation and Configuration Guide*.

Step 4. Enable all ISL ports (if not enabled) on both switches in the fabric:

```
FC_switch_A_1# config t
```

Enter the configuration commands, one per line. End with CTRL-Z.

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

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

```
show topology isl
```

Step 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    16g   64g
  1   0x11 cisco9 fc1/40  fc1/40 cisco9 0xbc    1    1   15 up   up    16g   64g
  1   0x11 cisco9 fc1/44  fc1/44 cisco9 0xbc    1    1   15 up   up    16g   64g
  1   0x11 cisco9 fc1/48  fc1/48 cisco9 0xbc    1    1   15 up   up    16g   64g
```

Chapter 6. Performing IP switch maintenance and replacement

If necessary, you can nondisruptively upgrade, downgrade, or replace the IP switches in a MetroCluster configuration.

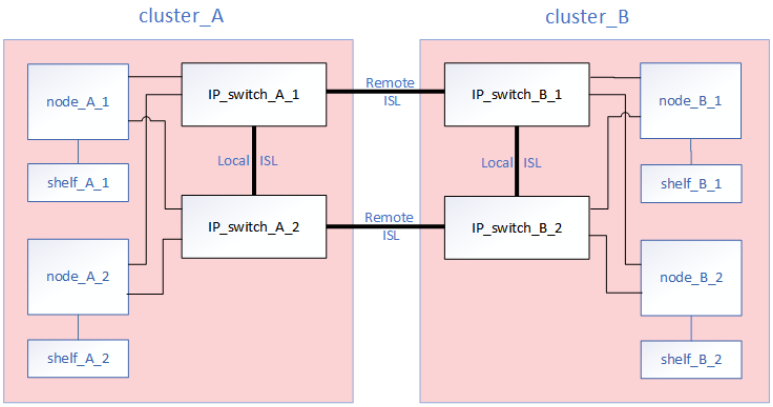
Replacing an IP switch

You might need to replace a failed switch, or upgrade or downgrade a switch. The new switch can be the same as the old switch when a switch has failed, or you can change the switch type (upgrade or downgrade the switch).

About this task

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

This procedure is for Cisco or Broadcom switches. If you want to change the switch vendor, further steps are required.



If you upgrade or downgrade the networks, you must repeat this procedure for the second network.

Step 1. Check the health of the configuration.

a. Check that the MetroCluster is configured and in normal mode on each cluster: `metrocluster show`

Example

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

© Copyright Lenovo 2019, 2023

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

Example

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

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

Example

```
cluster_A::> metrocluster check run

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

Component           Result
------------------- ---------
nodes               ok
lifs                ok
config-replication  ok
aggregates          ok
4 entries were displayed.

Command completed. Use the "metrocluster check show -instance" command or sub-commands in "metrocluster
 check" directory for detailed results.
To check if the nodes are ready to do a switchover or switchback operation, run "metrocluster switchover
-simulate" or "metrocluster switchback -simulate", respectively.
```

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

Step 2. Configure the new switch before installation.

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

Follow the steps in the section *Configuring the IP switches* in the *MetroCluster IP Installation and Configuration Guide.*

MetroCluster IP installation and configuration

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

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

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

Step 4. Disconnect the ISL connections from the remote switch that connect to the old switch.

Example

You should disconnect the ISL connections from the ports on the IP_switch_A_1 that connect to IP_switch_B_1.

Step 5. Power off the switch, remove the cables and physically remove IP_switch_B_1.

Step 6. Install the new switch.

Cable the new switch first (including the ISLs) according to the steps in the *Cabling the IP switches* section in the *MetroCluster IP Installation and Configuration Guide.*

MetroCluster IP installation and configuration

Note: The used ports might be different from those on the old switch if the switch type is different. If you are upgrading or downgrading the switches, do **NOT** cable the local ISLs. Only cable the local ISLs if you are upgrading or downgrading the switches in the second network and both switches at one site are the same type.

Step 7. Power up the switch or switches.
If the new switch is the same, power up the new switch. If you are upgrading or downgrading the switches, then power up both switches. The configuration can operate with two different switches at each site until the second network is updated.

Step 8. Verify that the MetroCluster configuration is healthy by repeating step 1.

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

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

Upgrading firmware on MetroCluster IP switches

Step 1. Check the health of the configuration.

a. Check that the MetroCluster is configured and in normal mode on each cluster: `metrocluster show`

Example

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

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

Example

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

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

Example

```
cluster_A::> metrocluster check run

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

Component           Result
------------------- ---------
nodes               ok
lifs                ok
config-replication  ok
aggregates          ok
4 entries were displayed.

Command completed. Use the "metrocluster check show -instance" command or sub-commands in "metrocluster
 check" directory for detailed results.
To check if the nodes are ready to do a switchover or switchback operation, run "metrocluster switchover
 -simulate" or "metrocluster switchback -simulate", respectively.
```

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

Step 2. Install the software on the first switch.
Follow the steps for installing switch software in the relevant section in the *MetroCluster IP Installation and Configuration Guide* depending on whether the switch type is Cisco or Broadcom:

- *Downloading and installing the Cisco switch NX-OS software*
- *Downloading and installing the Broadcom switch EFOS software*

MetroCluster IP installation and configuration

Repeat this step for each of the switches.

Note: In the sections referenced, skip the step on downloading the reference configuration (RCF) files.

Step 3. Repeat step 1 to check the health of the configuration.

Upgrading RCF files on MetroCluster IP switches

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

Before you begin

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

About this task

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

Step 1. Verify the health of the configuration.

a. Verify that the MetroCluster components are healthy: `metrocluster check run`

Example

```
dpgqa-mcc-funct-8040-0403_siteA::*> metrocluster check run

----
```

Result

The operation runs in the background.

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

Result

After approximately five minutes, the following results are displayed:

```
----------
dpgqa-mcc-funct-8040-0403_siteA::*> metrocluster check show

Last Checked On: 4/7/2019 21:15:05

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

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

d. Verify that there are no health alerts: `system health alert show`

Step 2. Prepare the IP switches for the application of the new RCF files.Follow the steps in the section for your switch vendor from the *MetroCluster IP Installation and Configuration guide*.

MetroCluster IP installation and configuration

- *Resetting the Broadcom IP switch to factory defaults*
- *Resetting the Cisco IP switch to factory defaults*

Step 3. Download and install the IP RCF file depending on your switch vendor.Follow the steps in the section for your switch vendor from the *MetroCluster IP Installation and Configuration guide*.

MetroCluster IP installation and configuration

- *Downloading and installing the Broadcom IP RCF files*
- *Downloading and installing the Cisco IP RCF files*

a. Update the switches in the following order: Switch_A_1, Switch_B_1, Switch_A_2, Switch_B_2.

Renaming a Cisco IP switch

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

About this task

In the examples in this task, the switch name is changed from myswitch to IP_switch_A_1.

Step 1. Enter global configuration mode:
`configure terminal`

Example

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

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

Step 2. Rename the switch:

`switchname` *new-switch-name*

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

Example

The CLI prompt changes to reflect the new name:

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

Step 3. Exit configuration mode:

`exit`

Example

The top-level switch prompt is displayed:

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

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

`copy running-config startup-config`

Step 5. Verify that the switch name change is visible from the ONTAP cluster prompt.
Note that the new switch name is shown, and the old switch name (`myswitch`) does not appear.

a. Enter advanced privilege mode, pressing `y` when prompted:
`set -privilege advanced`

b. Display the attached devices:
`network device-discovery show`

c. Return to admin privilege mode:
`set -privilege admin`

Example

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

```
cluster_A::storage show> set advanced

Warning: These advanced commands are potentially dangerous; use them only when directed to do so by Lenovo personnel.
Do you want to continue? {y|n}: y

cluster_A::storage show*> network device-discovery show
Node/       Local  Discovered
Protocol    Port   Device                    Interface          Platform
----------- ------ ------------------------- ----------------  ----------------
node_A_2/cdp
            e0M    LF01-410J53.mycompany.com(SAL18516DZY)
                                             Ethernet125/1/28  N9K-C9372PX
            e1a    IP_switch_A_1(FOC21211RBU)
                                             Ethernet1/2       N3K-C3232C
            e1b    IP_switch_A_1(FOC21211RBU)
                                             Ethernet1/10      N3K-C3232C
.
.
.                                             Ethernet1/18      N9K-C9372PX
node_A_1/cdp
            e0M    LF01-410J53.mycompany.com(SAL18516DZY)
                                             Ethernet125/1/26  N9K-C9372PX
            e0a    IP_switch_A_2(FOC21211RB5)
                                             Ethernet1/1       N3K-C3232C
```

```
            e0b    IP_switch_A_2(FOC21211RB5)
                                               Ethernet1/9       N3K-C3232C
            e1a    IP_switch_A_1(FOC21211RBU)
.
.
.
16 entries were displayed.
```

Chapter 7. Adding, removing, or changing ISL ports nondisruptively

You might need to add, remove, or change ISL ports. You can convert dedicated ISL ports to shared ISL ports, or change the speed of ISL ports on a switch. ISLs traveling greater distances might require lower speeds to improve quality. You also might need to increase the speed of ISL ports to take advantage of system upgrades and add ISLs to a Cisco IP switch if you are adding or upgrading hardware.

Before you begin

If you are converting dedicated ISL ports to shared ISL ports, ensure the new ports meet the requirements for shared ISL ports.

About this task

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

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

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

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

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

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

Step 2. Remove the existing cables and transceivers.

Step 3. Change the ISL port as required.

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

© Copyright Lenovo 2019, 2023

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

```
switch_A_1# conf t
```

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

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

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

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

```
switch_A_1# show int eth1/17,eth1/18 brief
```

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

```
switch_A_1# show int eth1/17,eth1/18 brief
--------------------------------------------------------------------------------
Ethernet      VLAN    Type Mode   Status  Reason                   Speed     Port
Interface
--------------------------------------------------------------------------------
Eth1/17       10         eth     access    up      none                      100G(D) 10
Eth1/18       10         eth     access    up      none                      100G(D) 10

switch_B_1# show int eth1/17,eth1/18 brief
--------------------------------------------------------------------------------
Ethernet      VLAN    Type Mode   Status  Reason                   Speed     Port
Interface
--------------------------------------------------------------------------------
Eth1/17       10         eth     access    up      none                      100G(D) 10
Eth1/18       10         eth     access    up      none                      100G(D) 10
```

Step 6. Repeat the procedure for fabric 2.

Chapter 8. Identifying storage in a MetroCluster IP configuration

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

Identification of local and remote shelves

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

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

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

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

Example of sysconfig -a output

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

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

```
node_A_1> run local sysconfig -a

Lenovo Release 9.6P5: Fri Dec 13 13:23:47 EST 2019
System ID: 1111111111 (node_A_1); partner ID: 2222222222 (node_A_2)
System Serial Number: serial-number (node_A_1)
.
.
.
slot 0: NVMe Disks
                0    : NETAPP   X4001S172A1T9NTE NA01 1831.1GB 4160B/sect (S3NBNX0J500528)
                1    : NETAPP   X4001S172A1T9NTE NA01 1831.1GB 4160B/sect (S3NBNX0J500735)
                2    : NETAPP   X4001S172A1T9NTE NA01 1831.1GB 4160B/sect (S3NBNX0J501165)
.
.
.
slot 3: SAS Host Adapter 3a (PMC-Sierra PM8072 rev. C, SAS, <UP>)
MFG Part Number:    Microsemi Corp. 110-03801 rev. A0
Part number:        111-03801+A0
Serial number:      7A1063AF14B
Date Code:          20170320
Firmware rev:       03.08.09.00
Base WWN:           5:0000d1:702e69e:80
Phy State:          [12] Enabled, 12.0 Gb/s
                    [13] Enabled, 12.0 Gb/s
                    [14] Enabled, 12.0 Gb/s
                    [15] Enabled, 12.0 Gb/s
Mini-SAS HD Vendor:         Molex Inc.
Mini-SAS HD Part Number:    112-00436+A0
Mini-SAS HD Type:           Passive Copper (unequalized) 0.5m ID:00
Mini-SAS HD Serial Number: 614130640
                75.0 : NETAPP   X438_S1633400AMD NA04 381.3GB 520B/sect (S20KNYAG501805)
                75.1 : NETAPP   X438_S1633400AMD NA04 381.3GB 520B/sect (S20KNYAG502050)
```

© Copyright Lenovo 2019, 2023

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

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

slot 3: SAS Host Adapter 3c (PMC-Sierra PM8072 rev. C, SAS, <UP>)
MFG Part Number:    Microsemi Corp. 110-03801 rev. A0
Part number:        111-03801+A0
Serial number:      7A1063AF14B
Date Code:          20170320
Firmware rev:       03.08.09.00
Base WWN:           5:0000d1:702e69e:88
Phy State:          [0] Enabled, 12.0 Gb/s
                    [1] Enabled, 12.0 Gb/s
                    [2] Enabled, 12.0 Gb/s
                    [3] Enabled, 12.0 Gb/s
Mini-SAS HD Vendor:         Molex Inc.
Mini-SAS HD Part Number:    112-00436+A0
Mini-SAS HD Type:           Passive Copper (unequalized) 0.5m ID:00
Mini-SAS HD Serial Number: 614130691
                75.0 : NETAPP   X438_S1633400AMD NA04 381.3GB 520B/sect (S20KNYAG501805)
                75.1 : NETAPP   X438_S1633400AMD NA04 381.3GB 520B/sect (S20KNYAG502050)
                75.2 : NETAPP   X438_PHM2400MCTO NA04 381.3GB 520B/sect (25M0A03WT2KA)
                75.3 : NETAPP   X438_S1633400AMD NA04 381.3GB 520B/sect (S20KNYAG501793)
.
.
.
Shelf 75: DS224-12  Firmware rev. IOM12 A: 0220  IOM12 B: 0220
Shelf 76: DS224-12  Firmware rev. IOM12 A: 0220  IOM12 B: 0220

slot 3: SAS Host Adapter 3d (PMC-Sierra PM8072 rev. C, SAS, <UP>)
MFG Part Number:    Microsemi Corp. 110-03801 rev. A0
Part number:        111-03801+A0
Serial number:      7A1063AF14B
Date Code:          20170320
Firmware rev:       03.08.09.00
Base WWN:           5:0000d1:702e69e:8c
Phy State:          [4] Enabled, 12.0 Gb/s
                    [5] Enabled, 12.0 Gb/s
                    [6] Enabled, 12.0 Gb/s
                    [7] Enabled, 12.0 Gb/s
Mini-SAS HD Vendor:         Molex Inc.
Mini-SAS HD Part Number:    112-00436+A0
Mini-SAS HD Type:           Passive Copper (unequalized) 0.5m ID:01
Mini-SAS HD Serial Number: 614130690
                75.0 : NETAPP   X438_S1633400AMD NA04 381.3GB 520B/sect (S20KNYAG501805)
                75.1 : NETAPP   X438_S1633400AMD NA04 381.3GB 520B/sect (S20KNYAG502050)
                75.2 : NETAPP   X438_PHM2400MCTO NA04 381.3GB 520B/sect (25M0A03WT2KA)
.
.
.
Shelf 75: DS224-12  Firmware rev. IOM12 A: 0220  IOM12 B: 0220
Shelf 76: DS224-12  Firmware rev. IOM12 A: 0220  IOM12 B: 0220

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

```
                0.5  : NETAPP   X4001S172A1T9NTE NA01 1831.1GB 4160B/sect (S3NBNX0J500598)
                0.12 : NETAPP   X4001S172A1T9NTE NA01 1831.1GB 4160B/sect (S3NBNX0J501094)
                0.13 : NETAPP   X4001S172A1T9NTE NA01 1831.1GB 4160B/sect (S3NBNX0J500519)
.
.
.
Shelf 0: FS4483PSM3E  Firmware rev. PSM3E A: 0103  PSM3E B: 0103
Shelf 35: DS224-12  Firmware rev. IOM12 A: 0220  IOM12 B: 0220
Shelf 36: DS224-12  Firmware rev. IOM12 A: 0220  IOM12 B: 0220

node_A_1::>
```

Chapter 9. Hot-adding storage to a MetroCluster FC configuration

You can add storage to a MetroCluster FC configuration without bringing down the system. The procedure you use depends on the type of MetroCluster FC configuration and whether you are adding a single disk shelf to an existing stack or an entire stack.

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.

About this task

- This procedure is nondisruptive and takes approximately two hours to complete.
- You need the admin password and access to an FTP or SCP server.

This task applies to a MetroCluster FC configuration in which the storage is connected directly to the storage controllers with SAS cables. It does not apply to MetroCluster FC configurations using FC-to-SAS bridges or FC switch fabrics.

Step 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-adding SAS storage to a bridge-attached MetroCluster FC configuration

You can hot-add either an individual SAS disk shelf or a stack and set of bridges to an existing MetroCluster system.

These procedures apply to MetroCluster FC configurations using FC-to-SAS bridges that are either directly attached to the storage controllers or attached to an FC fabric.

These procedures are nondisruptive and take approximately four hours to complete.

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

Hot-adding a stack of SAS disk shelves to an existing pair of FibreBridge 7600N bridges

You can hot-add a stack of SAS disk shelves to an existing pair of FibreBridge 7600N bridges that have available ports.

Before you begin

© Copyright Lenovo 2019, 2023

- All of the disk shelves in the MetroCluster configuration (both the new shelves and existing shelves) must be running the same firmware version.
- The FibreBridge 7600N 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 “Requirements for using other interfaces to configure and manage FibreBridge bridges” on page 42.

Attention: 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.

Step 1. Properly ground yourself.

Step 2. From the console of either controller, verify that your system has disk autoassignment enabled:
`storage disk option show`

Result

The `Auto Assign` column indicates whether disk autoassignment is enabled.

```
Node        BKg. FW. Upd.  Auto Copy   Auto Assign  Auto Assign Policy
----------  -------------  ----------  -----------  ------------------
node_A_1             on           on           on           default
node_A_2             on           on           on           default
2 entries were displayed.
```

Step 3. Disable the switch ports for the new stack.

Step 4. 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” on page 8.

Step 5. 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.

Step 6. Save the configuration and reboot each bridge:
`SaveConfiguration Restart`

Step 7. Cable the disk shelves to the bridges:

a. Daisy-chain the disk shelves in each stack.

The *Hardware Maintenance and Installation Guide* for your expansion enclusre provides detailed information about cascading enclosure together.

ThinkSystem Storage DM Series

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

Result

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.

Note: The bridge SAS port B is disabled.

Step 8. 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	1. 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. 2. 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. 3. Click **Advanced** in the menu. 4. View the connected devices: `sastargets` 5. 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.

Note: 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.

Example

The following output shows that 10 disks are connected:

```
Tgt VendorID ProductID        Type        SerialNumber
  0 NETAPP   X410_S15K6288A15 DISK        3QP1CLE30000994OUHJV
  1 NETAPP   X410_S15K6288A15 DISK        3QP1ELF60000994OV1BV
  2 NETAPP   X410_S15K6288A15 DISK        3QP1G3EW0000994OU2MO
  3 NETAPP   X410_S15K6288A15 DISK        3QP1EWMP0000994OU1X5
  4 NETAPP   X410_S15K6288A15 DISK        3QP1FZLE0000994OG8YU
  5 NETAPP   X410_S15K6288A15 DISK        3QP1FZLF0000994OTZKZ
  6 NETAPP   X410_S15K6288A15 DISK        3QP1CEB40000939MGXL
  7 NETAPP   X410_S15K6288A15 DISK        3QP1G7A90000939FNTT
  8 NETAPP   X410_S15K6288A15 DISK        3QP1FYOT0000994OG8PA
```

```
9 NETAPP   X410_S15K6288A15 DISK        3QP1FXW600009940VERQ
```

Step 9. 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 Step 8 on page 87 for each remaining bridge.
Not correct	1. Check for loose SAS cables or correct the SAS cabling by repeating Step Step 7 on page 86. 2. Repeat Step Step 8 on page 87.

Step 10. 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:

Notes: The Brocade and Cisco switches use different port numbering, as shown in the following tables

- On Brocade switches, the first port is numbered “0”.
- On Cisco switches, the first port is numbered “1”.

Configurations using FibreBridge 7600N using both FC ports (FC1 and FC2)										
DR GROUP 1										
			Lenovo G6505		**Lenovo B6510**		**Lenovo DB620S, Lenovo DB630S**		**Lenovo DB720**	
Component		**Port**	**Switch 1**	**Switch 2**	**Switc-h 1**	**Switc-h 2**	**Switch 1**	**Switch 2**	**Swit-ch 1**	**Swit-ch 2**
Stack 1	bridge_x_ 1a	FC1	8		8		8		10	
		FC2	-	8	-	8	-	8	-	10
	bridge_x_ 1B	FC1	9	-	9	-	9	-	11	-
		FC2	-	9	-	9	-	9	-	11
Stack 2	bridge_x_ 2a	FC1	10	-	10	-	10	-	14	-
		FC2	-	10	-	10	-	10	-	14
	bridge_x_ 2B	FC1	11	-	11	-	11	-	17	-
		FC2	-	11	-	11	-	11	-	17
Stack 3	bridge_x_ 3a	FC1	12	-	12	-	12	-	18	-
		FC2	-	12	-	12	-	12	-	18
	bridge_x_ 3B	FC1	13	-	13	-	13	-	19	-
		FC2	-	13	-	13	-	13	-	19
Stack *y*	bridge_x_ *y*a	FC1	14	-	14	-	14	-	20	-
		FC2	-	14	-	14	-	14	-	20
	bridge_x_ *y*b	FC1	15	-	15	-	15	-	21	-
		FC2	--	15	--	15	-	15	-	21
Note: Additional bridges can be cabled to ports 16, 17, 20 and 21 in DB620S and DB630S switches.										

Configurations using FibreBridge 7600N using both FC ports (FC1 and FC2)								
DR GROUP 2								
			Lenovo DB620S, Lenovo DB630S		Lenovo B6510		Lenovo DB720	
Component		Port	Switch 1	Switch 2	Switch 1	Switch 2	Switch 1	switch 2
Stack 1	bridge_x_51a	FC1	26	-	32	-	32	-
		FC2	-	26	-	32	-	32
	bridge_x_51b	FC1	27	-	33	-	33	-
		FC2	-	27	-	33	-	33
Stack 2	bridge_x_52a	FC1	30	-	34	-	34	-
		FC2	-	30	-	34	-	34
	bridge_x_52b	FC1	31	-	35	-	35	-
		FC2	-	31	-	35	-	35
Stack 3	bridge_x_53a	FC1	32	-	36	-	36	-
		FC2	-	32	-	36	-	36
	bridge_x_53b	FC1	33	-	37	-	37	-
		FC2	-	33	-	37	-	37
Stack *y*	bridge_x_5*y*a	FC1	34	-	38	-	38	-
		FC2	-	34	-	38	-	38
	bridge_x_5*y*b	FC1	35	-	39	-	39	-
		FC2	-	35	-	39	-	39
Note: Additional bridges can be cabled to ports 36 - 39 in DB620S and DB630S switches.								

Configurations using FibreBridge 7600N using one FC port (FC1 or FC2) only									
DR GROUP 1									
		Lenovo B6505		Lenovo B6510		Lenovo DB620S, Lenovo DB630S		Lenovo DB720	
Com-po-nent	Port	Switch 1	Switch 2	Switch 1	Switch 2	Switch 1	Switch 2	Swit-ch 1	Swit-ch 2
Stack 1	bridge_x_1a	8		8		8		10	
	bridge_x_ 1b	-	8	-	8	-	8	-	10
Stack 2	bridge_x_2a	9	-	9	-	9	-	11	-
	bridge_x_ 2b	-	9	-	9	-	9	-	11
Stack 3	bridge_x_3a	10	-	10	-	10	-	14	-
	bridge_x_ 4b	-	10	-	10	-	10	-	14

Configurations using FibreBridge 7600N using one FC port (FC1 or FC2) only									
DR GROUP 1									
		Lenovo B6505		Lenovo B6510		Lenovo DB620S, Lenovo DB630S		Lenovo DB720	
Com-po-nent	**Port**	**Switch 1**	**Switch 2**	**Switch 1**	**Switch 2**	**Switch 1**	**Switch 2**	**Swit-ch 1**	**Swit-ch 2**
Stack y	bridge_x_*ya*	11	-	11	-	11	-	15	-
	bridge_x_*yb*	-	11	-	11	-	11	-	15
Note: Additional bridges can be cabled to ports 12 - 17, 20 and 21 in DB620 and DB630 switches. Additional bridges can be cabled to ports 16 - 17, 20 and 21 DB720 switches.									

Configurations using FibreBridge 7600N using one FC port (FC1 or FC2) only							
DR GROUP 2							
		Lenovo DB720		Lenovo DB620S, Lenovo DB630S		Lenovo B6510	
Stack 1	bridge_x_51a	32	-	26	-	32	-
	bridge_x_51b	-	32	-	26	-	32
Stack 2	bridge_x_52a	33	-	27	-	33	-
	bridge_x_52b	-	33	-	27	-	33
Stack 3	bridge_x_53a	34	-	30	-	34	-
	bridge_x_54b	-	34	-	30	-	34
Stack *y*	bridge_x_*ya*	35	-	31	-	35	-
	bridge_x_*yb*	-	35	-	31	-	35
Note: Additional bridges can be cabled to ports 32 - 39 in DB620S and DB630S switches. Additional bridges can be cabled to ports 36 - 39 in DB720 switches.							

Step 11. If your system does not have disk autoassignment enabled, assign disk drive ownership.

Disk and Aggregate Management Guide

Note: 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.

Step 12. Enable the switch ports for the new stack.

Step 13. 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
```

Step 14. 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 either FibreBridge 7600N or 6500N 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 "Requirements for using other interfaces to configure and manage FibreBridge bridges" on page 42

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.
- 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 Guide
- Each FC switch must have one FC port available for one bridge to connect to it.

  Note: 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.

Step 1. View the following document from the Lenovo Support Site:

The *Hardware Maintenance and Installation Guide* for your expansion enclusre: ThinkSystem Storage DM Series

Step 2. Download content from the ATTO website and from the Lenovo 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 "Hot-adding the stack of shelves" on page 92.
   - Make a copy of the ATTO FibreBridge Firmware Download page and release notes for reference later.

Step 3. 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.

Step 4. 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.

Step 5. 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 Lenovo, 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.

Before you begin

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

About this task

- Hot-adding a stack of SAS disk shelves and bridges is a nondisruptive procedure if all of the interoperability requirements are met.

- 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.

Note: Starting 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.

Attention: 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.

Step 1. Properly ground yourself.

Step 2. From the console of either controller module, check whether your system has disk autoassignment enabled:
```
storage disk option show
```

Result

The `Auto Assign` column indicates whether disk autoassignment is enabled.

```
Node         BKg. FW. Upd.  Auto Copy   Auto Assign  Auto Assign Policy
----------   -------------  ----------  -----------  ------------------
node_A_1              on           on           on           default
node_A_2              on           on           on           default
2 entries were displayed.
```

Step 3. Disable the switch ports for the new stack.

Step 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.

Step 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.

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.

Step 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.

Note: 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, enable SNMP on the bridge:
   `set SNMP enabled`

   If running ONTAP 9.5 or later, ONTAP will use in-band monitoring via the FC connection to the bridge, so you should leave SNMP disabled.

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

   Note: 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.

   Note: 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*

Note: 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.

Step 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`

Note: 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: `SaveConfigurationFirmwareRestart`You are prompted to restart the bridge.

Step 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” on page 8.

Step 10. Cable the disk shelves to the bridges:

a. Daisy-chain the disk shelves in each stack.

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

Result

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.

Note: The bridge SAS port B is disabled.

Step 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	1. 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. 2. 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. 3. Click **Advanced** in the menu. 4. View the connected devices: `sastargets` 5. 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.

Note: 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.

Example

The following output shows that 10 disks are connected:

```
Tgt VendorID ProductID         Type        SerialNumber
  0 NETAPP   X410_S15K6288A15 DISK        3QP1CLE300009940UHJV
  1 NETAPP   X410_S15K6288A15 DISK        3QP1ELF600009940V1BV
  2 NETAPP   X410_S15K6288A15 DISK        3QP1G3EW00009940U2M0
  3 NETAPP   X410_S15K6288A15 DISK        3QP1EWMP00009940U1X5
  4 NETAPP   X410_S15K6288A15 DISK        3QP1FZLE00009940G8YU
  5 NETAPP   X410_S15K6288A15 DISK        3QP1FZLF00009940TZKZ
  6 NETAPP   X410_S15K6288A15 DISK        3QP1CEB400009939MGXL
  7 NETAPP   X410_S15K6288A15 DISK        3QP1G7A900009939FNTT
  8 NETAPP   X410_S15K6288A15 DISK        3QP1FYOT00009940G8PA
  9 NETAPP   X410_S15K6288A15 DISK        3QP1FXW600009940VERQ
```

Step 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 for each remaining bridge.
Not correct	1. Check for loose SAS cables or correct the SAS cabling by repeating Step . 2. Repeat Step .

Step 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:

Notes: Brocade and Cisco switches use different port numbering, as shown in the following tables.

- On Brocade switches, the first port is numbered “0”.
- On Cisco switches, the first port is numbered “1”.

Step 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.

Step 15. If your system does not have disk autoassignment enabled, assign disk drive ownership.

Disk and Aggregate Management Guide

Note: 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.

Step 16. Enable the switch ports for the new stack.

Step 17. 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`

Step 18. If applicable, repeat this procedure for the partner site.

Hot-adding a SAS disk shelf to a stack of SAS disk shelves

You can hot-add one or more SAS disk shelves to a stack of SAS disk shelves.

Preparing to hot-add SAS disk shelves

Before you begin

- Your system must be a supported configuration and must be running a supported version of ONTAP.
- All disk drives and disk shelves in the system must be running the latest firmware version.

  Upgrade, Revert or Downgrade Guide

Hot-adding a disk shelf

You can hot-add a disk shelf when you want to increase storage without any reduction in performance.

Before you begin

- The system must meet all of the requirements in "Preparing to hot-add SAS disk shelves" on page 97.
- Your environment must meet one of the following scenarios to hot-add a shelf:
  - You have two FibreBridge 7600N bridges connected to a stack of SAS disk shelves.
  - You have one FibreBridge 7600N bridge and one FibreBridge 6500N bridge connected to a stack of SAS disk shelves.
  - You have two FibreBridge 6500N bridges 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.

  Note: Each pair of FibreBridge 7600N bridges can support up to four stacks.

Attention: 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.

Step 1. Properly ground yourself.

Step 2. Verify disk shelf connectivity from the system console of either controller:

```
sysconfig -v
```

Result

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 7600N bridges results in the following output:

```
FC-to-SAS Bridge:
cisco_A_1-1:9.126L0: ATTO  FibreBridge7600N 2.10  FB7600N100189
cisco_A_1-2:1.126L0: ATTO  FibreBridge7600N 2.10  FB7600N100162
```

- Each disk shelf on a separate line under each FC port to which it is visible:

```
Shelf    0: IOM6  Firmware rev. IOM6 A: 0173 IOM6 B: 0173
Shelf    1: IOM6  Firmware rev. IOM6 A: 0173 IOM6 B: 0173
```

- Each disk drive on a separate line under each FC port to which it is visible:

```
cisco_A_1-1:9.126L1   : NETAPP   X421_HCOBD450A10 NA01 418.0GB (879097968 520B/sect)
cisco_A_1-1:9.126L2   : NETAPP   X421_HCOBD450A10 NA01 418.0GB (879097968 520B/sect)
```

Step 3. Check whether your system has disk autoassignment enabled from the console of either controller:

```
storage disk option show
```

Result

The autoassignment policy is shown in the Auto Assign column.

```
Node         BKg. FW. Upd.  Auto Copy   Auto Assign  Auto Assign Policy
-----------  -------------  ----------  -----------  ------------------
node_A_1              on           on           on           default
node_A_2              on           on           on           default
2 entries were displayed.
```

Step 4. If your system does not have disk autoassignment 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 Guide

Note: 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.

Note: Disk drives must not be added to aggregates or volumes until the disk drive and disk shelf firmware have been updated.

Step 5. Install and cable the disk shelf:

Note the following considerations:

- For FibreBridge 6500N bridges:

  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 FibreBridge 7600N bridges:

  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.

   Note: 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).

Step 6. 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`

Step 7. If you are hot-adding more than one disk shelf, repeat the steps, starting with step , for each disk shelf that you are hot-adding.

Chapter 10. 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.

Before you begin

- 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.

  Attention: 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 Guide
- You must have removed drive ownership after removing the aggregates from the drives in the shelves that you are removing.

  Disk and Aggregate Management Guide
- 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.

About this task

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

Step 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`

© Copyright Lenovo 2019, 2023

Step 2. Set the privilege level to advanced:
`set -privilege advanced`

Step 3. Verify that no mailbox drive is on the shelves:
`storage failover mailbox-disk show`

Step 4. Remove the shelf according to the steps for the relevant scenario.

Scenario 1: To remove an aggregate when the shelf contains either unmirrored, mirrored, or both types of aggregate...

Steps:

1. Use the `storage aggregate delete -aggregate` *aggregate name* command to remove the aggregate.
2. Use the standard procedure to remove ownership of all drives in that shelf, and then physically remove the shelf.

   Follow the instructions in the Service Guide for DM120S, DM240S, and DM600S to hot-remove expansion enclosures.

Scenario 2: To remove an aggregate when the shelf contains either unmirrored, mirrored, or both types of aggregate...

Steps: To remove a plex from a mirrored aggregate, you need to unmirror the aggregate.

1. Identify the plex that you want to remove by using the `run -node local sysconfig -r` command.

   In the following example, you can identify the plex from the line **`Plex/dpg_mcc_8020_13_a1_aggr1/plex0`**. In this case, the plex to specify is *plex0*.

```
dpgmcc_8020_13_a1a2::storage aggregate> run -node local sysconfig -r
*** This system has taken over dpg-mcc-8020-13-a1
Aggregate dpg_mcc_8020_13_a1_aggr1 (online, raid_dp, mirrored) (block checksums)
 Plex /dpg_mcc_8020_13_a1_aggr1/plex0 (online, normal, active, pool0)
  RAID group /dpg_mcc_8020_13_a1_aggr1/plex0/rg0 (normal, block checksums)
   RAID Disk Device                          HA  SHELF BAY CHAN Pool Type  RPM  Used (MB/blks)    Phys (MB/blks)
   --------- ------                          ------------- ---- ---- ---- ----- --------------    --------------
   dparity   mcc-cisco-8Gb-fab-2:1-1.126L16  0c    32  15  FC:B   0   SAS 15000 272000/557056000  274845/562884296
   parity    mcc-cisco-8Gb-fab-2:1-1.126L18  0c    32  17  FC:B   0   SAS 15000 272000/557056000  274845/562884296
   data      mcc-cisco-8Gb-fab-2:1-1.126L19  0c    32  18  FC:B   0   SAS 15000 272000/557056000  274845/562884296
   data      mcc-cisco-8Gb-fab-2:1-1.126L21  0c    32  20  FC:B   0   SAS 15000 272000/557056000  274845/562884296
   data      mcc-cisco-8Gb-fab-2:1-1.126L22  0c    32  21  FC:B   0   SAS 15000 272000/557056000  274845/562884296

 Plex /dpg_mcc_8020_13_a1_aggr1/plex1 (online, normal, active, pool1)
  RAID group /dpg_mcc_8020_13_a1_aggr1/plex1/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-3:1-1.126L37  0d    34  10  FC:A   1   SAS 15000 272000/557056000  280104/573653840
   parity    mcc-cisco-8Gb-fab-3:1-1.126L14  0d    33  13  FC:A   1   SAS 15000 272000/557056000  280104/573653840
   data      mcc-cisco-8Gb-fab-3:1-1.126L41  0d    34  14  FC:A   1   SAS 15000 272000/557056000  280104/573653840
   data      mcc-cisco-8Gb-fab-3:1-1.126L15  0d    33  14  FC:A   1   SAS 15000 272000/557056000  280104/573653840
   data      mcc-cisco-8Gb-fab-3:1-1.126L45  0d    34  18  FC:A   1   SAS 15000 272000/557056000  280104/573653840
```

2. Use the `storage aggregate plex delete -aggregate` *aggr_name* `-plex` *plex_name* command to remove the plex.

   plex defines the plex name, such as *plex3* or *plex6*.

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

   Follow the instructions in the Service Guide for DM120S, DM240S, and DM600S to hot-remove expansion enclosures.

Chapter 11. 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.

About this task

Note: 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.

Before you begin

Check the overall health of the configuration. If the system does not appear healthy, address the issue first before proceeding.

Step 1. Offline the affected plexes on the controller where they are hosted:
`aggr offline` *plex_name*

Example

The example shows the commands for offlining plexes for a controller running cMode.

```
Cluster_A_1::> storage aggregate plex offline -aggr aggr1 -plex plex2
Cluster_A_1::> storage aggregate plex offline -aggr aggr2 -plex plex6
Cluster_A_1::> storage aggregate plex offline -aggr aggr3 -plex plex1
```

Step 2. Verify that the plexes are offline:
`aggr status -r`*aggr_name*

Example

The example shows the commands for verifying that the aggregates are offline for a controller running cMode.

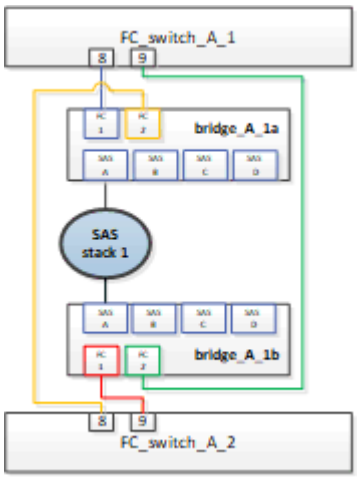
```
Cluster_A_1::> storage aggregate show -aggr aggr1
Cluster_A_1::> storage aggregate show -aggr aggr2
Cluster_A_1::> storage aggregate show -aggr aggr3
```

Step 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:

© Copyright Lenovo 2019, 2023



Switch ports 8 and 9 on each switch connect the bridges to the network.

The following example shows ports 8 and 9 being disabled on a Brocade switch.

```
FC_switch_A_1:admin> portDisable 8
FC_switch_A_1:admin> portDisable 9

FC_switch_A_2:admin> portDisable 8
FC_switch_A_2:admin> portDisable 9
```

The following example shows port 8 and 9 being disabled on a Cisco switch.

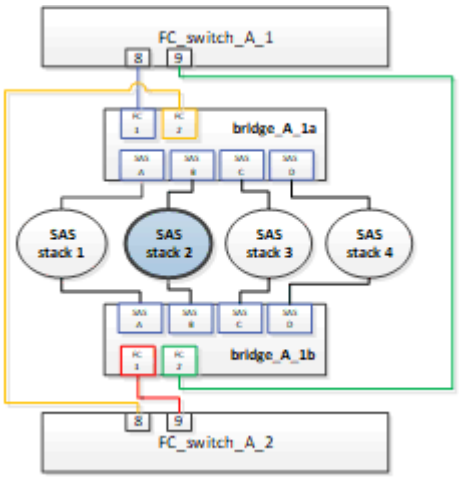
```
FC_switch_A_1# conf t
FC_switch_A_1(config)# int fc1/8
FC_switch_A_1(config)# shut
FC_switch_A_1(config)# int fc1/9
FC_switch_A_1(config)# shut
FC_switch_A_1(config)# end

FC_switch_A_2# conf t
FC_switch_A_2(config)# int fc1/8
FC_switch_A_2(config)# shut
FC_switch_A_2(config)# int fc1/9
FC_switch_A_2(config)# shut
FC_switch_A_2(config)# end
```

If the bridges are connecting two or more SAS stacks, disable the SAS ports connecting the bridges to the target shelf:
`SASportDisable` *port number*

Note: This is not applicable to FibreBridge 6500 bridges.

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.
```

Step 4. If you previously disabled the switch ports, verify that they are disabled:
`switchShow`

Example

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
```

Example

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
```

Step 5. Wait for ONTAP to realize that the disk is missing.

Step 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.

Step 1. Remove all disks and disconnect all cables from the shelf that is being replaced.

Step 2. Remove the shelf modules.

Step 3. Insert the new shelf.

Step 4. Insert the new disks into the new shelf.

Step 5. Insert the shelf modules.

Step 6. Cable the shelf (SAS or Power).

Step 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.

Step 1. Verify that the shelf powers properly and the links on the IOM modules are present.

Step 2. Enable the switch ports or SAS port according to the following scenarios:

Option	Step

If you previously disabled switch ports	1. Enable the switch ports: `portEnable` *port number*
	
	The example shows the switch port being enabled on a Brocade switch.
	
	`Switch_A_1:admin> portEnable 6`
	`Switch_A_2:admin> portEnable 6`
	
	The example shows the switch port being enabled on a Cisco switch.
	
	`Switch_A_1# conf t`
	`Switch_A_1(config)# int fc1/6`
	`Switch_A_1(config)# no shut`
	`Switch_A_1(config)# end`
	
	`Switch_A_2# conf t`
	`Switch_A_2(config)# int fc1/6`
	`Switch_A_2(config)# no shut`
	`Switch_A_2(config)# end`

If you previously disabled a SAS port	1. 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.`
	
	Note: This is not applicable to ATTO 6500 FibreBridges.

Step 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`

Example

The example shows the `switchShow` command for verifying that a Brocade switch is online.

```
Switch_A_1:admin> SwitchShow
Switch_A_2:admin> SwitchShow
```

Example

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
```

Note: After several minutes, ONTAP detects that new disks have been inserted and displays a message for each new disk.

Step 4. Verify that the disks have been detected by ONTAP:
`sysconfig -a`

Step 5. Online the plexes that were offline earlier:
`aggr online`*plex_name*

Example

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.

Note: You can monitor the progress of resynchronization using the `aggr status -r`*aggr_name* command.

Chapter 12. When to migrate root volumes to a new destination

You might need to move root volumes to another root aggregate within a four-node MetroCluster configuration.

Migrating root volumes within a four-node MetroCluster configuration

To migrate root volumes to a new root aggregate within a four-node MetroCluster configuration, you can use the system node migrate-root command while meeting the following requirements.

- You can use system node migrate-root to move root aggregates within a four-node MetroCluster configuration.
- All root aggregates must be mirrored.
- You can add new shelves on both sites with smaller SSDs to host the root aggregate.
- You must check the SSD limits that the platform supports before attaching new disks.

  Lenovo Press
- If you move the root aggregate to smaller disks, you need to accommodate the minimum root volume size of the platform to ensure all core files are saved.

Note: The four-node procedure can also be applied to an eight-node configuration.

© Copyright Lenovo 2019, 2023

Chapter 13. Moving a metadata volume in MetroCluster configurations

You can move a metadata volume from one aggregate to another aggregate in a MetroCluster configuration. You might want to move a metadata volume when the source aggregate is decommissioned or unmirrored, or for other reasons that make the aggregate ineligible.

Before you begin

- You must have cluster administrator privileges to perform this task.
- The target aggregate must be mirrored and should not be in the degraded state.
- The available space in the target aggregate must be larger than the metadata volume that you are moving.

Step 1. Set the privilege level to advanced:
`set -privilege advanced`

Step 2. Identify the metadata volume that should be moved:
`volume show MDV_CRS*`

Example

```
Cluster_A::*> volume show MDV_CRS*
Vserver   Volume       Aggregate    State      Type       Size  Available Used%
--------- ------------ ------------ ---------- ---- ---------- ---------- -----
Cluster_A
          MDV_CRS_14c00d4ac9f311e7922800a0984395f1_A
                       Node_A_1_aggr1
                                    online     RW         10GB     9.50GB    5%
Cluster_A
          MDV_CRS_14c00d4ac9f311e7922800a0984395f1_B
                       Node_A_2_aggr1
                                    online     RW         10GB     9.50GB    5%
Cluster_A
          MDV_CRS_15035e66c9f311e7902700a098439625_A
                       Node_B_1_aggr1
                                    -          RW             -          -     -
Cluster_A
          MDV_CRS_15035e66c9f311e7902700a098439625_B
                       Node_B_2_aggr1
                                    -          RW             -          -     -
4 entries were displayed.

Cluster_A::>
```

Step 3. Identify an eligible target aggregate:
`metrocluster check config-replication show-aggregate-eligibility`

Example

The following command identifies the aggregates in cluster_A that are eligible to host metadata volumes:

```
Cluster_A::*> metrocluster check config-replication show-aggregate-eligibility

Aggregate Hosted Config Replication Vols Host Addl Vols Comments
------------ ------------------------------------------ -------------- --------
Node_A_1_aggr0 - false Root Aggregate
Node_A_2_aggr0 - false Root Aggregate
Node_A_1_aggr1 MDV_CRS_1bc7134a5ddf11e3b63f123478563412_A true -
```

© Copyright Lenovo 2019, 2023

```
Node_A_2_aggr1 MDV_CRS_1bc7134a5ddf11e3b63f123478563412_B true -
Node_A_1_aggr2 - true
Node_A_2_aggr2 - true
Node_A_1_Aggr3 - false Unable to determine available space of aggregate
Node_A_1_aggr5 - false Unable to determine mirror configuration
Node_A_2_aggr6 - false Mirror configuration does not match requirement
Node_B_1_aggr4 - false NonLocal Aggregate
```

Note: In the previous example, Node_A_1_aggr2 and Node_A_2_aggr2 are eligible.

Step 4. Start the volume move operation:

`volume move start -vserver` *svm_name* `-volume` *metadata_volume_name* `-destination-aggregate` *destinataion_aggregate_name*

Example

The following command moves metadata volume **MDV_CRS_14c00d4ac9f311e7922800a0984395f1** from **aggregateNode_A_1_aggr1** to **aggregateNode_A_1_aggr2**:

```
Cluster_A::*> volume move start -vserver svm_cluster_A -volume MDV_CRS_14c00d4ac9f311e7922800a0984395f1
-destination-aggregate aggr_cluster_A_02_01

Warning: You are about to modify the system volume
         "MDV_CRS_9da04864ca6011e7b82e0050568be9fe_A".  This may cause severe
         performance or stability problems.  Do not proceed unless directed to
         do so by support.  Do you want to proceed? {y|n}: y
[Job 109] Job is queued: Move "MDV_CRS_9da04864ca6011e7b82e0050568be9fe_A" in Vserver
"svm_cluster_A" to aggregate "aggr_cluster_A_02_01".
Use the "volume move show -vserver svm_cluster_A -volume
MDV_CRS_9da04864ca6011e7b82e0050568be9fe_A" command to view the status of this operation.
```

Step 5. Verify the state of the volume move operation:

`volume move show -volume vol_constituent_name`

Step 6. Return to the admin privilege level:

`set -privilege admin`

Chapter 14. Renaming a cluster in MetroCluster configurations

Renaming a cluster in a MetroCluster configuration involves making the changes, and then verifying on both the local and remote clusters that the change took effect correctly.

Step 1. View the cluster names using the `metrocluster node show` command:

Example

```
cluster_1::*> metrocluster node show
DR                               Configuration  DR
Group Cluster Node               State          Mirroring Mode
----- ------- ------------------ -------------- --------- --------------------
1     cluster_1
              node_A_1           configured     enabled   normal
              node_A_2           configured     enabled   normal
      cluster_2
              node_B_1           configured     enabled   normal
              node_B_2           configured     enabled   normal
4 entries were displayed.
```

Step 2. Rename the cluster:
`cluster identity modify -name` *new_name*

Example

In the following example, the *cluster_1* cluster is renamed *cluster_A*:

```
cluster_1::*> cluster identity modify -name cluster_A
```

Step 3. Verify on the local cluster that the renamed cluster is running normally:
`metrocluster node show`

Example

In the following example, the newly renamed *cluster_A* is running normally:

```
cluster_A::*> metrocluster node show
DR                               Configuration  DR
Group Cluster Node               State          Mirroring Mode
----- ------- ------------------ -------------- --------- --------------------
1     cluster_A
              node_A_1           configured     enabled   normal
              node_A_2           configured     enabled   normal
      cluster_2
              node_B_1           configured     enabled   normal
              node_B_2           configured     enabled   normal
4 entries were displayed.
```

Step 4. Rename the remote cluster:
`cluster peer modify-local-name -name` *cluster_2* `-new-name` *cluster_B*

Example

In the following example, *cluster_2* is renamed *cluster_B*:

```
cluster_2::> cluster peer modify-local-name -name cluster_2 -new-name
cluster_B
```

Step 5. Verify on the remote cluster that the local cluster was renamed and is running normally:
`metrocluster node show`

Example

© Copyright Lenovo 2019, 2023

In the following example, the newly renamed *cluster_B* is running normally:

```
cluster_B::*> metrocluster node show
DR                               Configuration  DR
Group Cluster Node               State          Mirroring Mode
----- ------- ------------------ -------------- --------- --------------------
1     cluster_B
              node_B_1           configured     enabled   normal
              node_B_2           configured     enabled   normal
      cluster_A
              node_A_1           configured     enabled   normal
              node_A_2           configured     enabled   normal
4 entries were displayed.
```

Step 6. Repeat these steps for each cluster that you want to rename.

Chapter 15. Powering off and powering on a data center

You must know how to power off and power on a data center for the purpose of site maintenance or to relocate a site to another location.

About this task

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

Note: This procedure is for use in both MetroCluster IP and FC configurations.

Powering off a MetroCluster site

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

Step 1. Before you begin, check that any non-mirrored aggregates at the site are offline.

Step 2. Verify the operation of the MetroCluster configuration in ONTAP:

a. Check whether the system is multipathed:
`node run -node` *node-name* `sysconfig -a`

b. Check for any health alerts on both clusters:
`system health alert show`

c. Confirm the MetroCluster configuration and that the operational mode is normal:
`metrocluster show`

d. Perform a MetroCluster check:
`metrocluster check run`

e. Display the results of the MetroCluster check:
`metrocluster check show`

f. Check for any health alerts on the switches (if present):
`storage switch show`

Step 3. Enter the following command to implement the switchover:
`metrocluster switchover`
The operation can take several minutes to complete.

Attention: In MetroCluster FC configurations, 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 may 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.

Step 4. Monitor and verify the completion of the switchover:
`metrocluster operation show`

Example

```
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
```

© Copyright Lenovo 2019, 2023

```
  Operation: Switchover
Start time: 10/4/2012 19:04:13
     State: successful
  End time: 10/4/2012 19:04:22
    Errors: -
```

Step 5. If you are in a MetroCluster IP configuration running ONTAP 9.6 or later, wait for the disaster site plexes to come online and the healing operations to automatically complete. In MetroCluster IP configurations running earlier versions of ONTAP, the disaster site nodes do not automatically boot to ONTAP and the plexes remain offline.

Step 6. Depending on your configuration and ONTAP version, identify and move offline affected plexes that are located at the disaster site.

Notes: This step is required in the following configurations:

- All MetroCluster FC configurations.
- MetroCluster IP configurations running ONTAP version 9.6 or later.

You should move the following plexes offline:

- Non-mirrored plexes residing on disks located at the disaster site. If you do not move the non-mirrored plexes at the disaster site offline, an outage might result when the disaster site is later powered off.
- Mirrored plexes residing on disks located at the disaster site for aggregate mirroring.

Once they are moved offline, the plexes are inaccessable.

Example

1. 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.

   Note: “Home” refers to the “home” owner, not the “current” owner.

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

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

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

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

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

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

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

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

```
Cluster_A::>
```

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

Node	Disks in pool	Should the disks be set offline?	Example of plexes to be moved offline
Node _A_1 and Node _A_2	Disks in pool 0	No. Disks are local to cluster A.	-
	Disks in pool 1	Yes. Disks are remote to cluster A.	Node_A_1_aggr0/plex4 Node_A_1_aggr1/plex1 Node_A_2_aggr0/plex4 Node_A_2_aggr1/plex1
Node _B_1 and Node _B_2	Disks in pool 0	Yes. Disks are remote to cluster A.	Node_B_1_aggr1/plex0 Node_B_1_aggr0/plex0 Node_B_2_aggr0/plex0 Node_B_2_aggr1/plex0
	Disks in pool 1	No. Disks are local to cluster A.	-

2. Move the affected plexes offline:
   ```
   storage aggregate plex offline
   storage aggregate plex offline -aggregate Node_B_1_aggr0 -plex plex0
   ```

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

Step 7. Persistently offline the switchports according to the switch type.

Note: This step is only required for MetroCluster FC configurations. Skip this step if your configuration is a MetroCluster IP configuration or a stretched MetroCluster configuration with FC backend switches.

Switch type	Action

If the FC switches are Brocade switches...	1. Use the `portpersistentdisable` *port* command to persistently disable the ports as shown in the following example. This must be done on both switches at the surviving site.
	
	`Switch_A_1:admin> portcfgpersistentdisable 14`
	`Switch_A_1:admin> portcfgpersistentdisable 15`
	`Switch_A_1:admin>`
	
	2. Verify that the ports are disabled using the `switchshow` command shown in the following example:
	
	`Switch_A_1:admin> switchshow`
	`switchName:Switch_A_1`
	`switchType:109.1`
	`switchState:Online`
	`switchMode:Native`
	`switchRole:Principal`
	`switchDomain:2`
	`switchId:fffc02`
	`switchWwn:10:00:00:05:33:88:9c:68`
	`zoning:ON (T5_T6)`
	`switchBeacon:OFF`
	`FC Router:OFF`
	`FC Router BB Fabric ID:128`
	`Address Mode:0`
	
	`Index Port Address Media Speed State Proto`
	`==========`
	`...`
	`14 14 020e00 id 16G No_Light FC Disabled (Persistent)`
	`15 15 020f00 id 16G No_Light FC Disabled (Persistent)`
	`...`
	`Switch_A_1:admin>`

If the FC switches are Cisco switches...	1. Use the `interface` command to persistently disable the ports. The following example shows ports 14 and 15 being disabled:
	
	`Switch_A_1# conf t`
	`Switch_A_1(config)# interface fc1/14-15`
	`Switch_A_1(config)# shut`
	
	`Switch_A_1(config-if)# end`
	`Switch_A_1# copy running-config startup-config`
	
	2. 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`

Step 8. Power off the site.
The following equipment needs to be turned off in no specific order:

Configuration type	Equipment to be powered off
In a MetroCluster IP configuration, power off...	• MetroCluster IP switches • Storage controllers • Storage shelves
In a MetroCluster FC configuration, power off...	• MetroCluster FC switches • Storage controllers • Storage shelves • Atto FibreBridges (if present)

Relocating the powered-off site of the MetroCluster

Once 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.

About this task

- 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.

Step 1. Make sure that the cabling for all components is carefully recorded so that it can be correctly reconnected at the new location.

Step 2. Physically relocate all the hardware, storage controllers, FC and IP switches, FibreBridges, and storage shelves.

Step 3. Configure the ISL ports and verify the intersite connectivity.

a. Power on the FC and IP switches.

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

Step 4. Enable the ports.

Note: This step is only required in MetroCluster FC configurations. You can skip this step if your configuration is a MetroCluster IP configuration.

Enable the ports according to the correct switch types in the following table:

Switch type	Command

If the FC Switches are Brocade switche-s...	1. Use the `portcfgpersistentenable` *port number* command to persistently enable the port. This must be done on both switches at the surviving site.
	The following example shows ports 14 and 15 being enabled on Switch_A_1.
	`switch_A_1:admin> portcfgpersistentenable 14`
	`switch_A_1:admin> portcfgpersistentenable 15`
	`switch_A_1:admin>`
	2. Verify that the switch port is enabled: `switchshow` The following example shows that ports 14 and 15 are enabled:
	`switch_A_1:admin> switchshow`
	`switchName:Switch_A_1`
	`switchType:109.1`
	`switchState:Online`
	`switchMode:Native`
	`switchRole:Principal`
	`switchDomain:2`
	`switchId:fffc02`
	`switchWwn:10:00:00:05:33:88:9c:68`
	`zoning:ON (T5_T6)`
	`switchBeacon:OFF`
	`FC Router:OFF`
	`FC Router BB Fabric ID:128`
	`Address Mode:0`
	`Index Port Address Media Speed State Proto`
	`==========`
	`...`
	`14 14 020e00 id 16G Online FC E-Port 10:00:00:05:33:86:89:cb "Switch_A_1"`
	`15 15 020f00 id 16G Online FC E-Port 10:00:00:05:33:86:89:cb "Switch_A_1" (downstream)`
	`...`
	`switch_A_1:admin>`

If the FC Switches are Cisco switche-s...	1. Enter the `interface` command to enable the port.
	The following example shows ports 14 and 15 being enabled on Switch_A_1.
	`switch_A_1# conf t`
	`switch_A_1(config)# interface fc1/14-15`
	`switch_A_1(config)# no shut`
	`switch_A_1(config-if)# end`
	`switch_A_1# copy running-config startup-config`
	2. Verify that the switch port is enabled: `show interface brief`
	`switch_A_1# show interface brief`
	`switch_A_1#`

Step 5. Use tools on the switches (as they are available) to verify the intersite connectivity.

Note: You should only proceed if the links are properly configured and stable.

Step 6. Disable the links again if they are found to be stable.
Disable the ports based on whether you are using Brocade or Cisco switches as shown in the following table:

Switch type	Command

If the FC Switches are Brocade switches...	1. Enter the `portcfgpersistentdisable` *port number* command to persistently disable the port.
	This must be done on both switches at the surviving site. The following example shows ports 14 and 15 being disabled on Switch_A_1:
	`switch_A_1:admin> portpersistentdisable 14`
	`switch_A_1:admin> portpersistentdisable 15`
	`switch_A_1:admin>`
	2. 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>`

If the FC Switches are Cisco switches...	1. Disable the port using the `interface` command.
	The following example shows ports fc1/14 and fc1/15 being disabled on Switch A_1:
	`switch_A_1# conf t`
	`switch_A_1(config)# interface fc1/14-15`
	`switch_A_1(config)# shut`
	`switch_A_1(config-if)# end`
	`switch_A_1# copy running-config startup-config`
	2. Verify that the switch port is disabled using the `show interface brief` command.
	`switch_A_1# show interface brief`
	`switch_A_1#`

Powering on the MetroCluster configuration and returning to normal operation

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

Step 1. Power on the switches.

Switches should be powered on 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.

Step 2. Power on the shelves and allow enough time for them to power on completely.

Step 3. Power on the FibreBridge bridges.

Note: You can skip this step if your configuration is a MetroCluster IP configuration.

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 command-line interface (CLI).

Step 4. Enable the ISLs on the FC switches.

Note: Skip this step if your configuration is a MetroCluster IP configuration.

Enable the ports based on whether you are using Brocade or Cisco switches as shown in the following table:

<table>
<tr><th>Switch type</th><th>Command</th></tr>
<tr><td>If the FC Switches are Brocade switches...</td><td>
1. Enter the <code>portcfgpersistentenable</code> port command to persistently enable the ports. This must be done on both switches at the surviving site.
The following example shows ports 14 and 15 being enabled on Switch_A_1:
<pre>Switch_A_1:admin> portcfgpersistentenable 14
Switch_A_1:admin> portcfgpersistentenable 15
Switch_A_1:admin></pre>
2. Verify that the switch port is enabled using the <code>switchshow</code> command:
<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 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></pre>
</td></tr>
<tr><td>If the FC Switches are Cisco switches...</td><td>
1. Use the <code>interface</code> command to enable the ports.
The following example shows port fc1/14 and fc1/15 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 startup-config</pre>
2. Verify that the switch port is disabled:
<pre>switch_A_1# show interface brief
switch_A_1#</pre>
</td></tr>
</table>

Step 5. Verify that the storage is now visible.

Select the appropriate method of determining whether the storage is visible based on whether you have a MetroCluster IP or FC configuration:

Configuration	Step
If your configuration is a MetroCluster IP...	Verify that the local storage is visible from the node Maintenance mode.
If your configuration is a MetroCluster FC...	Verify that the storage is visible from the surviving site. Put the offline plexes back online. This restarts the resync operations and reestablishes the SyncMirror.

Step 6. Reestablish the MetroCluster configuration.

Follow the instructions in the *MetroCluster Disaster and Recovery Guide* to perform healing and switchback operations according to your MetroCluster configuration.

MetroCluster management and disaster recovery

Chapter 16. Powering off an entire MetroCluster IP configuration

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

About this task

Note: Starting 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.

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

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

b. Run the following command: `metrocluster interconnect show`

c. Confirm connectivity to the disks by entering the following command on any one of the MetroCluster nodes: `run local sysconfig -v`

d. Run the following command: `storage port show`

e. Run the following command: `storage switch show`

f. Run the following command: `network interface show`

g. Run the following command: `network port show`

h. Run the following command: `network device-discovery show`

i. Perform a MetroCluster check: `metrocluster check run`

j. Display the results of the MetroCluster check: `metrocluster check show`

k. Run the following command: `metrocluster configuration-settings interface show`

Step 2. If necessary, disable AUSO by modifying the AUSO Failure Domain to `auso-disabled`

Example

```
cluster_A_site_A::*>metrocluster modify -auto-switchover-failure-domain auso-disabled
```

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

Step 3. Verify the change using the command `metrocluster operation show`

Example

```
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: -
```

Step 4. Halt the nodes: `halt`

Example

© Copyright Lenovo 2019, 2023

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

Step 5. Power off the following equipment at the site:

- Storage controllers
- MetroCluster IP switches
- Storage shelves

Step 6. Wait for thirty minutes and then power on all the controllers.

Step 7. After the controllers are powered on, verify the MetroCluster configuration from both sites.To verify the configuration, repeat step 1.

Chapter 17. 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

Note: Starting 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.

Step 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`

Step 2. Disable AUSO by modifying the AUSO Failure Domain to `auso-disabled`

Example

```
cluster_A_site_A::*>metrocluster modify -auto-switchover-failure-domain auso-disabled
```

Step 3. Verify the change using the command `metrocluster operation show`

Example

```
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: -
```

Step 4. Halt the nodes by using the following command: `halt`

For a four-node or eight-node MetroCluster configuration, use the parameter `inhibit-takeover`
`system node halt -node node1_SiteA -inhibit-takeover true -ignore-quorum-warnings true`

For a two-node MetroCluster configuration, use the command:
`system node halt -node node1_SiteA -ignore-quorum-warnings true`

Step 5. Power off the following equipment at the site:

- Storage controllers

© Copyright Lenovo 2019, 2023

- MetroCluster FC switches
- ATTO FibreBridges
- Storage shelves

Step 6. Wait for thirty minutes and then power on all the controllers.

Step 7. After the controllers are powered on, verify the MetroCluster configuration from both sites.To verify the configuration, repeat step 1.

Chapter 18. Reconfiguring the FC switch layout

Before you begin

You must identify the FC switches present in the fabric domain.

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

About this task

This procedure is nondisruptive and takes approximately four hours to complete (excluding rack and stack) when disks are zeroed.

Sending a custom AutoSupport message prior to reconfiguring switches

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

About this task

This task must be performed on each MetroCluster site.

Step 1. Log in to the cluster.

Step 2. 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`

Step 3. Repeat these steps on the partner site.

Verifying the health of the MetroCluster configuration

You should check the health of the MetroCluster configuration to verify proper operation.

Step 1. Verify that the MetroCluster components are healthy:
`metrocluster check run`

Example

```
dpgqa-mcc-funct-8040-0403_siteA::*> metrocluster check run

----
```

Result

The operation runs in the background.

Step 2. After the `metrocluster check run` operation completes, run `metrocluster check show` to view the results.

Result

© Copyright Lenovo 2019, 2023

After approximately five minutes, the following results are displayed:

```
-----------
dpgqa-mcc-funct-8040-0403_siteA::*> metrocluster check show

Last Checked On: 4/7/2019 21:15:05

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

Step 3. To check the status of the running MetroCluster check operation, use the command:
`metrocluster operation history show -job-id 38`

Step 4. Verify that there are no health alerts:
`system health alert show`

Persistently disabling the switches

You must disable the switches in the fabric persistently so that you can modify its configuration.

About this task

You disable the switches by running the commands on the switch command line; the commands used for this are not ONTAP commands.

Step 1. Persistently disable the switch:

For Brocade switches, use the `switchCfgPersistentDisable` command.

For Cisco switches, use the `suspend` command.

Example

The following command disables a Brocade switch persistently:

```
FC_switch_A_1:admin> switchCfgPersistentDisable
```

Example

The following command disables a Cisco switch:

```
vsan [vsna #] suspend
```

Determining the new cabling layout

You must determine the cabling for the new controller modules and any new disk shelves to the existing FC switches.

About this task

This task must be performed at each MetroCluster site.

Step 1. Use the *Fabric-attached MetroCluster Installation and Configuration Guide* to determine the cabling layout for your switch type, using the port usage for an eight-node MetroCluster configuration.

The FC switch port usage must match the usage described in the guide so that the Reference Configuration Files (RCFs) can be used.

Fabric-attached MetroCluster installation and configuration

Note: If your environment cannot be cabled in a way that RCFs can be used, then contact technical support. Do not use this procedure if the cabling cannot use RCFs.

Applying RCF files and recabling the switches

You must apply the appropriate reference configuration (RCF) files to reconfigure your switches to accommodate the new nodes. After you apply the RCF files, you can recable the switches.

Before you begin

The FC switch port usage must match the usage described in the *Fabric-attached MetroCluster Installation and Configuration Guide* so that the RCFs can be used.

Fabric-attached MetroCluster installation and configuration

Step 1. Locate the RCF files for your configuration.

You must use the RCF files that match your switch model.

Lenovo Data Center Support Site

Step 2. Apply the RCF files, following the directions on the Download page and adjusting the ISL settings as needed.

Step 3. Verify that the switch configuration is saved.

Step 4. Cable both of the FC-to-SAS bridges to the FC switches, using the cabling layout you created in the “Determining the new cabling layout” section.

Step 5. Verify that the ports are online:

For Brocade switches, use the `switchshow` command.

For Cisco switches, use the `show interface brief` command.

Step 6. Cable the FC-VI ports from the controllers to the switches.

Step 7. From the existing nodes, verify that the FC-VI ports are online:

```
metrocluster interconnect adapter show
metrocluster interconnect mirror show
```

Persistently enable the switches

You must enable the switches in the fabric persistently.

Step 1. Persistently enable the switch:

For Brocade switches, use the `switchCfgPersistentenable` command.

For Cisco switches, use the `no suspend` command.

Example

The following command persistently enables a Brocade switch:

```
FC_switch_A_1:admin> switchCfgPersistentenable
```

Example

The following command enables a Cisco switch:

```
vsan [vsna #]no suspend
```

Verifying switchover, healing, and switchback

You should verify the switchover, healing, and switchback operations of the MetroCluster configuration.

Step 1. Use the procedures for negotiated switchover, healing, and switchback that are mentioned in the *MetroCluster Management and Disaster Recovery Guide*.
MetroCluster management and disaster recovery

Chapter 19. Port assignments for FC switches

You need to verify that you are using the specified port assignments when you cable the FC switches.

Port assignments for systems using two initiator ports

You can configure DM7000H and AFA DM7000F systems using a single initiator port for each fabric and two initiator ports for each controller.

You can follow the cabling for the FibreBridge 6500N bridge or FibreBridge 7600N bridge using only one FC port (FC1 or FC2). Instead of using four initiators, connect only two initiators and leave the other two that are connected to the switch port empty.

You must apply the correct RCF file for the FibreBridge 6500N bridge's configuration.

If zoning is performed manually, then follow the zoning used for a FibreBridge 6500N or a FibreBridge 7600N bridge using one FC port (FC1 or FC2). In this scenario, one initiator port rather than two is added to each zone member per fabric.

You can change the zoning or perform an upgrade from a FibreBridge 6500 to a FibreBridge 7600 using the procedure in the section “Hot-swapping a FibreBridge 6500N bridge with a FibreBridge 7600N bridge” on page 24.

The following table shows port assignments for FC switches when using ONTAP 9.4 and later.

Configurations using FibreBridge 6500N bridges or FibreBridge 7600N using one FC port (FC1 or FC2) only			
MetroCluster 1 or DR Group 1			
		Lenovo switch models B6505, B6510, DB620S, and DB610S	
Component	Port	Connects to FC switch...	Connects to switch port...
controller_x_1	FC-VI port a	1	0
	FC-VI port b	2	0
	FC-VI port c	1	1
	FC-VI port d	2	1
	HBA port a	1	2
	HBA port b	2	2
	HBA port c	-	-
	HBA port d	-	-
Stack 1	bridge_x_1a	1	8
	bridge_x_1b	2	8
Stack y	bridge_x_*y*a	1	11
	bridge_x_*y*b	2	11

The following table shows port assignments for FC switches when using ONTAP 9.4.

© Copyright Lenovo 2019, 2023

MetroCluster two-node configuration			
Component	Port	Lenovo B6505 or B6510 FC SAN switch	
		FC_switch_x_1	FC_switch_x_2
controller_x_1	FC-VI port a	0	-
	FC-VI port b	-	0
	HBA port a	1	-
	HBA port b	-	1
	HBA port c	2	-
	HBA port d	-	2

Port assignments for FC switches

You need to verify that you are using the specified port assignments when you cable the FC switches.

Ports that are not used for attaching initiator ports, FC-VI ports, or ISLs can be reconfigured to act as storage ports. However, if the supported RCFs are being used, the zoning must be changed accordingly.

If the supported RCFs are used, ISL ports might not connect to the same ports shown here and might need to be reconfigured manually.

Configurations running ONTAP releases should use the port assignments shown here.

Overall cabling guidelines

You should be aware of the following guidelines when using the cabling tables:

- The Brocade and Cisco switches use different port numbering:
  - On Brocade switches, the first port is numbered 0.
  - On Cisco switches, the first port is numbered 1.
- The cabling is the same for each FC switch in the switch fabric.
- AFA DM7000F and DM7000H storage systems can be ordered with one of two options for FC-VI connectivity:
  - Onboard ports 0e and 0f configured in FC-VI mode.
  - Ports 1a and 1b on an FC-VI card in slot 1.
- AFA DM7100F and DM7100H storage systems use ports 2a and 2b for FC-VI connectivity.
- If you have two MetroCluster configurations sharing ISLs, use the same port assignments as that for an eight-node MetroCluster cabling.

  The number of ISLs you cable may vary depending on your site's requirements.

  See the section on ISL considerations.

Brocade port usage for controllers in a MetroCluster configuration

The following tables show port usage on Brocade switches. The tables show the maximum supported configuration, with eight controller modules in two DR groups. For smaller configurations, ignore the rows for the additional controller modules. Note that eight ISLs are supported only on the Lenovo B6510 FC SAN switch, DB620S, DB630S, and DB720 switches.

Note: Port usage for the Lenovo B6505 FC SAN switch and Lenovo ThinkSystem DB610S 32G FC SAN switch switches in an eight-node MetroCluster configuration is not shown. Due to the limited number of ports, port assignments must be made on a site-by-site basis depending on the controller module model and the number of ISLs and bridge pairs in use.

Configurations using FibreBridge 6500N bridges or FibreBridge 7600N using one FC port (FC1 or FC2) only				
MetroCluster 1 or DR Group 1				
		Lenovo switch models B6505, B6510, DB610S, DB620S, DB630S		Lenovo switch model DB720
Component	Port	Connects to FC switch...	Connects to switch port...	Connects to switch port...
Stack 1	bridge_x_1a	1	8	10
	bridge_x_1b	2	8	10
Stack 2	bridge_x_2a	1	9	11
	bridge_x_2b	2	9	11
Stack 3	bridge_x_3a	1	10	14
	bridge_x_4b	2	10	14
Stack y	bridge_x_*y*a	1	11	15
	bridge_x_*y*b	2	11	15

Notes:

- On DB620S and DB630S switches, additional bridges can be cabled to ports 12 - 17, 20 and 21.
- On DB610S switches, additional bridges can be cabled to ports 12 - 19.
- On DB720 switches, additional bridges can be cabled to ports 16 - 17, 20 and 21.

Configurations using FibreBridge 6500N bridges or FibreBridge 7500N or 7600N using one FC port (FC1 or FC2) only					
MetroCluster 2 or DR Group 2					
			Lenovo switch model		
Component	Port	Connects to FC_switch...	B6510	DB620S, DB630S	DB720
controller_x_3	FC-VI port a	1	24	18	18
	FC-VI port b	2	24	18	18
	FC-VI port c	1	25	19	19
	FC-VI port d	2	25	19	19
	HBA port a	1	26	24	26
	HBA port b	2	26	24	26
	HBA port c	1	27	25	27
	HBA port d	2	27	25	27
controller_x_4	FC-VI port a	1	28	22	22
	FC-VI port b	2	28	22	22

Configurations using FibreBridge 6500N bridges or FibreBridge 7500N or 7600N using one FC port (FC1 or FC2) only					
MetroCluster 2 or DR Group 2					
			Lenovo switch model		
Component	Port	Connects to FC_switch...	B6510	DB620S, DB630S	DB720
	FC-VI port c	1	29	23	23
	FC-VI port d	2	29	23	23
	HBA port a	1	30	28	30
	HBA port b	2	30	28	30
	HBA port c	1	31	29	31
	HBA port d	2	32	29	31
Stack 1	bridge_x_51a	1	32	26	32
	bridge_x_51b	2	32	26	32
Stack 2	bridge_x_52a	1	33	27	33
	bridge_x_52b	2	33	27	33
Stack 3	bridge_x_53a	1	34	30	34
	bridge_x_54b	2	34	30	34
Stack *y*	bridge_x_*y*a	1	35	31	35
	bridge_x_*y*b	2	35	31	35
Notes: • On DB720 switches, additional bridges can be cabled to ports 36-39.					

Configurations using FibreBridge 7500N or 7600N using both FC ports (FC1 and FC2)					
MetroCluster 1 or DR Group 1					
			Lenovo switch models B6505, B6510, DB610S, DB620S, DB630S		Lenovo switch DB720
Component		Port	Connects to FC_ switch...	Connects to switch port...	Connects to switch port...
Stack 1	bridge_x_1a	FC1	1	8	10
		FC2	2	8	10
	bridge_x_1B	FC1	1	9	11
		FC2	2	9	11
Stack 2	bridge_x_2a	FC1	1	10	14
		FC2	2	10	14
	bridge_x_2B	FC1	1	11	15
		FC2	2	11	15
Stack 3	bridge_x_3a	FC1	1	12*	16

Configurations using FibreBridge 7500N or 7600N using both FC ports (FC1 and FC2)					
MetroCluster 1 or DR Group 1					
Component		Port	Lenovo switch models B6505, B6510, DB610S, DB620S, DB630S		Lenovo switch DB720
			Connects to FC_ switch...	Connects to switch port...	Connects to switch port...
		FC2	2	12*	16
	bridge_x_3B	FC1	1	13*	17
		FC2	2	13*	17
Stack *y*	bridge_x_*y*a	FC1	1	14*	20
		FC2	2	14*	20
	bridge_x_*y*b	FC1	1	15*	21
		FC2	2	15*	21
* - Ports 12 through 15 are reserved for the second MetroCluster or DR group on the Brocade 7840 switch. **Note:** Additional bridges can be cabled to ports 16, 17, 20 and 21 in DB620S and DB630S switches.					

Configurations using FibreBridge 7500N or 7600N using both FC ports (FC1 and FC2)					
MetroCluster 2 or DR Group 2					
Component	Port	Brocade switch model			
		Connects to FC_switch...	B6510	DB620S, DB630S	DB720
controller_x_3	FC-VI port a	1	24	18	18
	FC-VI port b	2	24	18	18
	FC-VI port c	1	25	19	19
	FC-VI port d	2	25	19	19
	HBA port a	1	26	24	26
	HBA port b	2	26	24	26
	HBA port c	1	27	25	27
	HBA port d	2	27	25	27
controller_x_4	FC-VI port a	1	28	22	22
	FC-VI port b	2	28	22	22
	FC-VI port c	1	29	23	23
	FC-VI port d	2	29	23	23

Configurations using FibreBridge 7500N or 7600N using both FC ports (FC1 and FC2)						
MetroCluster 2 or DR Group 2						
Component		Port	Brocade switch model			
			Connects to FC_switch...	B6510	DB620S, DB630S	DB720
		HBA port a	1	30	28	30
		HBA port b	2	30	28	30
		HBA port c	1	31	29	31
		HBA port d	2	31	29	31
Stack 1	bridge_x_51a	FC1	1	32	26	32
		FC2	2	32	26	32
	bridge_x_51b	FC1	1	33	27	33
		FC2	2	33	27	33
Stack 2	bridge_x_52a	FC1	1	34	30	34
		FC2	2	34	30	34
	bridge_x_52b	FC1	1	35	31	35
		FC2	2	35	31	35
Stack 3	bridge_x_53a	FC1	1	36	32	36
		FC2	2	36	32	36
	bridge_x_53b	FC1	1	37	33	37
		FC2	2	37	33	37
Stack *y*	bridge_x_5*y*a	FC1	1	38	34	38
		FC2	2	38	34	38
	bridge_x_5*y*b	FC1	1	39	35	39
		FC2	2	39	35	39
Note: Additional bridges can be cabled to ports 36 to 39 in DB620S and DB630S switches.						

Brocade port usage for ISLs in a MetroCluster configuration running ONTAP 9.4 or later

The following table shows ISL port usage for the Brocade switches.

Switch model	ISL port	Switch port
Lenovo B6505 FC SAN switch	ISL port 1	20
	ISL port 2	21
	ISL port 3	22
	ISL port 4	23
Lenovo B6510 FC SAN switch	ISL port 1	40
	ISL port 2	41
	ISL port 3	42
	ISL port 4	43

Switch model	ISL port	Switch port
	ISL port 5	44
	ISL port 6	45
	ISL port 7	46
	ISL port 8	47
Brocade 7810	ISL port 1	ge2 (10-Gbps)
	ISL port 2	ge3 (10-Gbps)
	ISL port 3	ge4 (10-Gbps)
	ISL port 4	ge5 (10-Gbps)
	ISL port 5	ge6 (10-Gbps)
	ISL port 6	ge7 (10-Gbps)
Brocade 7840 **Note:** The Brocade 7840 switch supports either two 40 Gbps VE-ports or up to four 10 Gbps VE-ports per switch for the creation of FCIP ISLs.	ISL port 1	ge0 (40-Gbps) or ge2 (10-Gbps)
	ISL port 2	ge1 (40-Gbps) or ge3 (10-Gbps)
	ISL port 3	ge10 (10-Gbps)
	ISL port 4	ge11 (10-Gbps)
Brocade G610	ISL port 1	20
	ISL port 2	21
	ISL port 3	22
	ISL port 4	23
Lenovo ThinkSystem DB620S 32G FC SAN switch and Lenovo ThinkSystem DB630S 32G FC SAN switch, DB620S-1, G630-1, DB720	ISL port 1	40
	ISL port 2	41
	ISL port 3	42
	ISL port 4	43
	ISL port 5	44
	ISL port 6	45
	ISL port 7	46
	ISL port 8	47

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

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

Cisco 9396S			
Component	**Port**	**Switch 1**	**Switch 2**
controller_x_3	FC-VI port a	49	
	FC-VI port b	-	49

Cisco 9396S			
Component	**Port**	**Switch 1**	**Switch 2**
	FC-VI port c	50	
	FC-VI port d	-	50
	HBA port a	51	
	HBA port b	-	51
	HBA port c	52	
	HBA port d	-	52
controller_x_4	FC-VI port a	53	-
	FC-VI port b	-	53
	FC-VI port c	54	-
	FC-VI port d	-	54
	HBA port a	55	-
	HBA port b	-	55
	HBA port c	56	-
	HBA port d	-	56

Cisco 9148S			
Component	**Port**	**Switch 1**	**Switch 2**
controller_x_3	FC-VI port a	25	
	FC-VI port b	-	25
	FC-VI port c	26	-
	FC-VI port d	-	26
	HBA port a	27	-
	HBA port b	-	27
	HBA port c	28	-
	HBA port d	-	28
controller_x_4	FC-VI port a	29	-
	FC-VI port b	-	29
	FC-VI port c	30	-
	FC-VI port d	-	30
	HBA port a	31	-
	HBA port b	-	31
	HBA port c	32	-
	HBA port d	-	32

Cisco 9132T			
MDS module 1			
Component	**Port**	**Switch 1**	**Switch 2**
MDS module 2			
Component	Port	Switch 1	Switch 2
controller_x_3	FC-VI port a	1	-
	FC-VI port b	-	1
	FC-VI port c	2	-
	FC-VI port d	-	2
	HBA port a	3	-
	HBA port b	-	3
	HBA port c	4	-
	HBA port d	-	4
controller_x_4	FC-VI port a	5	-
	FC-VI port b	-	5
	FC-VI port c	6	-
	FC-VI port d	-	6
	HBA port a	7	-
	HBA port b	-	7
	HBA port c	8	-
	HBA port d	-	8

Cisco 9148 or 9250i*			
Component	**Port**	**Switch 1**	**Switch 2**
controller_x_3	FC-VI port a	7	-
	FC-VI port b	-	7
	HBA port a	8	-
	HBA port b	-	8
	HBA port c	9	-
	HBA port d	-	9
controller_x_4	FC-VI port a	10	-
	FC-VI port b	-	10
	HBA port a	11	-
	HBA port b	-	11
	HBA port c	13	-
	HBA port d	-	13

* - The Cisco 9250i switch is not supported for eight-node MetroCluster configurations.

Cisco port usage for FC-to-SAS bridges in a MetroCluster configuration running ONTAP 9.4 or later

Cisco 9396S			
FibreBridge 7600 using two FC ports	**Port**	**Switch 1**	**Switch 2**
bridge_x_1a	FC1	9	-
	FC2	-	9
bridge_x_1b	FC1	10	-
	FC2	-	10
bridge_x_2a	FC1	11	-
	FC2	-	11
bridge_x_2b	FC1	12	-
	FC2	-	12
bridge_x_3a	FC1	13	-
	FC2	-	13
bridge_x_3b	FC1	14	-
	FC2	-	14
bridge_x_4a	FC1	15	-
	FC2	-	15
bridge_x_4b	FC1	16	-
	FC2	-	16
		Additional bridges can be attached using ports 17 through 40 and 57 through 88 following the same pattern.	

Cisco 9148S			
FibreBridge 7600 using two FC ports	**Port**	**Switch 1**	**Switch 2**
bridge_x_1a	FC1	9	-
	FC2	-	9
bridge_x_1b	FC1	10	-
	FC2	-	10
bridge_x_2a	FC1	11	-
	FC2	-	11
bridge_x_2b	FC1	12	-
	FC2	-	12
bridge_x_3a	FC1	13	-
	FC2	-	13
bridge_x_3b	FC1	14	-
	FC2	-	14
bridge_x_4a	FC1	15	-

Cisco 9148S			
FibreBridge 7600 using two FC ports	Port	Switch 1	Switch 2
	FC2	-	15
bridge_x_4b	FC1	16	-
	FC2	-	16
		Additional bridges for a second DR group or second MetroCluster configuration can be attached using ports 33 through 40 following the same pattern.	

Cisco 9132T			
FibreBridge 7600 using two FC ports	Port	Switch 1	Switch 2
bridge_x_1a	FC1	9	-
	FC2	-	9
bridge_x_1b	FC1	10	-
	FC2	-	10
bridge_x_2a	FC1	11	-
	FC2	-	11
bridge_x_2b	FC1	12	-
	FC2	-	12
		Additional bridges for a second DR group or second MetroCluster configuration can be attached using the same port numbers on the second MDS module.	

Cisco 9148 or 9250i			
FibreBridge 7600 using two FC ports	Port	Switch 1	Switch 2
bridge_x_1a	FC1	14	-
	FC2	-	14
bridge_x_1b	FC1	15	-
	FC2	-	15
bridge_x_2a	FC1	17	-
	FC2	-	17
bridge_x_2b	FC1	18	-
	FC2	-	18
bridge_x_3a	FC1	19	-
	FC2	-	19
bridge_x_3b	FC1	21	-
	FC2	-	21
bridge_x_4a	FC1	22	-

Cisco 9148 or 9250i			
FibreBridge 7600 using two FC ports	Port	Switch 1	Switch 2
	FC2	-	22
bridge_x_4b	FC1	23	-
	FC2	-	23
		Additional bridges for a second DR group or second MetroCluster configuration can be attached using ports 25 through 48 following the same pattern.	

The following tables show bridge port usage when using FibreBridge 6500 bridges or FibreBridge 7600 bridges using one FC port (FC1 or FC2) only. For FibreBridge 7600 bridges using one FC port, either FC1 or FC2 can be cabled to the port indicated as FC1. Additional bridges can be attached using ports 25-48.

FibreBridge 6500 bridge or FibreBridge 7600 using one FC port	Port	Cisco 9396S	
		Switch 1	Switch 2
bridge_x_1a	FC1	9	-
bridge_x_1b	FC1	-	9
bridge_x_2a	FC1	10	-
bridge_x_2b	FC1	-	10
bridge_x_3a	FC1	11	-
bridge_x_3b	FC1	-	11
bridge_x_4a	FC1	12	-
bridge_x_4b	FC1	-	12
bridge_x_5a	FC1	13	-
bridge_x_5b	FC1	-	13
bridge_x_6a	FC1	14	-
bridge_x_6b	FC1	-	14
bridge_x_7a	FC1	15	-
bridge_x_7b	FC1	-	15
bridge_x_8a	FC1	16	-
bridge_x_8b	FC1	-	16
		Additional bridges can be attached using ports 17 through 40 and 57 through 88 following the same pattern.	

FibreBridge 6500 bridge or FibreBridge 7600 using one FC port	Port	Cisco 9148S	
		Switch 1	Switch 2
bridge_x_1a	FC1	9	-
bridge_x_1b	FC1	-	9

FibreBridge 6500 bridge or FibreBridge 7600 using one FC port	Port	Cisco 9148S	
		Switch 1	Switch 2
bridge_x_2a	FC1	10	-
bridge_x_2b	FC1	-	10
bridge_x_3a	FC1	11	-
bridge_x_3b	FC1	-	11
bridge_x_4a	FC1	12	-
bridge_x_4b	FC1	-	12
bridge_x_5a	FC1	13	-
bridge_x_5b	FC1	-	13
bridge_x_6a	FC1	14	-
bridge_x_6b	FC1	-	14
bridge_x_7a	FC1	15	-
bridge_x_7b	FC1	-	15
bridge_x_8a	FC1	16	-
bridge_x_8b	FC1	-	16
		Additional bridges for a second DR group or second MetroCluster configuration can be attached using ports 25 through 48 following the same pattern.	

Cisco 9148 or 9250i			
FibreBridge 6500 bridge or FibreBridge 7600 using one FC port	Port	Switch 1	Switch 2
bridge_x_1a	FC1	14	-
bridge_x_1b	FC1	-	14
bridge_x_2a	FC1	15	-
bridge_x_2b	FC1	-	15
bridge_x_3a	FC1	17	-
bridge_x_3b	FC1	-	17
bridge_x_4a	FC1	18	-
bridge_x_4b	FC1	-	18
bridge_x_5a	FC1	19	-
bridge_x_5b	FC1	-	19
bridge_x_6a	FC1	21	-
bridge_x_6b	FC1	-	21
bridge_x_7a	FC1	22	-
bridge_x_7b	FC1	-	22
bridge_x_8a	FC1	23	-

Cisco 9148 or 9250i			
FibreBridge 6500 bridge or FibreBridge 7600 using one FC port	**Port**	**Switch 1**	**Switch 2**
bridge_x_8b	FC1	-	23
		Additional bridges can be attached using ports 25 through 48 following the same pattern.	

Cisco port usage for ISLs in an eight-node configuration in a MetroCluster configuration running ONTAP 9.4 or later

The following table shows ISL port usage. ISL port usage is the same on all switches in the configuration.

Switch model	ISL port	Switch port
Cisco 9396S	ISL 1	44
	ISL 2	48
	ISL 3	92
	ISL 4	96
Cisco 9148 with 24 port license or 9250i	ISL 1	12
	ISL 2	16
	ISL 3	20
	ISL 4	24
Cisco 9148, 9148S	ISL 1	20
	ISL 2	24
	ISL 3	44
	ISL 4	48
Cisco 9132T	ISL 1	MDS module 1 port 13
	ISL 2	MDS module 1 port 14
	ISL 3	MDS module 1 port 15
	ISL 4	MDS module 1 port 16

Chapter 20. Where to find additional information

You can learn more about configuring, operating, and monitoring a MetroCluster configuration in Lenovo's extensive documentation library.

MetroCluster and miscellaneous guides

Guide	Content
ONTAP 9 Documentation Center	• All MetroCluster guides
Fabric-attached MetroCluster installation and configuration	• Fabric-attached MetroCluster architecture • Cabling the configuration • Configuring the FC-to-SAS bridges • Configuring the FC switches • Configuring the MetroCluster in ONTAP
MetroCluster IP installation and configuration	• MetroCluster IP architecture • Cabling the MetroCluster IP configuration • Configuring the MetroCluster in ONTAP
Unified Manager for DM Series documentation Unified Manager for DM Series	• Monitoring the MetroCluster configuration and performance
Tiebreaker software installation and configuration	• Monitoring the MetroCluster configuration with the MetroCluster Tiebreaker software

© Copyright Lenovo 2019, 2023

Appendix A. Contacting Support

You can contact Support to obtain help for your issue.

You can receive hardware service through a Lenovo Authorized Service Provider. To locate a service provider authorized by Lenovo to provide warranty service, go to https://datacentersupport.lenovo.com/serviceprovider and use filter searching for different countries. For Lenovo support telephone numbers, see https://datacentersupport.lenovo.com/supportphonelist for your region support details.

© Copyright Lenovo 2019, 2023

Appendix B. Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area.

Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service.

Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document is not an offer and does not provide a license under any patents or patent applications. You can send inquiries in writing to the following:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

© Copyright Lenovo 2019, 2023

Trademarks

LENOVO, LENOVO logo, and THINKSYSTEM are trademarks of Lenovo. All other trademarks are the property of their respective owners. © 2023 Lenovo.