IBM System Storage N series

IBM

Fabric-attached MetroCluster Systems: Brocade Switch Configuration Guide

Contents

Preface

Supported features

IBM System Storage N series storage systems are driven by NetApp Data ONTAP software. Some features described in the product software documentation are neither offered nor supported by IBM. Please contact your local IBM representative or reseller for further details.

Information about supported features can also be found on the N series support website (accessed and navigated as described in *Websites* on page 5).

Websites

IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. The following web pages provide N series information:

- A listing of currently available N series products and features can be found at the following web page:

  www.ibm.com/storage/nas/
- The IBM System Storage N series support website requires users to register in order to obtain access to N series support content on the web. To understand how the N series support web content is organized and navigated, and to access the N series support website, refer to the following publicly accessible web page:

  www.ibm.com/storage/support/nseries/

  This web page also provides links to AutoSupport information as well as other important N series product resources.
- IBM System Storage N series products attach to a variety of servers and operating systems. To determine the latest supported attachments, go to the IBM N series interoperability matrix at the following web page:

  www.ibm.com/systems/storage/network/interophome.html
- For the latest N series hardware product documentation, including planning, installation and setup, and hardware monitoring, service and diagnostics, see the IBM N series Information Center at the following web page:

  publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

Getting information, help, and service

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This section contains

information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your IBM N series product, and whom to call for service, if it is necessary.

Before you call

Before you call, make sure you have taken these steps to try to solve the problem yourself:

- Check all cables to make sure they are connected.
- Check the power switches to make sure the system is turned on.
- Use the troubleshooting information in your system documentation and use the diagnostic tools that come with your system.
- Refer to the N series support website (accessed and navigated as described in *Websites* on page 5) for information on known problems and limitations.

Using the documentation

The latest versions of N series software documentation, including Data ONTAP and other software products, are available on the N series support website (accessed and navigated as described in *Websites* on page 5).

Current N series hardware product documentation is shipped with your hardware product in printed documents or as PDF files on a documentation CD. For the latest N series hardware product documentation PDFs, go to the N series support website.

Hardware documentation, including planning, installation and setup, and hardware monitoring, service, and diagnostics, is also provided in an IBM N series Information Center at the following web page:

publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

Hardware service and support

You can receive hardware service through IBM Integrated Technology Services. Visit the following web page for support telephone numbers:

www.ibm.com/planetwide/

Firmware updates

IBM N series product firmware is embedded in Data ONTAP. As with all devices, it is recommended that you run the latest level of firmware. Any firmware updates are posted to the N series support website (accessed and navigated as described in *Websites* on page 5).

Note: If you do not see new firmware updates on the N series support website, you are running the latest level of firmware.

Verify that the latest level of firmware is installed on your machine before contacting IBM for technical support.

How to send your comments

Your feedback helps us to provide the most accurate and high-quality information. If you have comments or suggestions for improving this document, please send them by e-mail to *starpubs@us.ibm.com*.

Be sure to include the following:

• Exact publication title
• Publication form number (for example, GC26-1234-02)
• Page, table, or illustration numbers
• A detailed description of any information that should be changed

Supported Brocade switches

To configure your fabric-attached MetroCluster configuration, you must use one of the supported Brocade switches running Brocade Fabric Operating System 6.0.x or later.

The following Brocade switches are supported:

- Brocade 200E (IBM 2005-B16)
- Brocade 5000 (IBM 2005-B5K)
- Brocade 300 (IBM SAN24B-4)
- Brocade 5100 (IBM SAN40B-4)

How ports are configured for use on these switches depends on how disk ownership is defined by Data ONTAP.
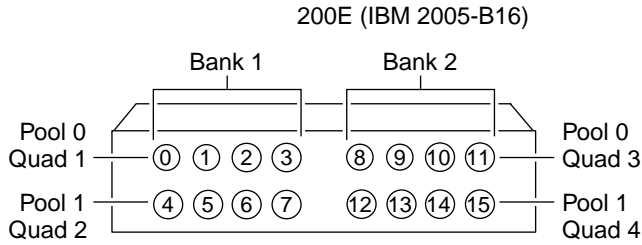
How switch ports are defined or grouped

How you define or group your switches depends on whether they support hardware-based or software-based disk ownership, as defined by Data ONTAP. All the currently supported switches support software-based disk ownership; the Brocade 200E (IBM 2005-B16) and Brocade 5000 (IBM 2005-B5K) switches also support hardware-based disk ownership.

Hardware-based disk ownership	The ports in Fibre Channel switches that support hardware-based disk ownership are divided into two *banks* that divide the switch into equal parts. The banks are further divided into *quadrants* of 2, 4, or 8 ports each (depending on the number of ports in the switch). The quadrants are combined to form two *pools*, which function identically in a fabric-attached MetroCluster configuration. Quadrants 1 and 3 belong to Pool 0; quadrants 2 and 4 belong to Pool 1.
Software-based disk ownership	The ports in Fibre Channel switches that support only software-based disk ownership are divided into port groupings of eight. The number of port groups available depends on the total number of ports you enabled with your Ports on Demand (POD) license.

Number of ports in the Brocade 200E (IBM 2005-B16) switch

The Brocade 200E (IBM 2005-B16) switch has 16 ports, but can also be licensed as an 8-port switch.
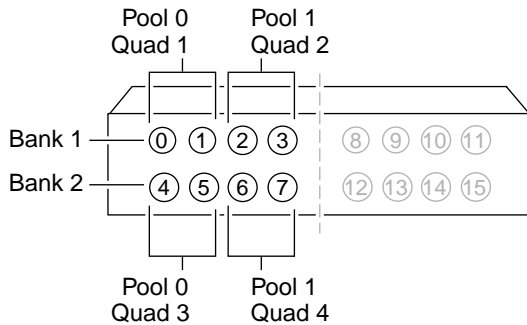
When licensed for 16 ports, the switch has ports 0 through 15 licensed and operational, as shown in the following diagram. The pool and quad assignments shown are for hardware ownership.

200E (IBM 2005-B16)

Bank 1   Bank 2

Pool 0
Quad 1 — ⓪ ① ② ③   ⑧ ⑨ ⑩ ⑪ — Pool 0
                                  Quad 3

Pool 1 — ④ ⑤ ⑥ ⑦   ⑫ ⑬ ⑭ ⑮ — Pool 1
Quad 2   Quad 4

When licensed for 8 ports, the Brocade 200E (IBM 2005-B16) switch has ports 0 through 7 licensed and operational, as shown in the following diagram.

> **Attention:** You can only install Small Form Pluggables (SFPs) in licensed ports. Installing SFPs in unlicensed ports causes Data ONTAP to erroneously assign disks to the unlicensed ports.
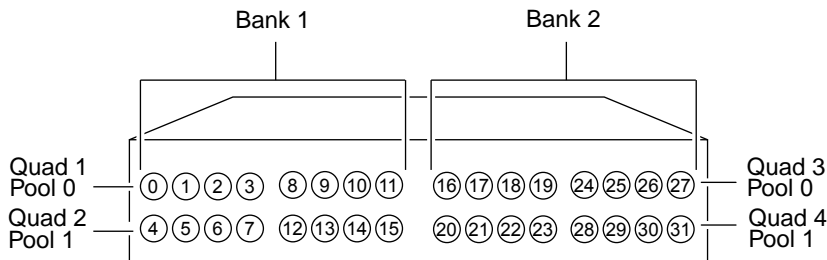
200E (IBM 2005-B16) in 8-port configuration

Pool 0   Pool 1
Quad 1   Quad 2

Bank 1 — ⓪ ① ② ③   ⑧ ⑨ ⑩ ⑪
Bank 2 — ④ ⑤ ⑥ ⑦   ⑫ ⑬ ⑭ ⑮

Pool 0   Pool 1
Quad 3   Quad 4

Number of ports in the Brocade 5000 (IBM 2005-B5K) switch

The Brocade 5000 (IBM 2005-B5K) switch has 32 licensed and operational ports.

The 32 ports on the switch are licensed and operational as ports 0 through 31, as shown in the following diagram.

Bank 1   Bank 2

Quad 1
Pool 0 — ⓪①②③ ⑧⑨⑩⑪   ⑯⑰⑱⑲ ㉔㉕㉖㉗ — Quad 3
                                          Pool 0
Quad 2
Pool 1 — ④⑤⑥⑦ ⑫⑬⑭⑮   ⑳㉑㉒㉓ ㉘㉙㉚㉛ — Quad 4
   Pool 1

Number of ports in the Brocade 300 (IBM SAN24B-4) switch

The 24-port Brocade 300 (IBM SAN24B-4) switch can be licensed as an 8-port, 16-port, or 24-port switch.

In the following diagram, ports 0 through 7 are enabled by default; ports 8 through 15, and ports 16 through 23 can be enabled through a ports on demand (POD) license.

Attention: You can only install Small Form Pluggables (SFPs) in licensed ports. Installing SFPs in unlicensed ports causes Data ONTAP to erroneously assign disks to the unlicensed ports.
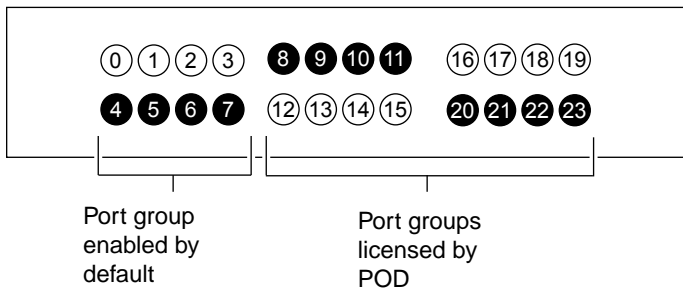

Number of ports in the Brocade 5100 (IBM SAN40B-4) switch

The 40-port Brocade 5100 (IBM SAN40B-4) switch can be licensed as a 24-port, 32-port, or 40-port switch.

In the following diagram, Ports 0 through 23 are enabled by default; Ports 24 through 31 and ports 32 through 39 can be enabled through a POD license.

Attention: You can only install Small Form Pluggables (SFPs) in licensed ports. Installing SFPs in unlicensed ports causes Data ONTAP to erroneously assign disks to the unlicensed ports.
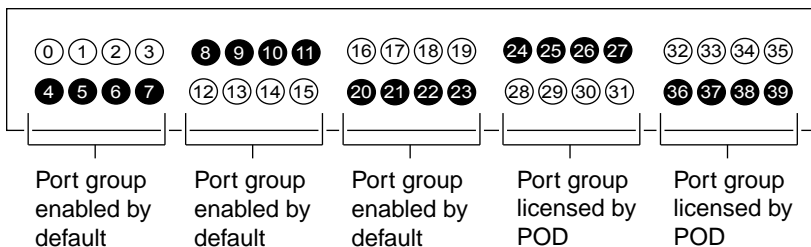
Configuring Brocade switches for a fabric-attached MetroCluster configuration

You must configure your Brocade switches to prepare them for your fabric-attached MetroCluster configuration. This set of procedures must be performed on each switch in the MetroCluster configuration.

About this task

If your Brocade switch is the Brocade 5100 (IBM SAN40B-4), Brocade 5000 (IBM 2005-B5K), Brocade 200E (IBM 2005-B16) or Brocade 300 (IBM SAN24B-4), you need to keep the following points in mind when configuring your switch.

Note: The ports on all of these listed switches automatically sense link speeds and support transfer rates of 1, 2, 4, and 8 Gbits per second. The 8 Gbit SFPs support transfer rates of 2, 4, or 8 Gbits per second; 4 Gbit SFPs support transfer rates of 1, 2, or 4 Gbits per second.

Attention: Do not use the auto-sensing speed mode with switch ports that are connected to FC-VI adapter ports, FC initiator adapter ports, E-ports, or storage loops. You can use the `portcfgspeed` command to match the speed of these switch ports to the speed supported by the port at the other end of the link.

Configuring a switch initially

You must do an initial configuration of the switches to prepare them for your fabric-attached MetroCluster configuration. These procedures must be performed on each switch in the MetroCluster configuration. To configure a switch initially, complete the following steps:

Steps

1. If you have not configured the switch before, log in to the switch using the console port and your Windows HyperTerminal application (not available in Windows Vista versions).

2. Set the following values, as described in the Brocade switch hardware reference manual for your switch:

   • IP address
   • Default account passwords (administrator and user)
   • Date and time

   Note: Use the `telnet` command for subsequent switch login sessions, using the switch IP address.

3. Display the licenses installed on the switch by entering the following command:

```
licenseshow
```

> **Note:** Use the `telnet` command for subsequent switch login sessions, using the switch IP address.

You should have the following licenses:

- Fabric license (or two Domain licenses)
- Extended Fabric license (for ISL distances over 10 km)

If you do not have these licenses, contact your sales representative before proceeding.

4. Set the switch name, if needed, by entering the following command:

```
switchname switch_name
```

Use a name that is meaningful, such as *node name* and *switch*.

5. Determine the switch firmware by entering the following command:

```
version
```

The currently installed switch firmware is displayed.

6. Access the *MetroCluster Compatibility Matrix* from the IBM N series interoperability matrix at the following web page: www.ibm.com/systems/storage/network/interophome.html

7. If your switch firmware is not the supported version, complete the following substeps:

   a. Download the switch firmware from the IBM N series support site and install it, as described in the switch hardware reference manual for your switch.

   b. Reboot the switch.

Setting the switch parameters

You must customize the switch parameters to the domain of your fabric-attached MetroCluster configuration. These procedures must be performed on each switch in the MetroCluster configuration.

Steps

1. Disable the switch by entering the following command:

```
switchdisable
```

2. Clear any preexisting configuration by entering the following commands in sequence:

```
cfgclear
```

```
cfgdisable
```

cfgsave

3. Configure the switch with default settings by entering the following command:

   configdefault

4. Enter the following command:

   configure

5. Set the following switch parameters:

   a. Set `Fabric parameters` as follows:

      y

      Note: For For N series gateway systems deployed in a MetroCluster configuration, set this parameter to `N`.

      Additional parameters appear. Only define the following parameters:

      - Set `Domain_id = XXX`.
      - Set `Disable device probing = 1`.

      Note: `Domain_id` might be specified by your system administrator. It is a three-digit number unique to the switches in the fabric. If not, you can use any unique number, for example, 1, 2, 3, or 4.

   b. Set `Arbitrated Loop parameters` as follows:

      y

      Additional parameters appear. Define the following parameter:

      - Set `Send FAN frames = 0`.

6. Exit the configuration utility by entering the following command:

   ctrl-d

7. Reboot the switch so that the new settings take effect by entering the following command:

   fastboot

Completing the switch configuration

You must complete the switch configuration by considering the Fabric OS version and the licenses on it and ISL lengths of your fabric-attached MetroCluster configuration. These procedures must be performed on each switch in the MetroCluster configuration.

Steps

1. Log in to the switch and disable it by entering the following command:

switchdisable

2. If the switches are running Fabric OS 6.3.1c or later and the switch is of type 300 (SAN24B-4) or 5100 (SAN40B-4), set the FillWord on all ports to mode *3* by entering the following command:

   portCfgFillWord #port 3

3. Depending on the Fabric OS version running on your switch, complete one of the following procedures:

If the switches are running...	Then...

Fabric OS 6.1.x and later	**a.** Set all ports attached to disk shelf loops to Locked L-Port and Locked Loop Fairness by entering the following command for each port number:
	
	portCfgLPort *disk_port*,1,0,2
	
	    Attention: If you are using the FibreBridge 6500N for SAS disk shelves, you must not perform this step.
	
	b. For ports connected to FC-VI adapters, FC initiator adapters, E-ports, or disk shelf loops, set the switch port speed to that of the adapter, remote E-port, or disk shelf by entering the following command:
	
	portcfgspeed port_number, speed
	
	c. Ensure that the disk shelf loop port is showing ON in both the Locked L_Port and Locked Loop Fairness fields by entering the following commands in sequence:
	
	portCfgShow
	
	portCfglport
	
	    Attention: If you are using the FibreBridge 6500N for SAS disk shelves, you must not perform the portCfglport command.
	
	ON is displayed in both the Locked Loop and Locked Loop Fairness fields.

Fabric OS 6.0.x and earlier	**a.** Set all ports attached to disk shelf loops to half duplex by entering the following command for each port number: **portCfgLPort *disk_port*,0,0,1** **b.** Set all ports attached to disk shelf loops to Locked L-Port by entering the following command for each port number: **portCfgLPort *disk_port*,1,0,1** **c.** Ensure that the disk shelf loop port is showing ON in both the Locked L_Port and Locked Loop HD fields by entering the following commands in sequence: **portCfgShow** **portCfglport** ON is displayed in both the Locked L_Port and Locked Loop HD fields.

4. If your switch has the trunking license installed, disable trunking on the ISL port by entering the following command:

   portcfgtrunkport *ISL_port#* 0

5. Configure the long-distance ISL port for your ISL length by entering the following command:

   portcfglongdistance *portnumber distance_level VC_Translation_Link_Init desired_distance*

   distance_level should be one of the following values:

   - For distances up to 10 km, use LE
   - For distances beyond 10 km, use LS or LD

   desired_distance is recommended to be 1.5 times the actual distance of the link. Specifying 1 activates long-distance link initialization sequence, and specifying 0 deactivates this mode.

   Example
   If your ISL length is 18 km and your ISL is connected to port 13, you use the following command:

   portCfgLongDistance 13 LS 1 27

6. Enable the switch by entering the following command:

   switchenable

7. Verify that the switch settings are correct by entering the following command:

   configshow

Setting the switch port for FC/SAS bridges (ATTO for IBM FibreBridge 6500N)

The FibreBridge 6500N is Feature Code (FC) 2100 and is an FC to SAS protocol converter that is designed to enable the use of the SAS disk shelves, such as the EXN3000 and EXN3500 Storage Expansion Units, in MetroCluster installations.

You must configure the Brocade switches to prepare them to use FibreBridge 6500N bridges.

About this task

- You must perform the steps on each switch in the MetroCluster configuration.
- For fabric-attached MetroCluster configurations, each FC switch must have one FC port available for one bridge to connect to it.
- The FC port data rate must be set to 8 Gb.
- The FC port connection mode must be set to *ptp*.

Steps

1. Set all the Brocade FC switch ports connected to FibreBridge 6500N to G_Port by entering the following command:

    `portcfgGport PortNumber 1`

    > **Note:** You can use the `portcfgGport PortNumber 0` command to release the G_Port lockdown.

    Example
    To set port 1 to G_Port, enter the following command:

    `portcfggport 1 1`

2. Verify the setting by entering the following command:

    `portcfgshow`

    You should see that "Locked G_Port" is set to *ON* for all ports connected to FibreBridge 6500N bridges.

Brocade's zoning feature

Brocade switches ship with the zoning license enabled by default. Two zoning options are available on switches running Fabric OS 6.0.0b or later: the basic zone option and traffic isolation (TI) zone option.

Brocade's basic zoning feature
Brocade's basic zoning feature enables you to arrange fabric-connected storage devices into two or more logical zones on the same switch. Zone members can only access other members of the same zone. For example, you can isolate systems with different uses or protect systems in a heterogeneous environment.

> **Note:** If you are configuring a MetroCluster configuration with more than 336 disks, you must implement zoning.

Brocade's Traffic Isolation (TI) zoning feature
Brocade's TI zoning feature is available for the Brocade switches and the equivalent IBM switches listed in this document that are connected to storage systems running Data ONTAP 7.2.5.1 or later. It enables you to control the flow of Inter-Switch links (ISLs) traffic by creating a dedicated path for traffic flowing from a specific set of source ports (N_Ports). The TI feature helps to separate Fibre Channel Virtual Interface (FCVI) messages, which are considered high-priority traffic, to avoid interruption or congestion caused by storage traffic.

Each TI zone is interpreted by each switch, and each switch considers only the routing required for its local ports. No consideration is given to the overall topology or to whether the TI zones accurately provide dedicated paths through the whole fabric. When you activate a TI zone, the fabric attempts to isolate all ISL traffic entering from a member of the zone to only those ISLs that are included in the zone. The fabric also attempts to exclude traffic not in the TI zone from using ISLs within that TI zone.

You use the TI feature in the following scenarios:

- When you want to dedicate an ISL to high-priority, host-to-target traffic such as cluster interconnect traffic.
- When you want to force high-volume, low-priority traffic onto a particular ISL to limit the effect of this high traffic pattern on the fabric.

> **Note:** Although the TI feature provides better resiliency and performance, it requires twice the number of ISLs and therefore twice the number of fibers between locations. If this overhead is possible, then it is recommended that the TI feature be used.

Rules for configuring zones

You need to know the rules for using Brocade's zoning feature to create a zone for a Fibre Channel Virtual Interface (FCVI) from a local node to a remote node, create a second zone for all FC initiator ports, and if necessary, create a third zone when implementing Traffic Isolation (TI).

Unless you are implementing traffic isolation, you need to create two zones on each switch: a zone for an FCVI from a local node to a remote node, and a second zone containing all FC initiator ports and all connected storage shelves. If you are implementing TI, you must create the TI zone before creating the basic zones.

The following rules apply to the software you are running:

- If you want to upgrade your switches to Brocade Fabric OS 6.3.1c or later, you must first implement single-loop storage zones.
- If your Brocade switches are running Fabric OS 6.1.1a or earlier, you must first upgrade them to Brocade Fabric OS 6.2.*x* before upgrading to Fabric OS 6.3.1c or later.

To create the basic zones, the following rules must be observed:

- You must configure the zones using Domain IDs and port numbers; you cannot use the PWWN.
- You must configure these zones on a primary and a secondary fabric.
- You must configure the primary and the secondary fabrics separately.
- You cannot include the ISL port in a zone.
  You can use any other port in any of the four quadrants. For more information, see the *Data ONTAP 7-Mode High-Availability Configuration Guide* for Data ONTAP 7.2.5.1 or later.

To create the TI zone, the following rules must be observed:

- If you have two ISLs, you must create a TI zone.
- You use the `zone` command with options to create and modify TI zones.

  Note: You cannot use other zoning commands, such as `zoneCreate`, `aliCreate`, and `cfgCreate` to manage TI zones.

- You use the `cfgenable` command to commit all TI zone commands.
- A TI zone must include a set (two or more) of E_Ports forming an end-to-end path.
  Inclusion of N_Ports is optional.
- A given N_Port can be a member of only a single TI zone.
  This rule is enforced during zone creation or modification.
- An E_Port can be a member of only a single TI zone.
  This rule is enforced during zone creation or modification.
- If multiple E_Ports on the lowest-cost route to a domain are configured, the various source ports for that zone are load balanced across the specified E_Ports.
- The TI zones appear in the defined zone configuration only and do not appear in the effective zone configuration.

A TI zone only provides Traffic Isolation and is not a "regular" zone.

Configuring zones on the primary fabric

You can configure zones on the primary fabric to control access to fabric segments and to establish barriers between operating environments.

Steps

1. Connect to a switch on the primary fabric by entering the following command:

```
telnet
```

Example
The following example uses factory defaults:

```
telnet 10.77.77.77
```

2. If you are implementing traffic isolation (TI), complete the following steps:

   a. Create the FCVI TI zone on the primary fabric by entering the following command:

   ```
   zone --create -t objtype -o optionlist zonename -p portlist
   ```

   Example
   The following example creates a TI zone called *FCVI_TI_1* with failover disabled and the TI zone activated. The example includes switch 1, port 0 and port 4; and switch 3, port 0 and port 4. The switch1, port 0 and switch 3, port 0 are the FC-VI adapter ports. Switch 1, port 4 and switch 3, port 4 are the ports of the first ISL between the two switches.

   ```
   zone --create -t ti FVCI_TI_1 -o n -p 1,0; 1,4; 3,0; 3,4
   ```

   b. Activate the TI zone by entering the following command:

   ```
   zone --activate zonename
   ```

   Example
   The following example activates the TI zone called *FCVI_TI_1*:

   ```
   zone --activate FVCI_TI_1
   ```

3. Create a zone for the FCVI traffic by entering the following command:

```
zonecreate "zonename","portlist"
```

You need to specify the name you are assigning the zone in `zonename` and specify the switch ID and port number you want to include in this zone in `portlist`.

Example
The following example assumes that software disk ownership is used and creates a zone called *FCVI_1* that includes switch 1, port 0 on the local node, and switch 3, port 0 on the remote node:

```
zonecreate "FCVI_1", "1,0; 3,0"
```

Note: This zone does not include the ISL ports.

4. Create the storage zone on the primary fabric.

If you are creating a...	Then...
Single storage zone	Enter the following command:
	`zonecreate "zonename", "portlist"`
	The following example creates a zone, *STOR_1*, that includes all storage initiators and disk shelf ports. Switch 1, port 5 and switch 3, port 5 are the ports of the second ISL between these two switches.
	`zonecreate "STOR_1","1,1; 1,2; 1,3; 1,6; 1,7; 1,8; 1,9; 1,10; 1,11; 1,12; 1,13; 1,14; 1,15; 1,16; 1,17; 1,18; 1,19; 1,20; 1,21; 1,22; 1,23; 3,1; 3,2; 3,3; 3,6; 3,7; 3,8; 3,9; 3,10; 3,11; 3,12; 3,13; 3,14; 3,15; 3,16; 3,17; 3,18; 3,19; 3,20; 3,21; 3,22; 3,23"`
Single-loop storage zone (for Fabric OS 6.3.1c and later)	Enter the following command for each disk loop port that includes the port and the ports containing the FC controller connections:
	`zonecreate "zonename", "portlist"`
	The following example creates the zones STOR_Switch#_Port# and includes a `zonecreate` command for each port that contains a disk loop. The first four ports of each command are the FC controller connections that are followed by the disk loop port.
	`zonecreate "STOR_1_6","1,1;1,2;3,1;3,2;1,6"`

Note: The only ports not listed in the example are the ISL ports. Even if you have insufficient storage to connect to all the ports on the switch, it is best practice to configure the ports now to avoid modifying the zones later when additional storage is added.

5. Create a configuration on the primary fabric.

If you are creating...	Then...
Single storage zone	Enter the following command to create the configuration:
	`cfgcreate "configuration_name","zonename; zonename"`
	The following example creates a configuration *Config_1* that includes the zones *FCVI_1* and *STOR_1*. It does not contain the FCVI_TI zone.
	`cfgcreate "Config_1", "FCVI_1; STOR_1"`

If you are creating...	Then...

Single-loop storage zones (for Fabric OS 6.3.1c and later)	Enter the following command to create configuration:
	
	`cfgcreate "configuration_name","zonename; zonename"`
	
	The following example creates a configuration *Config_1* that includes the zones *FCVI_1* and all the single-loop storage zones. It does not contain the FCVI_TI zone.
	
	`cfgcreate "Config_1", "FCVI_1; STOR_1_6; STOR_1_7;`
	`STOR_1_8; STOR_1_9; STOR_1_10; STOR_1_11; STOR_1_12;`
	`STOR_1_13; STOR_1_14; STOR_1_15; STOR_1_16; STOR_1_17;`
	`STOR_1_18; STOR_1_19; STOR_1_20; STOR_1_21; STOR_1_22;`
	`STOR_1_23; STOR_3_6; STOR_3_7; STOR_3_8; STOR_3_9;`
	`STOR_3_10; STOR_3_11; STOR_3_12; STOR_3_13; STOR_3_14;`
	`STOR_3_15; STOR_3_16; STOR_3_17; STOR_3_18; STOR_3_19;`
	`STOR_3_20; STOR_3_21; STOR_3_22; STOR_3_23"`

6. Save the configuration by entering the following command:

   `cfgsave`

7. Enable the configuration by entering the following command:

   `cfgenable "Configuration_name"`

   Example
   The following example saves and enables the *Config_1* configuration:

   `cfgenable "Config_1"`

8. Verify the configuration by entering the following command:

   `zone --show`

9. Verify the FCVI and STOR zones by entering the following command:

   `cfgshow`

Configuring zones on the secondary fabric

Having configured zones on the primary fabric, you must configure zones on the secondary fabric to control access to fabric segments and to establish barriers between operating environments.

Steps

1. Connect to a switch on the secondary fabric by entering the following command:

   `telnet`

   Example
   The following example uses factory defaults:

   `telnet 10.77.77.77`

2. If you are implementing traffic isolation (TI), complete the following substeps; otherwise, go to step 3:

   a. Create the FCVI TI zone on the secondary fabric by entering the following command:

      ```
      zone --create -t objtype -o optionlist zonename -p portlist
      ```

      Example
      The following example creates a TI zone called *FCVI_TI_2* with failover disabled and the TI zone activated. Switch 2, port 0 and switch 4, port 0 are the FC-VI adapter ports. Switch 2, port 4 and switch 4, port 4 are the ports of the first ISL between these two switches. The example includes switch 2, port 0 and port 4; and switch 4, port 0 and port 4.

      ```
      zone --create -t ti FVCI_TI_2 -o n -p 2,0; 2,4; 4,0; 4,4
      ```

   b. Activate the TI zone by entering the following command:

      ```
      zone --activate zonename
      ```

      Example
      The following example activates the TI zone called *FCVI_TI_2*:

      ```
      zone --activate FVCI_TI_2
      ```

3. Create the FCVI zone on the secondary fabric by entering the following command:

   ```
   zonecreate "zonename","portlist"
   ```

   You must specify the name you are assigning the zone in `zonename` and specify the switch ID and port number you want to include in this zone in `portlist`.

   Example
   The following example creates a zone `FCVI_2`. This zone includes the two FCVI adapter ports switch 2, port 0 on the local node, and switch 4, port 0 on the remote node. This zone does not contain the ISL ports.

   ```
   zonecreate "FVCI_2","2,0; 4,0"
   ```

4. Create the storage zone on the secondary fabric.

If you are creating...	Then...

A single storage zone	Enter the following command:
	
	`zonecreate "zonename","portlist"`
	
	The following example creates a zone called *STOR_2*. This zone includes all storage initiator and disk shelf ports. Switch 2, port 5 and Switch 4, port 5 are the ports of the second ISL between these two switches. This zone does not contain the ISL ports.
	
	`zonecreate "STOR_2","2,1; 2,2; 2,3; 2;6; 2,7; 2,8; 2,9; 2,10; 2,11; 2,12; 2,13; 2,14; 2,15; 2,16; 2,17; 2,18; 2,19; 2,20; 2,21; 2,22; 2,23; 4,1; 4,2; 4,3; 4;6; 4,7; 4,8; 4,9; 4,10; 4,11; 4,12; 4,13; 4,14; 4,15; 4,16; 4,17; 4,18; 4,19; 4,20; 4,21; 4,22; 4,23"`

If you are creating...	Then...

Single-loop storage zones (for Fabric OS 6.3.1c and later)	Enter the following command for each disk loop port that includes the port and the ports containing the FC controller connections:
	
	`zonecreate "zonename", "portlist"`
	
	The following example creates the zone STOR_Switch#_Port# and includes a `zonecreate` command for each port that contains a disk loop. The first four ports of each command are the FC controller connections that are followed by the disk loop port.
	
	`zonecreate "STOR_2_6", "2,1; 2,2; 4,1; 4,2; 2,6"`

Note: The only ports not listed in this example are the ISL ports. Even if you have insufficient storage to connect to all the ports on the switch, it is a best practice to configure the ports now to avoid modifying the zones later when additional storage is added.

5. Create a configuration zone on the secondary fabric.

If you are creating...	Then...

A single storage zone	Enter the following command to create the configuration:
	
	`cfgcreate "configuration_name", "zonename; zonename"`
	
	`configuration_name` is the name that you provide for the configuration.
	
	`zonename` are the zones that you want to include in the configuration.
	
	The following example creates a configuration named `Config_2` that includes the zones named `FCVI_2` and `STOR_2`:
	
	`cfgcreate "Config_2","FCVI_2; STOR_2"`

Single-loop storage zones	Enter the following command to create the configuration:
	
	`cfgcreate "configuration_name","zonename; zonename"`
	
	The following example creates a configuration `Config_2` that includes the zones `FCVI_2` and all the single-loop storage zones. It does not contain the FCVI_TI zone.
	
	`cfgcreate "Config_2","FCVI_2; STOR_2_6; STOR_2_7; STOR_2_8;`
	`STOR_2_9; STOR_2_10; STOR_2_11; STOR_2_12; STOR_2_13;`
	`STOR_2_14; STOR_2_15; STOR_2_16; STOR_2_17; STOR_2_18;`
	`STOR_2_19; STOR_2_20; STOR_2_21; STOR_2_22; STOR_2_23;`
	`STOR_4_6; STOR_4_7; STOR_4_8; STOR_4_9; STOR_4_10;`
	`STOR_4_11; STOR_4_12; STOR_4_13; STOR_4_14; STOR_4_15;`
	`STOR_4_16; STOR_4_17; STOR_4_18; STOR_4_19; STOR_4_20;`
	`STOR_4_21; STOR_4_22; STOR_4_23"`

6. Save the configuration by entering the following command:

`cfgsave`

7. Enable the configuration by entering the following command:

```
cfgenable "Configuration_name"
```

Example

The following example saves and enables the *Config_2* configuration:

```
cfgenable "Config_2"
```

8. Verify the configuration by entering the following command:

```
zone --show
```

9. Verify the FCVI and STOR zones by entering the following command:

```
cfgshow
```

Configuring zones on the shared-switches configuration

You must configure zones on the switches to achieve the shared-switches configuration. Configuring the zones involves steps such as, creating TI zones, FCVI zones, and storage zones. The sample illustration explains configuring the zones at the switches.

About this task

The following illustration shows a shared-switches configuration where two fabric-attached MetroCluster pairs share the switches and the ISLs between the switches. FMC1-1 and FMC1-2 form fabric-attached MetroCluster FMC1, and FMC2-1 and FMC2-2 form fabric-attached MetroCluster FMC2. The switches are named S1, S2, S3, and S4 with domain IDs 1, 2, 3, and 4, respectively.
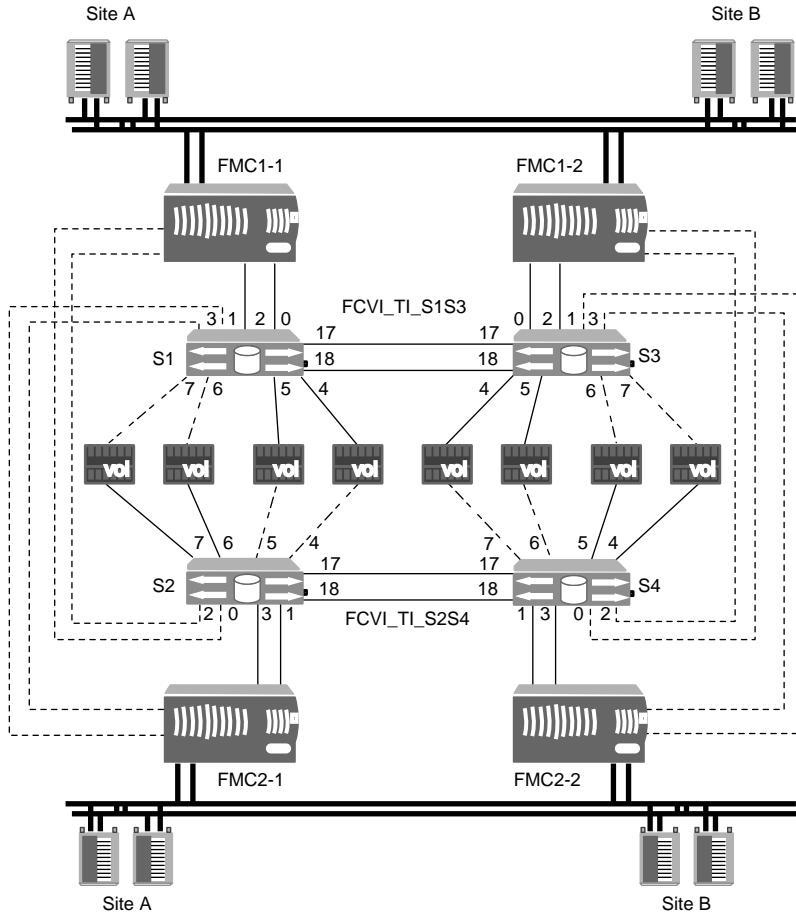
In this example, each storage system has two FCVI and HBA ports. Each port is connected to the primary and secondary switches.

- FMC1 connects the FCVI and HBA ports to switch S1 through port 0 and 2.
- FMC2 connects the FCVI and HBA ports to switch S2 through port 1 and 3.

The disk shelves are connected to the switch through ports 4, 5, 6, and 7. In addition, there are two ISLs on port 17 and 18 on all the switches.

Steps

1. Connect to a switch on the secondary fabric by entering the following command:

   ```
   telnet
   ```

   Example
   The following example uses factory defaults:

```
telnet 10.77.77.77
```

2. If you are implementing traffic isolation (TI), complete the following steps:

   a. Create the FCVI-TI zone by entering the following command:

   ```
   zone --create -t objtype -o optionlist zonename -p portlist
   ```

   Example
   The following example creates FCVI_TI zone *FCVI_TI_S1S3* on switch S1. The ports 0 and 1 are the FCVI zones on switch S1 and S3. Switch S1, port 17 and switch S3 port 17 form the ISLs.

   ```
   zone --create -t ti FCVI_TI_S1S3 -o n -p 1,0; 1,1; 1,17; 3,0; 3,1;
   3,17
   ```

   The following example creates FCVI zone *FCVI_TI_S2S4* on switch S2. The ports 0 and 1 are the FCVI ports at switch S2 and S4. Switch S2, port 17 and switch S4 port 17 form the ISLs.

   ```
   zone --create -t ti FCVI_TI_S2S4 -o n -p 2,0; 2,1; 2,17; 4,0; 4,1;
   4,17
   ```

   > **Note:** In a shared-switches configuration, you must set the TI failover as disabled. This is indicated by the option -o and -p.

   b. Activate the TI zone by entering the following command:

   ```
   zone --activate zonename
   ```

   Example
   The following example activates the TI zone called FCVI_TI_S1S3:

   ```
   zone --activate FCVI_TI_S1S3
   ```

3. Create the FCVI zone by entering the following command:

   ```
   zonecreate "zonename","portlist"
   ```

   Example
   Create the FCVI zone *FCVI_MC1_S1S3* and *FCVI_MC2_S1S3* at switch S1 by entering the following command:

   ```
   zonecreate "FCVI_MC1_S1S3", "1,0;3,0"
   ```

   ```
   zonecreate "FCVI_MC2_S1S3", "1,1;3,1"
   ```

   Similarly, create the FCVI zone *FCVI_MC1_S2S4* and *FCVI_MC2_S2S4* at switch S2:

   ```
   zonecreate "FCVI_MC1_S2S4","2,0;4,0"
   ```

   ```
   zonecreate "FCVI_MC2_S2S4", "2,1;4,1"
   ```

4. Create the storage zone at each MetroCluster configuration.

If you are creating...	Then...
A single storage zone	Create the storage zone by entering the following command:
	`zonecreate "zonename""portlist"`
	The following example creates storage zones, *STOR_MC1_S1S3* and *STOR_MC2_S1S3*, at switch S1:
	`zonecreate "STOR_MC1_S1S3", "1,2; 3,2; 1,4; 1,5; 3,4; 3,5"`
	`zonecreate "STOR_MC2_S1S3", "1,3; 3,3; 1,6; 1,7; 3,6; 3,7"`
	Similarly, create the storage zones, *STOR_MC1_S2S4* and *STOR_MC2_S2S4*, at switch S2:
	`zonecreate "STOR_MC1_S2S4", "2,2; 4,2; 2,4; 2,5; 4,4; 4,5"`
	`zonecreate "STOR_MC2_S2S4", "2,3; 4,3; 2,6; 2,7; 4,6; 4,7"`

If you are creating...	Then...

Single-loop storage zones (for Fabric OS 6.3.1c and later)	Create the storage zones by entering the following command:
	
	`zonecreate "zonename", "portlist"`
	
	The following example creates the storage zones, *STOR_S1_4*, *STOR_S1_5*, *STOR_S3_4*, *STOR_S3_5*, *STOR_S1_6*, *STOR_S1_7*, *STOR_S3_6*, and *STOR_S3_7* at switch S1:
	
	`zonecreate "STOR_S1_4", "1,2; 3,2; 1,4"`
	
	`zonecreate "STOR_S1_5", "1,2; 3,2; 1,5"`
	
	`zonecreate "STOR_S3_4", "1,2; 3,2; 3,4"`
	
	`zonecreate "STOR_S3_5", "1,2; 3,2; 3,5"`
	
	`zonecreate "STOR_S1_6", "1,3; 3,3; 1,6"`
	
	`zonecreate "STOR_S1_7", "1,3; 3,3; 1,7"`
	
	`zonecreate "STOR_S3_6", "1,3; 3,3; 3,6"`
	
	`zonecreate "STOR_S3_7", "1,3; 3,3; 3,7"`
	
	Similarly, create the storage zones for switch S2:
	
	`zonecreate "STOR_S2_4", "2,2; 4,2; 1,4"`
	
	`zonecreate "STOR_S2_5", "2,2; 4,2; 1,5"`
	
	`zonecreate "STOR_S4_4", "2,2; 4,2; 3,4"`
	
	`zonecreate "STOR_S4_5", "2,2; 4,2; 3,5"`
	
	`zonecreate "STOR_S2_6", "2,3; 4,3; 1,6"`
	
	`zonecreate "STOR_S2_7", "2,3; 4,3; 1,7"`
	
	`zonecreate "STOR_S4_6", "2,3; 4,3; 3,6"`
	
	`zonecreate "STOR_S4_7", "2,3; 4,3; 3,7"`

5. Create the configuration zone.

If you have...	Then...
A single storage zone	Enter the following command to create the configuration zones: **cfgcreate *"configuration_name""zonename; zonename"*** The following example creates a configuration *Config_MC_S1S3* that includes the FCVI zones *FCVI_MC1_S1S3* and *FCVI_MC2_S1S3*, and the storage zones *STOR_MC1_S1S3* and *STOR_MC2_S1S3* at site A: **cfgcreate "Config_MC_S1S3", "FCVI_MC1_S1S3; FCVI_MC2_S1S3; STOR_MC1_S1S3; STOR_MC2_S1S3"**
Single-loop storage zones (for Fabric OS 6.3.1c and later)	Enter the following command for each disk loop port that includes the port and the ports containing the FC controller connections: **zonecreate *"zonename", "portlist"*** The following example creates a configuration *Config_MC_S1S3* that includes the FCVI zones *FCVI_MC1_S1S3* and *FCVI_MC2_S1S3*, and all the single-loop storage zones at site A. **cfgcreate "Config_MC_S1S3", "FCVI_MC1_S1S3; FCVI_MC2_S1S3; STOR_S1_4; STOR_S1_5; STOR_S1_6; STOR_S1_7; STOR_S3_4; STOR_S3_5; STOR_S3_6; STOR_S3_7"**

6. Save the configuration by entering the following command:

   cfgsave

7. Enable the configuration by entering the following command:

   cfgenable *"Configuration_name"*

   Example
   The following example saves and enables the *Config_MC_S1S3* configuration:

   cfgenable "Config_MC_S1S3"

8. Verify the configuration by entering the following command:

   zone --show

   The configuration shows the TI zone *FCVI_T1_S1S3* and *FCVI_TI_S2S4* correctly configured.

9. Verify the FCVI and STOR zones by entering the following command:

   cfgshow

The configuration shows the FCVI and storage zones correctly configured.

Copyright and trademark information

(INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

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

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

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, FAServer, FastStak, FilerView, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service

Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe on any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, N.Y. 10504-1785
U.S.A.

For additional information, visit the web at:
http://www.ibm.com/ibm/licensing/contact/

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

Index

IBM ®