



**Best Practices in the World of Interopmode/ISL  
(Multivendor Environment Blade Servers)  
&  
Troubleshooting FSPF/VSAN & DOMAINS in Cisco MDS  
Environments**

**EMC Proven™ Professional Knowledge Sharing 2008**

Sanjay Sood  
Solutions Architect  
(EMC Resident)  
sood\_sanjay@emc.com

## **Table of Contents**

Objective .....Page 3

Introduction.....Page 4

**Chapter 1**.....Page 7

### **Fabric Switch Mode (ISL Technology/ No Access Mode)**

#### Cases

1- Brocade Embedded Switch (In Blade Server) <-ISL-> Brocade Edge Switches.  
(Page 7 – Page 22)

2- Legacy McData Embedded Switch (In Blade Server) <-ISL-> Legacy McData  
Edge Switches (Now Brocade). (Page 23 – Page 31)

3- ISL in Multivendor Environment: Brocade Embedded Switch (In Blade Server)  
<-ISL-> Cisco MDS Edge Switch. (Page 32 – Page 48)

**Chapter 2**.....Page 49

### **Cisco MDS Fabric VSAN, Domain and FSPF Troubleshooting**

Part A: Troubleshooting VSAN issues..... (Page 54 – Page 57).

Part B: Troubleshooting Domain ID Issues..... (Page 58 – Page 61).

Part C: Troubleshooting FSPF Issues..... (Page 62 – Page 71).

**Appendix-A**..... (Page 72 – Page 75).

*Disclaimer: The views, processes or methodologies published in this compilation are those of the authors. They do not necessarily reflect EMC Corporation's views, processes, or methodologies.*

## **Objective**

This article provides best practices on Interop Mode/InterSwitch Link (ISL) configuration topologies in Open Fabric Multivendor Environment (on Blade Servers), and also focuses on the importance of ISL in the embedded products server market. Server Consolidation/Virtualization conserves power, eases management and provides overall cost effectiveness due to Blade Servers available from various OEM vendors like Dell, IBM, and HP.

The concept of Interop/ISL on blade servers may be confusing as there are various OEM vendors like Brocade/McData (now Brocade), Qlogic, Emulex, and Cisco that offer fiber channel switches. On the other hand, such broad vendor choice offers customers the flexibility to customize blade servers per their requirements and their business affordability. This article will focus on the integration and best practices of some of the most popular embedded and EMC®-approved products available and also focus on bringing scattered information together.

This document is split into two parts.

**Chapter 1** – Best Practices on Interop Mode/ISL Technology on Blade Servers  
(Multivendor Environment) (Page 7 – Page 51)

**Chapter 2** – The Cisco MDS Fabric FSPF/VSAN & Domain Troubleshooting  
(Page 52 – Page 75)

## Introduction

Storage requirements, SAN fabric dependability and the pursuit of interoperability continue to grow. SAN solutions are similarly growing in both size and complexity. As the number of switches in the fabrics increase, fabric management becomes increasingly complex. SAN solution demand continues to grow as companies require more computing resources.

### **Key terms used:**

The following terms are used in this document.

Edge switch	Fabric switch that connects host, storage, or other devices, such as Brocade Embedded Switch, to the fabric.
E_Port	An InterSwitch link (ISL) port. A switch port that connects switches to form a fabric.
F_Port	A fabric port. A switch port that connects a host, host bus adaptor (HBA), or storage device to the SAN. On Brocade Access Gateway, the F_Port connects to a host only.
N_Port	A node port. A Fibre Channel host or storage port in a fabric or point-to-point connection. On Brocade Access Gateway, the N_Port connects to the edge switch.
N_Port ID virtualization (NPIV)	Allows a single Fibre Channel port to appear as multiple, distinct ports providing separate port identification and security zoning within the fabric for each operating system image as if each operating system image had its own unique physical port.
Access Gateway (AG)	Fabric OS mode for embedded switches that reduces storage area network (SAN) deployment complexity by leveraging NPIV.

You can convert an existing Brocade SAN Switch Module to Access Gateway mode through the Fabric OS command line interface (CLI) or Web Tools. SAN administrators can perform a firmware upgrade and run a simple CLI command or use the Brocade Web Tools to easily revert Access Gateway mode back to normal Fabric Mode.

#### Mapping

On Brocade Access Gateway, the configuration of F\_Port to N\_Port routes.

# Chapter 1

## Fabric Switch Mode (ISL Technology/No Access Mode)

### Focus:

The basic concept of ISL will always remain the same. Every OEM manufacturer has a proven track record when ISL is performed between the same brands of switches. But, when it comes to Interopmode parameters in Open Fabric multivendor environments, it is sometimes difficult to find a common ground. This portion of the article will focus on best practices when deploying ISL in a multi-vendor environment.

Consider an example between the most popular brands available in the industry. On one hand, there is a Brocade 4GB SAN solution available for Blade Servers which will be acting as an embedded switch. This will, in turn, be connected to the edge switches. This Brocade embedded switch will form an ISL with the edge switches. This article will cover the formation of ISL's connectivity by discussing the following three cases.

### Cases

- 1- **Brocade Embedded Switch (In Blade Server) <-ISL-> Brocade Edge Switches connectivity.** (Page 8 – Page 25)
- 2- **Legacy McData Embedded Switch (In Blade Server) <-ISL-> Legacy McData Edge Switches (now Brocade).** (Page 26 – Page 33)
- 3- **ISL in Multivendor Environment : Brocade Embedded Switch (In Blade Server) <-ISL-> Cisco Edge Switch** (Page 34 – Page 51)

In a complex scenario, the edge switches are connected to the core switches but the scope of this document will be limited to embedded switch <=> edge switches ISL connectivity and their parameters.

If the switch feature, Access Gateway (EMC Supported – refer to EMC Support Matrix) mode, is not used, the switch will work as a regular Fabric switch and you should be able to use all the ISL features. We will cover the ISL interconnectivity on a case-by-case basis between similar vendor and multivendor environments.

## 1. Brocade (Embedded Switch) ⇔ Brocade (Edge Switch) ISL Connectivity

### Connectivity (ISL Trunking) Best Practices

Focus - Blade Server Embedded Switch

ISL Trunking optimizes network performance, availability, and manageability by merging multiple ISLs into a single logical entity, called the Trunk Group. The trunk group offers these beneficial characteristics.

Scenario: This is a 4GB end-to-end SAN/Storage Solution on EMC /Brocade technology.

Best Performance:

- 4 Gbps end-to-end bandwidth from server to storage
- Multiple 12 Gbps Inter Switch Link (ISL) Trunks connecting the Blade Servers to the external SAN fabric providing up to 48 GB of total load balanced bandwidth.

(EMC Supports 8 ISL from Domain-to-Domain/Fabric)

*Note:* Refer to the Latest EMC Support Matrix or the Storage Vendor Support Matrix before using this solution in your production environment.

Simplified management

- Single point of SAN fabric management
- Advanced performance monitoring capability

High Availability/Redundancy and Monitoring

- Redundant SAN fabrics with no single point of failure
- EMC Multi-Pathing software support
- Switch firmware upgrades
- Monitoring tools monitor failures before component failures

### Security and Scalability

- Additional Blade Server chassis and storage arrays may be connected to the core switches without any disruption
- Encryption, authentication, and access control lists can be utilized to meet the strictest corporate security policies

### Lower Total Cost of Ownership (TCO)

- By improving SAN performance, availability, and manageability, ISL Trunking requires fewer ISLs, thus freeing FC ports and reducing complexity

Figure 1 shows the ISL formation between the Blade Servers Embedded Switches and the Edge Switches.

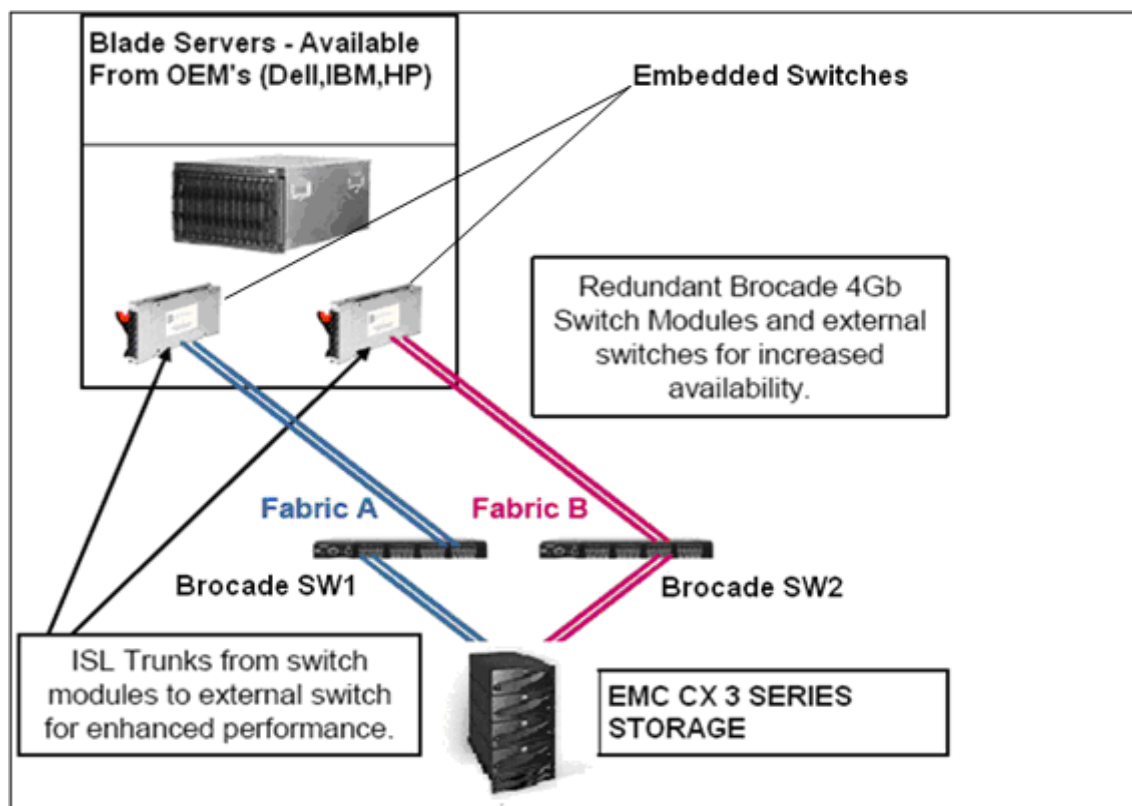


Figure 1: ISL between Blade Servers and Edge Switches

(Picture (Created) Resources -Online – See Annexure – A for More Details)



The Brocade 4Gb SAN Switch Embedded Module ISL Trunking feature allows two trunk groups with up to three ISL connections each between itself and any other Brocade switch that has an installed ISL Trunking license. Each trunk group allows **three ISLs** to merge logically into a **single 12 Gbps link** between switches (Figure 2). It is compatible with both short wavelength (SWL) and long wavelength (LWL) fiber optic cables and transceivers.

*Note:* This is a Brocade Embedded switch and External Edge/Core switch ISL. The concept is the same when two same-vendor or multi-vendor switches are connected to each other. (In multi-vendor environments there are a few other parameters like interop mode/Domain ID restriction that apply. We will focus on that later in this document)




Figure 2: Six ISLs become two 12Gb ISL Trunks

*(Picture Resources -Online – See Annexure – A for More Details)*

#### Dynamic Path Selection – (Enabled with License)

In Figure 2, the Brocade 4Gb SAN Switch Module also offers a performance enhancing feature called Dynamic Path Selection (DPS). DPS is a routing scheme that optimizes fabric-wide performance by automatically routing data to the most efficient available path in the fabric. While ISL Trunking can balance traffic at the most granular level (the FC frame), DPS balances loads at the FC Exchange level such as a SCSI read or write. DPS augments ISL Trunking to provide more effective load balancing in certain configurations, such as routing data between multiple trunk groups. As a result, a combination of DPS and ISL Trunking provides the greatest design flexibility and the highest degree of load balancing

### Trunking Prerequisite (Blade Server Embedded Switch)

On the Brocade 4Gb SAN Switch Module, all trunking ports must meet the following prerequisites:

- \_ There must be a direct connection between participating switches
- \_ Trunk ports must reside in the same port group
- \_ Trunk ports must run at the same speed (either 2 Gbps or 4 Gbps)
- \_ Trunk ports must be E\_Ports
- \_ Cable lengths for participating links should differ by no more than 500 meters
- \_ Trunk ports must be set to the same ISL mode (L0 is the default)

This feature is not supported in interoperability mode. For more information about ISL Trunking in general or Trunk Groups on other Brocade switch models, refer to the *Brocade Fabric OS Administrator's Guide*.

Figure 3 – displays the Trunk Grouping in Embedded Switch.

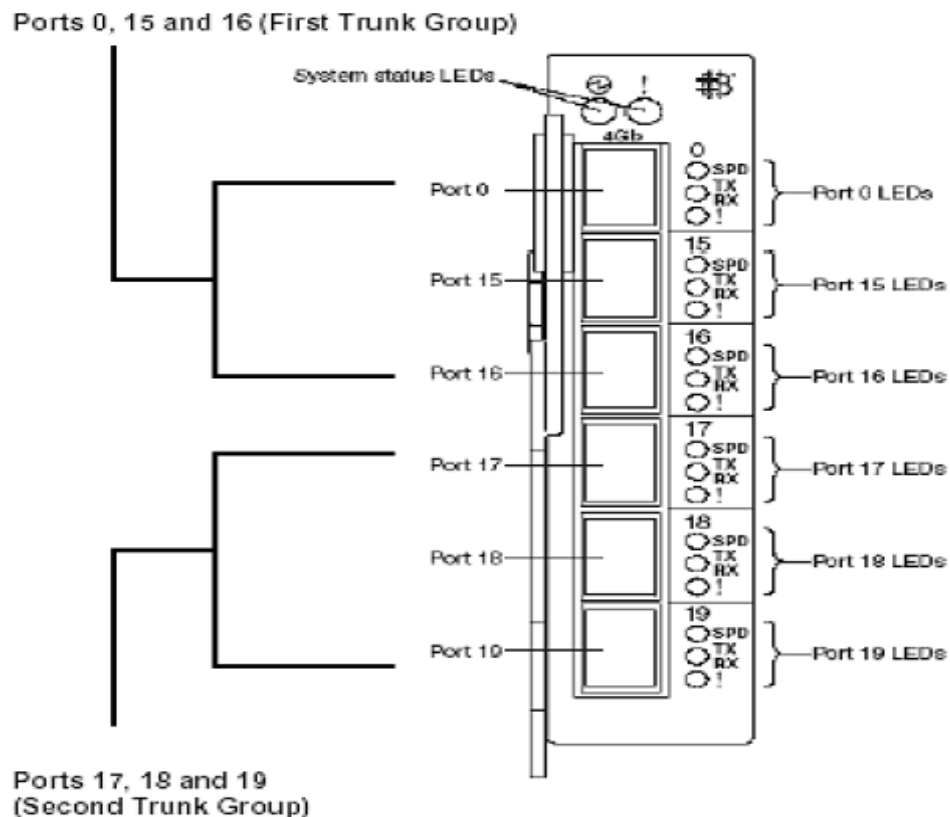


Figure 3: (Embedded Internal Switch) **ISL Trunk Groups** for 4Gb SAN Switch Module

(Picture (Modified) Resources Brocade 4gb SAN Solution – See Annexure – A for More Details)

Figure 4 shows the port groups that can be used to form the ISL with the Embedded Switch trunk groups shown in Figure 3.

### Example

Ports (0,15,16) in the Embedded switch (Figure 3) can be connected to 0-7 (First Trunk Group – see Figure 4) and Ports 17, 18, 19 can be connected to 8-15 (Second Trunk Group – see Figure 4). Generally, installations are basic and customers are not even aware of this feature. We will cover this in Trunking with screenshots.

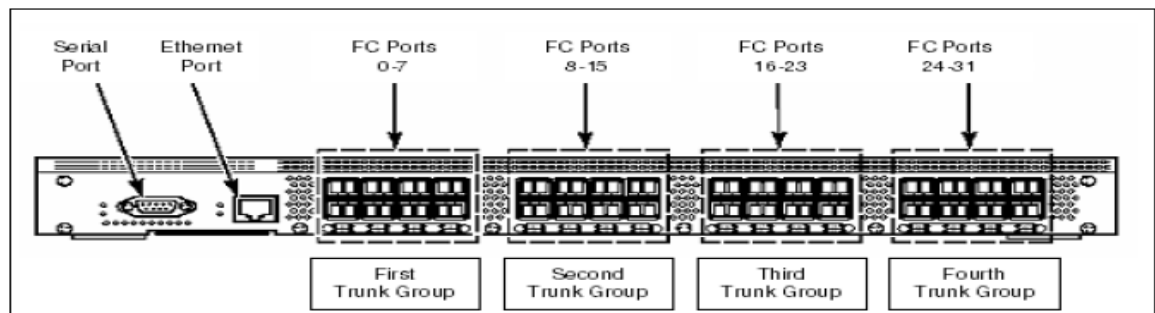


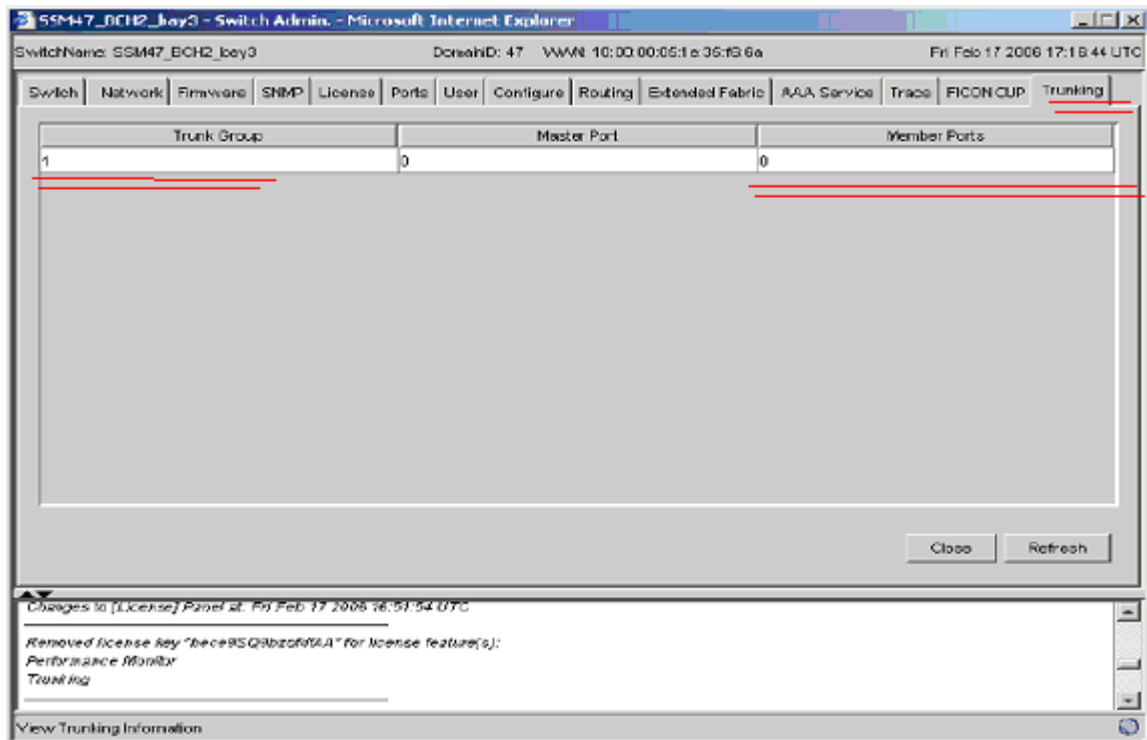
Figure 4: External Brocade Core/Edge Switch

- (Picture Resources Brocade Online – See Annexure – A for More Details)

-

If **ISL Trunking licenses** are installed on switches before the SAN fabric is built, follow these steps:

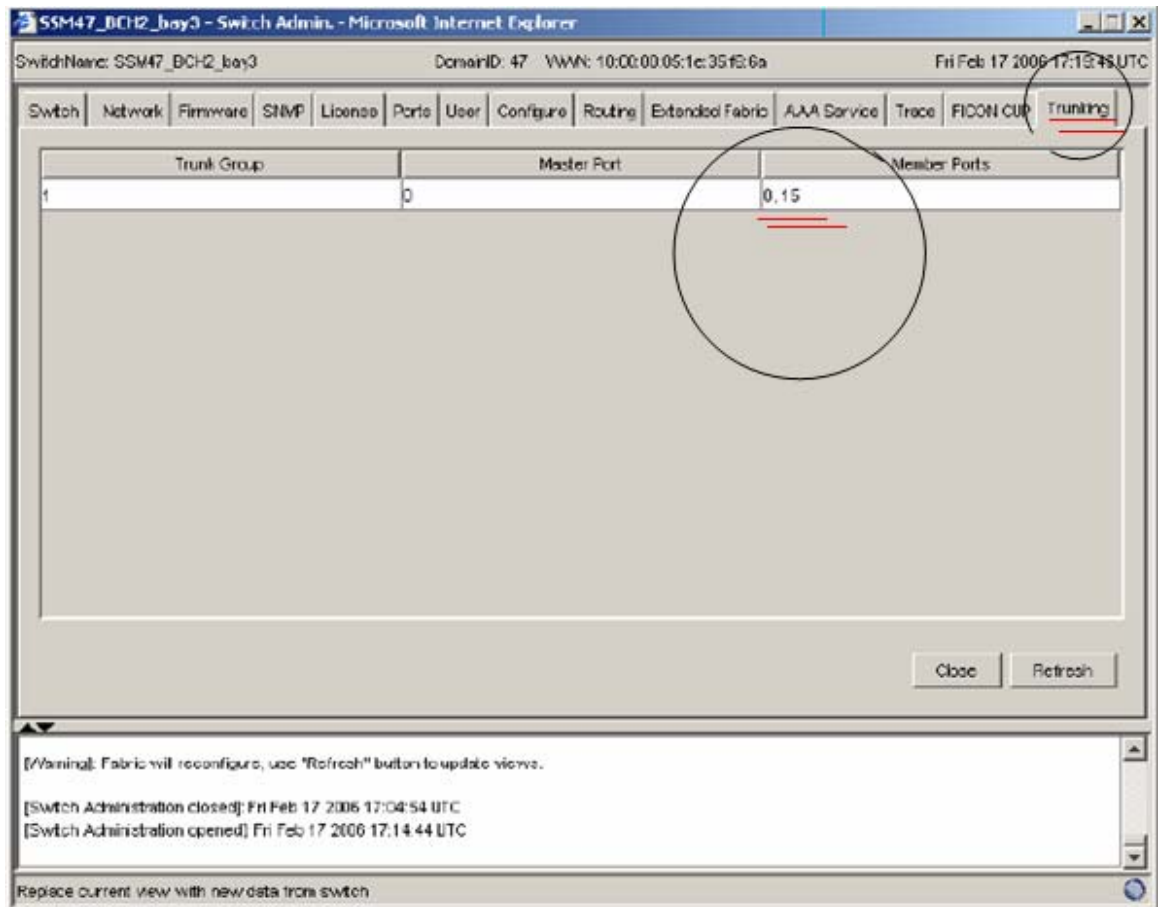
1. Connect the initial link between the switch module and external switch. This depends on how many ports are free on Edge Switches).
2. Verify that a Trunk group has been formed.  
In Web Tools, click the Admin button and log in.
3. Select the **Trunking tab** (Underlined red color) and verify that the trunk group has formed - shown in Figure 5.



*Figure 5: Trunking Tab to verify Trunk Groups.*

*(Picture (Modified) Resources Brocade 4GB SAN solution) – See Annexure – A for More Details*

4. Connect additional links between the switch module and external switch trunk port groups. Allow the links to establish.
5. On the Admin Trunking window, click the Refresh button and view the Member Ports column. Verify that the new link is now a member of the same Trunk group as shown in Figure 6.



*Figure 6: Multiple Ports forming Trunks*

*(Picture (Modified) Resources Brocade 4GB SAN solution) – See Annexure – A for More Details)*

6. **Repeat**, adding cables and verifying trunk members as necessary. The Edge Switch supports trunk groups up to **eight members (EMC Supported – refer to the latest support matrix)** and up to **32 Gb of bandwidth**. The Brocade Embedded 4GB SAN Switch supports up to two **three-member** trunk groups of **12 Gb each**.
7. That's it. This completes building trunks for Fabric A. Repeat the steps for Fabric B. You can now introduce the EMC Storage array or any other storage at this time into the Edge Switch.

### **Other configuring factors on Brocade 4Gb SAN Switch Module**

Please check your Blade Server OEM documentation for setting up the Blade Server as this portion only covers the switch (embedded) side.

### **Supported hardware and software**

This document is specific to Fabric OS v5.2.1 or higher running on the Brocade SilkWorm 4012, 4016, 4020, and 4024 embedded switches.

Starting from Fabric OS 5.0.2, the Blade Server-embedded products were qualified to connect with the edge switches. For best results, depending on your edge switches' fabric OS, refer to the EMC support matrix for ISL connectivity.

### **Switch Setup and Best Practices (via Web Tools)**

Launch Web Tools management software from the Management Module or directly by using a browser and the switch's IP address. This section will describe how to set up the browser for best performance and launch Web Tools external to the Management Module.

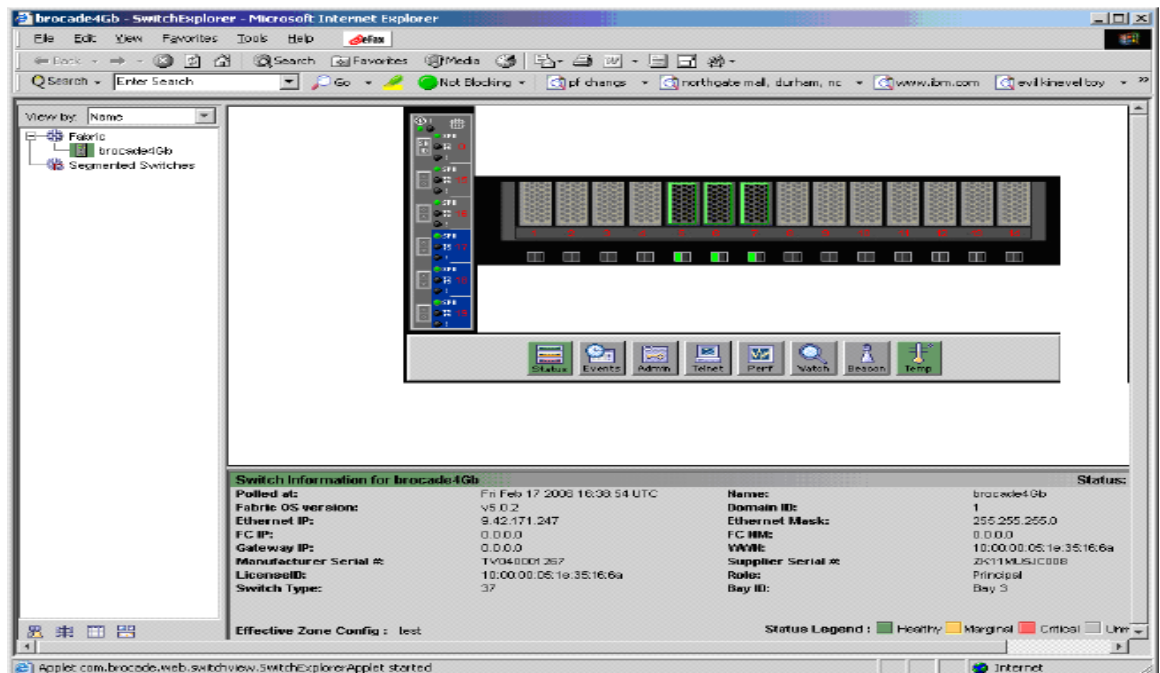
Web Tools requires any browser that conforms to HTML Version 4.0, JavaScript™ Version 1.0, and Java™ Plug-in 1.4.2\_06 or higher. You can launch Web Tools on any workstation with a compatible Web browser installed. Brocade Fabric Manager Administrator's Guide v5.2.0, can be downloaded from Brocade Connect via registration. Figure 7 shows the Embedded switch for Blade Servers.



*Figure 7: Embedded switch.*

*(Picture Resources Brocade 4GB SAN solution) – See Annexure – A for More Details)*

From the Internet Explorer or Mozilla browser, enter the switch IP address in the Address window box, then press Enter. A Web Tools session will open, similar to that shown in Figure 8. The Web Tools display for the SSM is somewhat different for external switches, since there are no Power Supply Fan buttons.



*Figure: 8 Web Tools run directly from Brocade switch*

*(Picture Resources Brocade/IBM 4GB SAN solution) – See Annexure – A for More Details)*

Figure 9 shows the External/Internal port mapping. Keep in mind that the external port connects to edge switches to make ISL, and internal ports make a connection to internal HBA's (Qlogic / Emulex, etc).

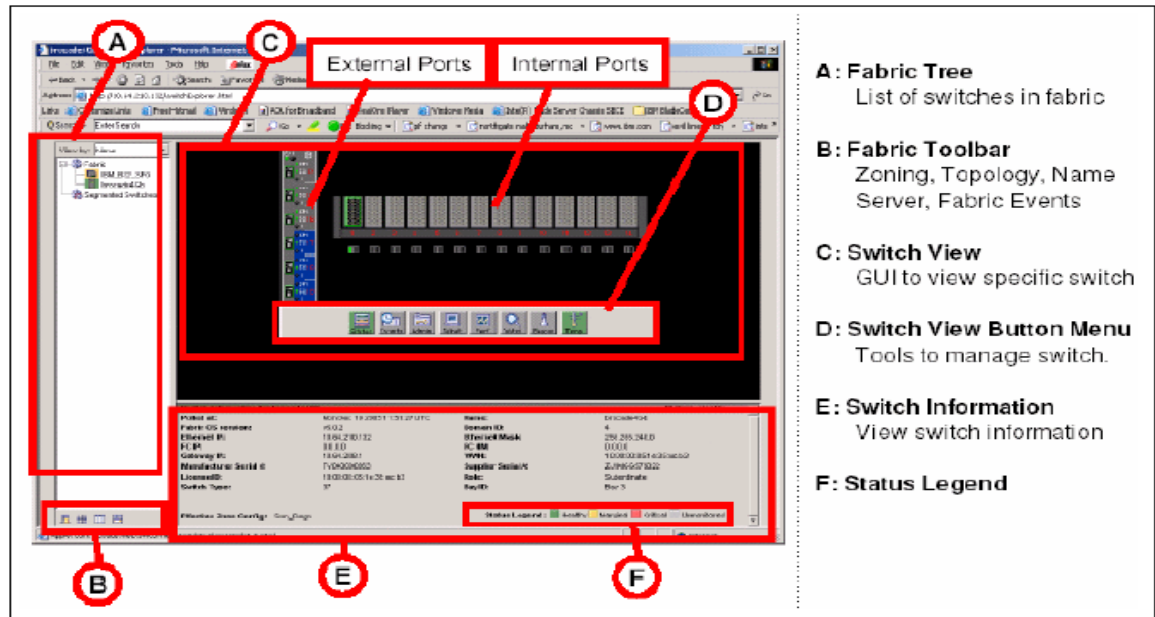


Figure 9: Internal – External Port Mapping

(Picture (Modified) Resources Brocade/IBM SAN solution) – See Annexure – A for More Details)

Check switch health

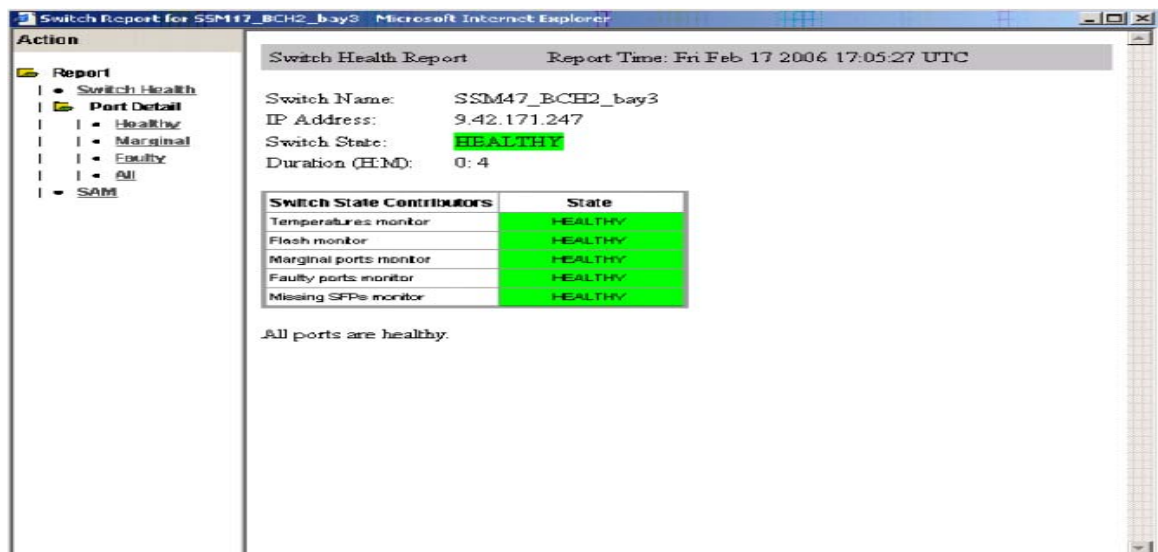


Figure 10: Switch Health

(Picture (Modified) Resources Brocade/IBM SAN solution) – See Annexure – A for More Details)

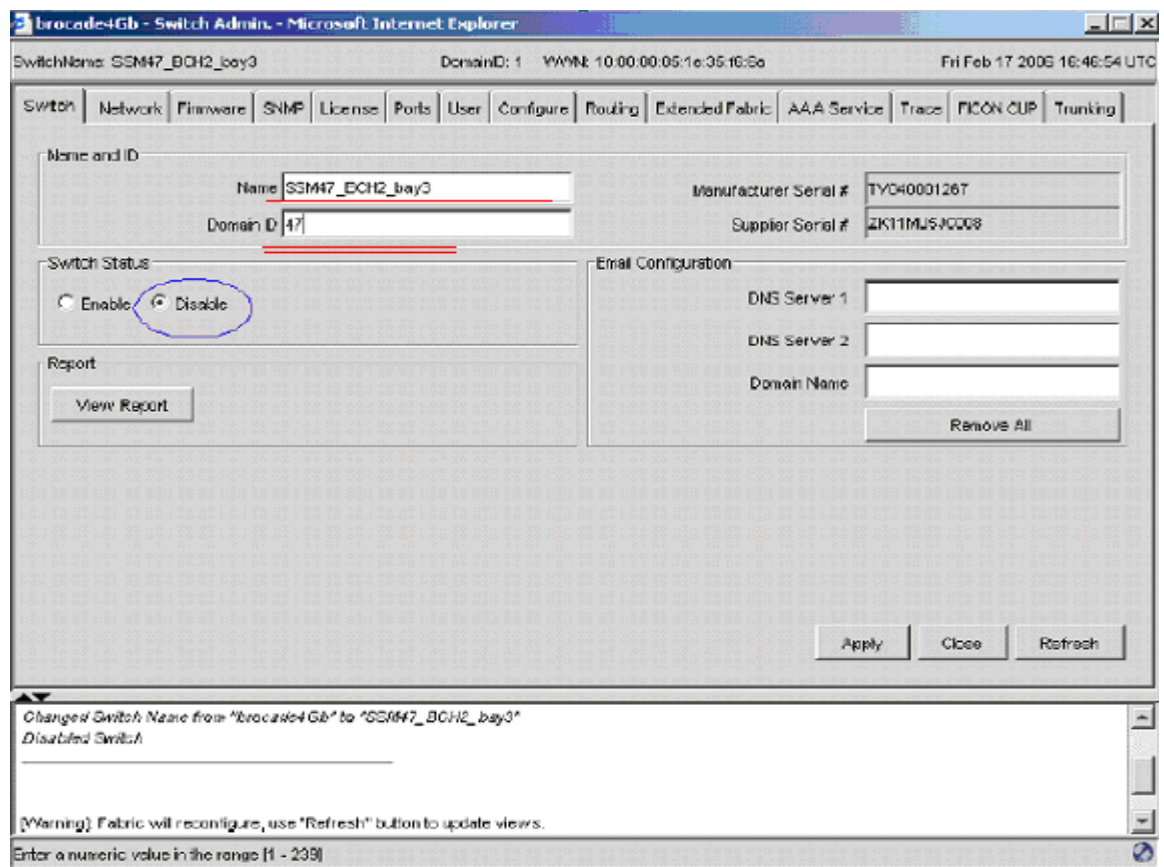


Select the Admin button from the Switch View Button menu. Log in to the switch using the following information:

Login USERID

Password PASSWORD (note that 0 is a zero)

Set the Switch Status to **Disable**, as shown in **Figure 11**. You must disable a switch before changing Domain ID.



*Figure 11: Disable switch before changing domain ID.*

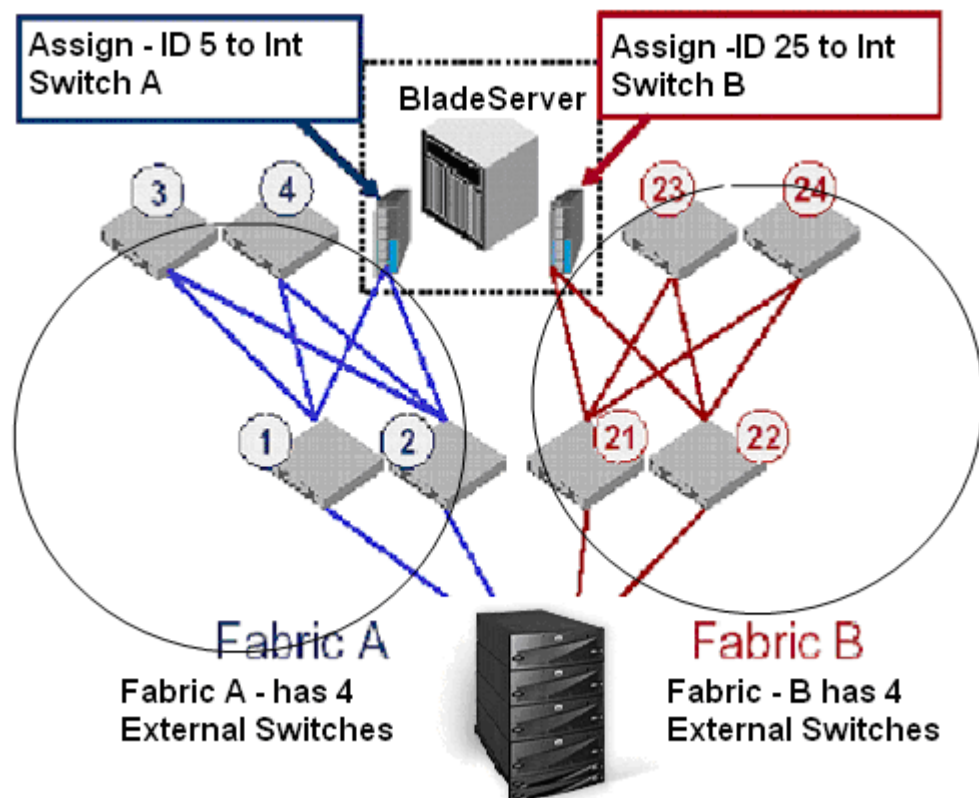
*(Picture (Modified) Resources Brocade/IBM SAN solution) – See Annexure – A for More Details)*

Select the Domain ID of the switch.

Guidelines for selecting Domain ID numbers:

- Set unique domain numbers for each switch in the SAN. This simplifies merging of fabrics if needed in the future.
- As a convention, consider setting the domain ID of each switch to the last octet of its IP address.
- Be aware that the highest allowed domain number is 239.

Figure 12 provides an example of both external switches and embedded switch modules.



*Figure 12: Selecting Domain ID numbers*

*(Picture (Created) Resources EMC/Brocade Online) – See Annexure – A for More Details)*

Enter the selected number into the Domain ID field.

In this same window (Figure 11), set the switch's name to some unique and descriptive name. The switch name can include up to 15 characters, must begin with a letter, and must consist of letters, digits, underscore characters, and no spaces. Setting the switch name is recommended to simplify locating switch module and fabric management.

### **Best Practices for New Integrations**

For new SAN installs, select the Configure tab (*Figure 13*) and verify that Switch PID Format is set to Format 1. Occasionally, when connecting to legacy external switches, you may need to change the Switch Port ID format, which is an EMC best practice. Refer to the EMC Support Matrix to ensure the list of parameters that you may wish to change.

Please pay attention to the following statements:

For new SAN installs (*with the new 5.2.1 FOS version there is no need to set PID format*), select the Configure tab (Figure 13) and verify that Switch PID Format is set to Format 1. You may need to change the Switch Port ID format on your edge switches as well. When integrating new embedded switches with SilkWorm 2000, 3200, and 3800 series switches, it is recommended to set the Core PID format to 1 on SilkWorm 2000, 3200, and 3800 series switches only.

*Highly Recommended: Set the Core PID on all switches running Fabric OS 2.x and 3.x (Refer to the Brocade SilkWorm Scalability Support Matrix (part number: 53-0000618-02 or later). I hope that this has minimized confusion! If not, read why we need a Core PID Format Change section.*

On new Brocade 4020 Embedded Switches there is no Port limitation (Refer to the latest Brocade Documentation from BrocadeConnect).

*Note:* Run the EMC support matrix and refer to the Brocade Switch Guide (because of the newer version of embedded switch connectivity to Brocade Legacy Switches in your environment, the PID format needs your attention). See Figure 13 marked in red.

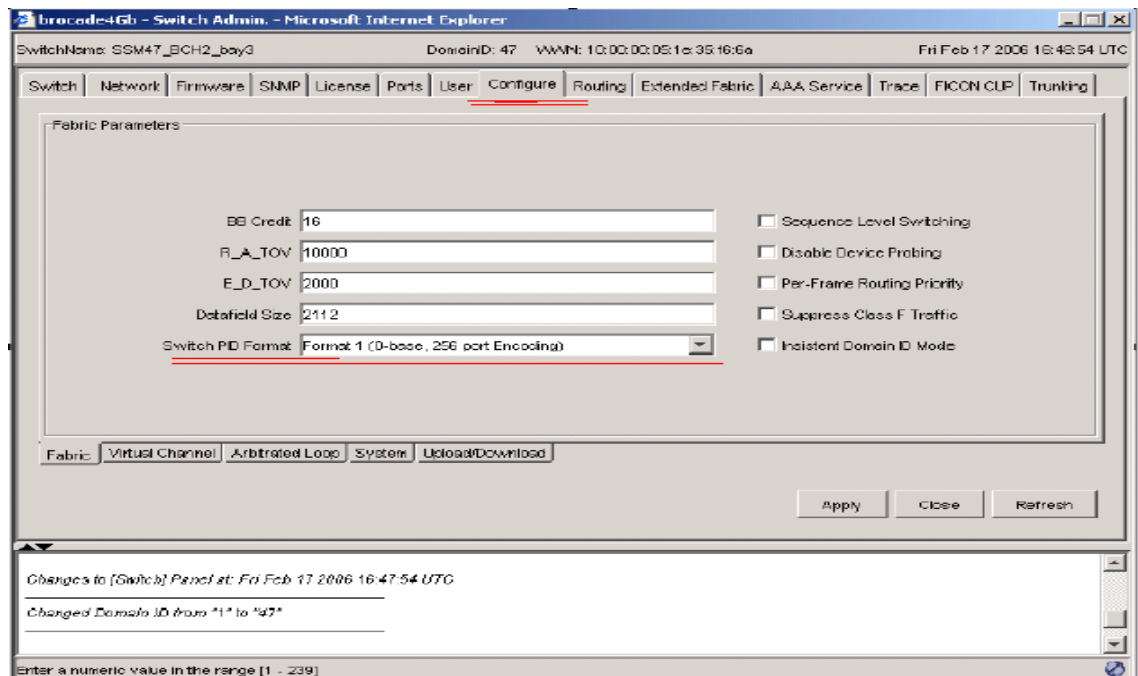


Figure - 13

*(Picture(Modified) Resources Brocade/IBM SAN solution ) – See Annexure – A for More Details)*

### Why we need Core PID Format Change

Many customers and even implementation personnel forget or become confused when setting Core PID values, so we will focus on PID importance.

When switches are added to the fabric, the Core PID format must be changed on lower port count switches. High port count Director switches, such as the SilkWorm 3900 and 12000, have a default Core PID format of 1.

Incompatibility with the Core PID format will segment the fabric until all lower port count switches have been changed to Core PID format of 1. Upgrading the new Core PID format on an existing switch running Fabric OS V2x/V3.x is a two-step process. When setting the Core PID format, the minimum Fabric OS versions are: Fabric OS 2.6.0c or greater for the SilkWorm 2000 series and Fabric OS 3.0.2c or greater for the SilkWorm 3200 and 3800.

You must upgrade any prior versions to the most recent available Fabric OS. After upgrading the Fabric OS, set the switch configuration. Both of these steps require disabling each switch in the fabric. Disabling and enabling a switch results in fabric disruption and may pause I/O. The fabric will remain segmented until all switches in the fabric are configured to the Core PID format of 1.

Please note the following:

The *default* Core PID setting for SilkWorm 3900 and 12000 switches is a *Core PID format of 1*.

To prevent the fabric from segmenting, set the *Core PID Format setting to 1* on the SilkWorm 3800, 3200, or 2000 Series switches in a fabric with a SilkWorm 3900 or 12000.

The *default* Core PID setting for SilkWorm 3800, 3200, and 2000 Series switches is Core PID Format-0. Verify the Core PID Format from the *configshow* output by referring to the “*fabric.ops.mode.pidFormat*” parameter.

An upgrade of the Fabric OS is required only if the current version is lower than the minimum specified below. However, it is always a best practice to upgrade the Fabric OS to the version recommended by your support provider.

- SilkWorm 2800 Fabric OS version 2.6.0c or higher is required
- SilkWorm 3800/3200 Fabric OS version 3.0.2c or higher is required

*Note:* Recommended minimum Fabric OS versions are 2.6.0c and 3.0.2c or higher.

### How can we do it using CLI mode, if required?

Once the Fabric OS has been upgraded to the recommended level, you can change the Core PID format.

- a. Disable the switch using the **switchdisable** command.
- b. Issue the **Configure** command. A list of configurable parameters will become available.
- c. Change the Core Switch PID Format from default 0 to 1, accept all other parameters as default.
- d. Enable the switch using the **switchenable** command.

Steps are shown below.

```
Security126: admin> switchDisable
Security126: admin> configure
Configure...
Fabric parameters (yes, y, no, n): [no] y
Domain: (1..239) [5]
BB credit: (1..27) [16]
R_A_TOV: (4000..120000) [10000]
E_D_TOV: (1000..5000)[2000]
Data field size: (256..2112) [2112]
Sequence Level Switching: (0..1)[0]
Disable Device Probing: (0..1) [0]
Suppress Class F Traffic: (0..1)[0]
SYNC IO mode: (0..1) [0]
VC Encoded Address Mode: (0..1) [0]
Core Switch PID Format: (0..1) [0] 1
Per-frame Route Priority: (0..1) [0]
Long Distance Fabric: (0..1) [0]
Virtual Channel parameters (yes, y, no, n): [no]
Switch Operating Mode (yes, y, no, n): [no]
Zoning Operation parameters (yes, y, no, n): [no]

Security126: admin> switchenable
```

## 2. McData (embedded in Blade Server) <-> OEM's Edge Switches McData

Here we will focus on how to setup a McData Embedded (4Gb Switch Modules) to attach to a McData Edge Switch.

Figure 14 presents the internal architecture of one of the OEM blade server manufacturers – in this instance, IBM Blade Center.

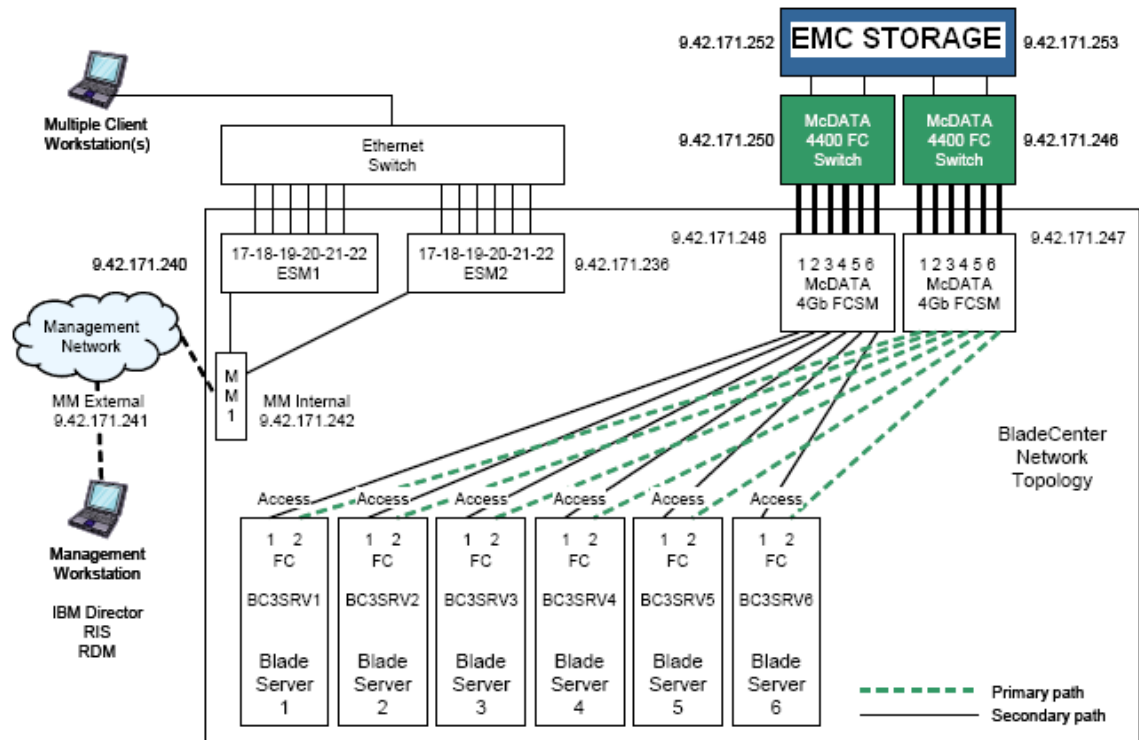


Figure 14: Internal Blade Center Architecture

(Picture (Modified) Resources Brocade (Legacy McData) /IBM SAN solution) – See Annexure – A for More Details)

## McDATA 4Gb Embedded Fibre Channel Switch Modules setup:

Follow the steps listed below to configure the McData Switch:

- a. Point your browser to the IP address of the McDATA EFCM, or use McData's default IP address to Open the Element Manager (10.1.1.10 / 255.0.0.0). You will need to make sure your current Java version is at or above 1.4.2.
- b. Log in to the device.
- c. Click the Topology window in the left pane, click the drop-down tree and double-click the IP address. Next, select your switch (in our example, McDATA4Gb). Select Switch → Switch Properties...as shown in Figure 15.

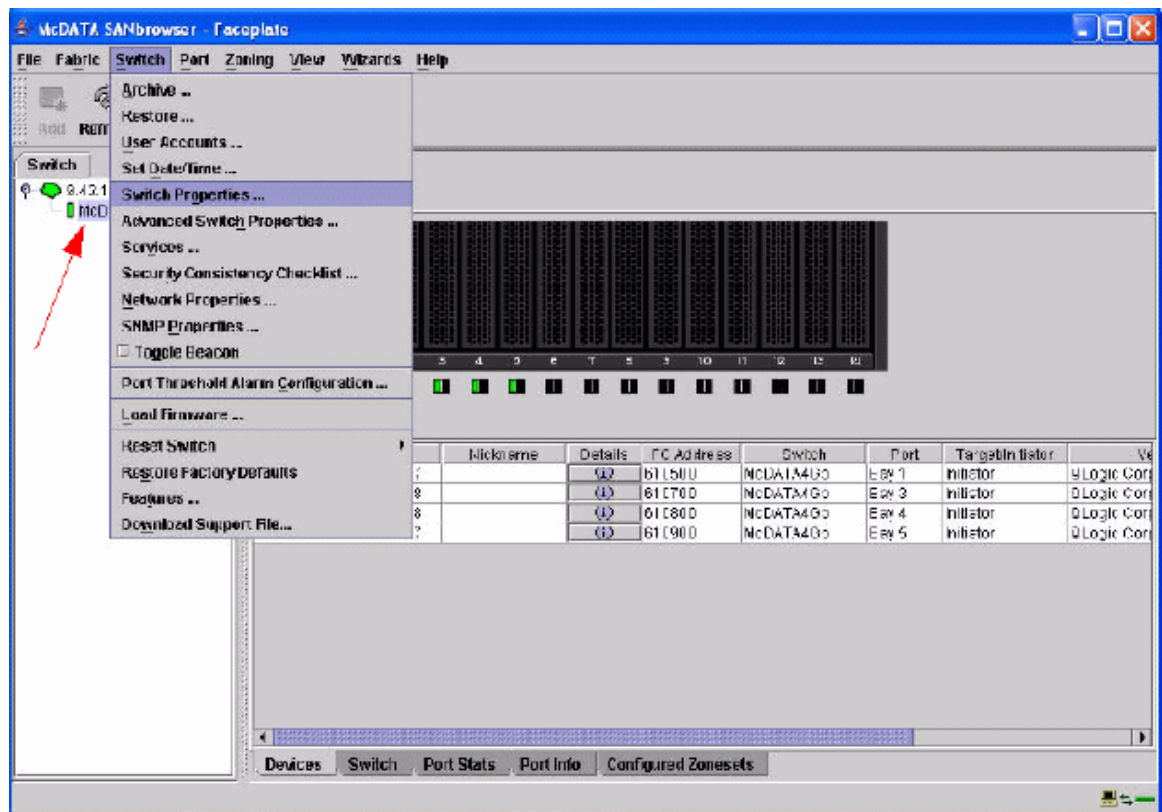


Figure 15: McData Switch properties setting.

(Picture Resources Legacy McData / SAN solution) – See Annexure – A For More Details)



- d. From the Switch Properties dialog (Figure 16), enter a symbolic name for your switch. Click the Enable radio button next to Domain ID lock (1). Enter a Unique Domain ID ranging from 1-31 (2). In this example, 2 has been selected for the DID. Click OK to continue.

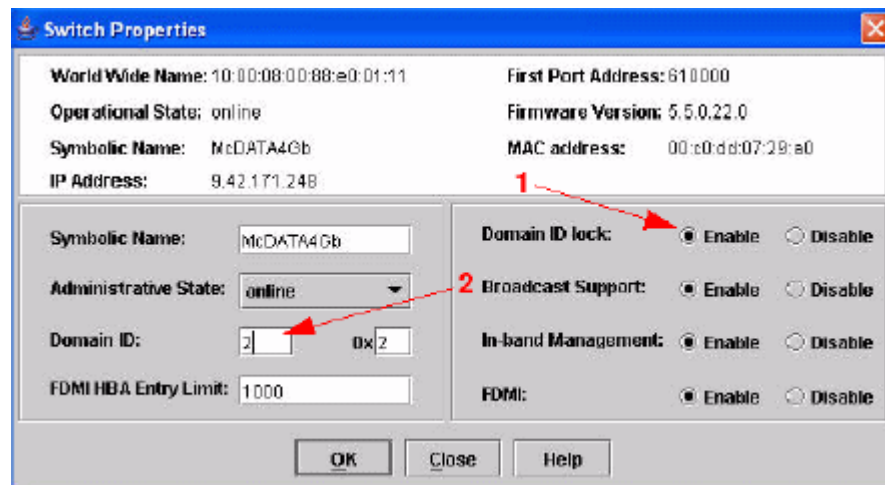


Figure 16:

(Picture Resources Legacy McData /IBM SAN solution – See Annexure – A for More Details)

*Note:* Domain ID lock on a McData BladeCenter switch is the same as the Insistent Domain ID setting on McData Sphereon switches and Intrepid Directors. This setting allows the switch to maintain its configured domain ID during a fabric merge or rebuild.

- i. Click OK when the Updating Switch Properties dialog box appears.
- ii. From the McData SANbrowser - Faceplate window, select your desired E\_Port. Next, select Port → Port Properties (Figure 17). We designated port 19. As you can see, the moment you turn on the internal blade server you will automatically log into the internal embedded switch. (Should you want to directly connect the EMC or any other OEM vendor storage, now is the time to start making zones/zonesets).

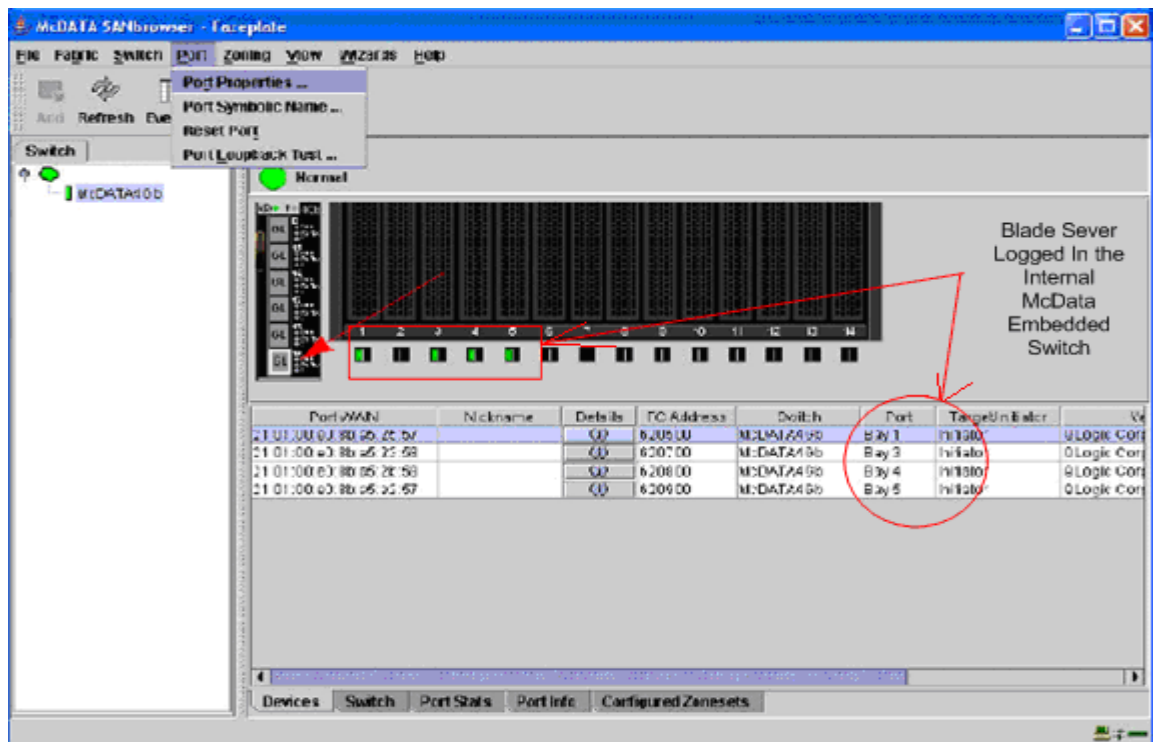


Figure 17: Blade Servers logged in the internal switch

(Picture (Modified) Resources Legacy McData /IBM SAN solution) – See Annexure – A for More Details)

- iii. From the Port Properties dialog box (Figure 18),  
Ensure the following settings are true for port 19:
  - a. Port State = Online
  - b. Port Speed = Auto-detect
  - c. Port Type = G-port (Generic Port)
  - d. I/O Stream Guard = Auto
  - e. Device Scan = EnableDd

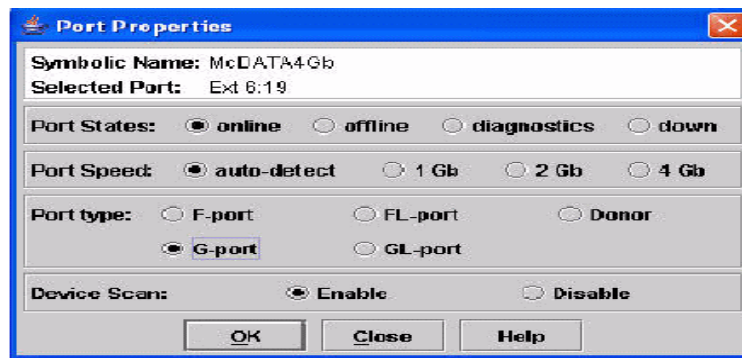



Figure 18: - Port Properties

(Picture Resources Legacy McData SAN solution) – See Annexure – A for More Details)

- iv. Click OK.
- v. From the Faceplate, select the Configured Zoneset tab (Figure 18). *Ensure there are no zones configured on the switch.* If there are zones configured, see the manufacturer user's guide for instructions on how to remove the zone set and zones. As per Best Practices, all zoning must be performed on the Edge Switch only.



*Figure 19: Confirming No Zones configured on Embedded Switch  
(Picture (Modified) Resources Legacy McData SAN solution – See Annexure – A  
for More Details)*

- vi. Close OK – Close the Window (proceed to the McData Edge Switch Configuration.)

### **McData Edge Switch Configuration**

Here we will cover only the most important connectivity options that must be verified before going live into the production environment.

1. Login into McData Edge Switch.
2. In the port configuration dialog window, you may optionally assign a name to the port and define the connection type. In this example, we assigned port 0 to an E\_Port type.

By assigning the port type to be specifically an E\_Port, no N\_Port are allowed to log in to the switch (see Figure 20.)

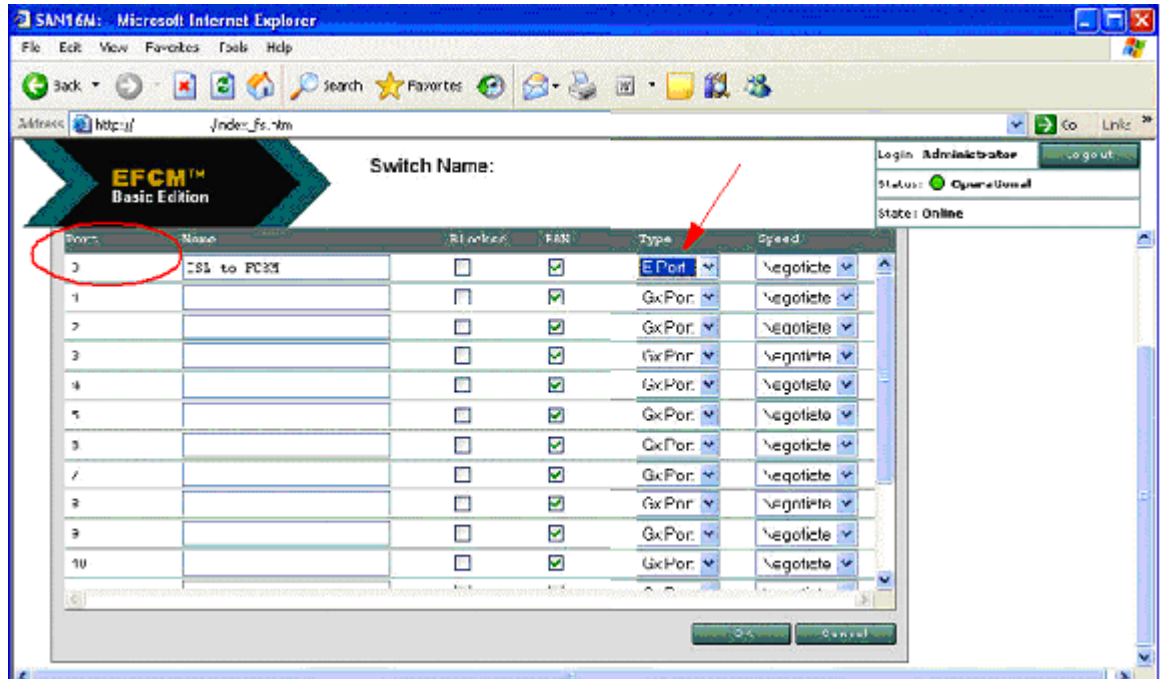


Figure 20: Confirming E\_Port-to-E\_Port connectivity

(Picture (Modified) Resources Legacy McData Online documentation – See Annexure – A for More Details)

3. Set the switch offline. Select Configure → Switch Online. Uncheck the check box. (Figure 21)

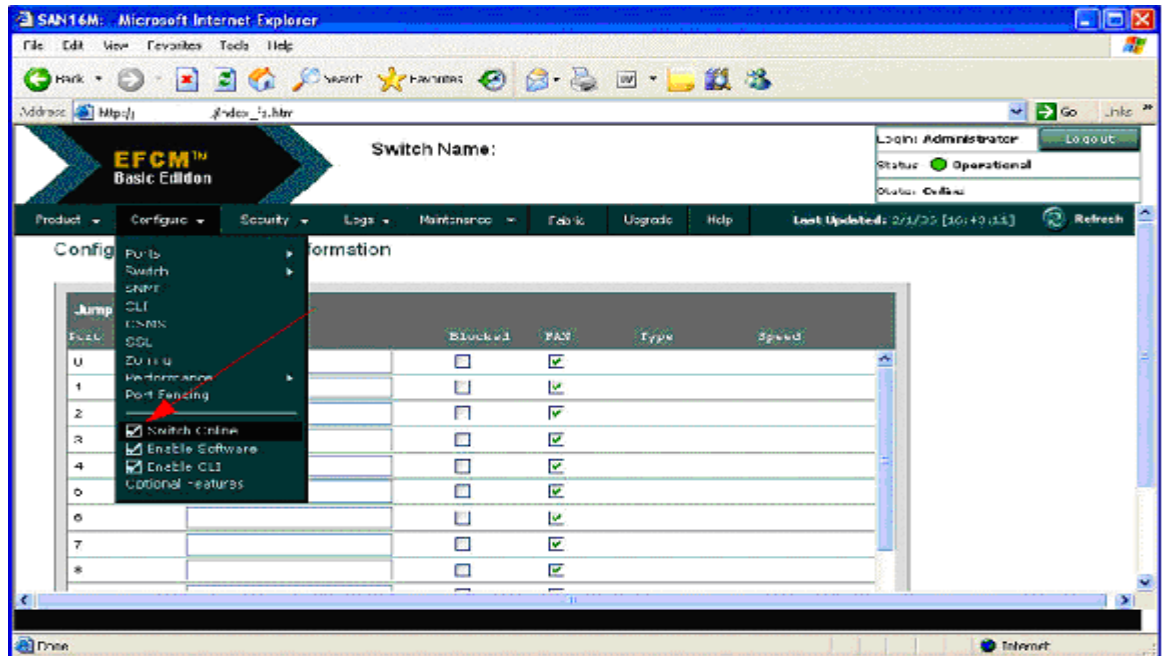


Figure – 21

(Picture Resources Legacy McData Online documentation – See Annexure – A for More Details)

4. From the Configure → Switch → Parameters settings window (Figure 22), select Insistent Domain ID (1). Next, define a unique ID ranging between 1 and 31 (2). In our example, we use Domain ID 3. Click OK.

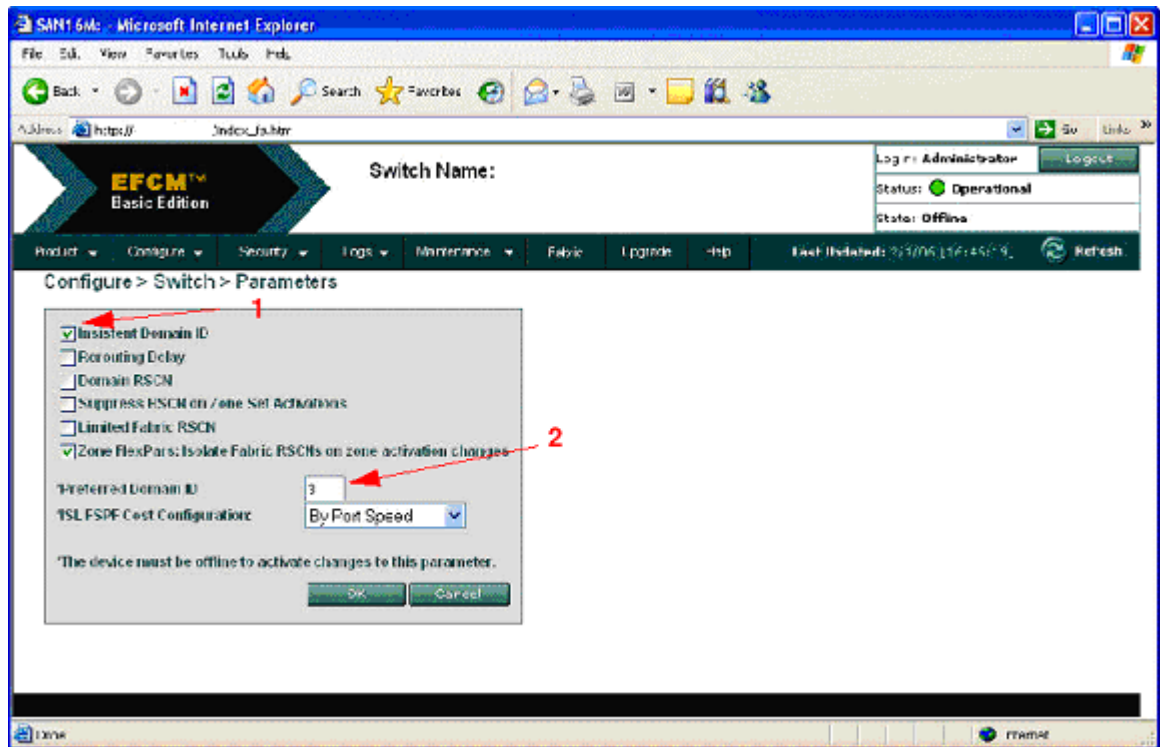


Figure 22: Domain ID Settings

(Picture (Modified) Resources Legacy McData Online documentation – See Annexure – A for More Details)

5. Go back and select Configure → Switch Online and check the check box.
6. You can now start connecting Storage units and should be ready to zone them with the Blade Server hosts.

### **3. ISL in a Multivendor Environment:**

#### **Brocade Embedded Switch (In Blade Server) <-ISL->**

#### **Cisco Edge Switch**


Though all parts of this article are significant, this part is the most interesting and most compelling topic in today's storage industry.

The concept of Interoperability will remain the same with either a Brocade Embedded switch (Embedded Switch 4010, 4016, 4020) or a regular Brocade Fabric Switch (like Brocade 4100 Switch).

*Note:* Embedded Switch can also function as in N\_Port ID Virtualization (NPIV) mode also called Access Gateway Mode - like Brocade 4020 Switch.

Brocade Access Gateway technology is the application of NPIV to an embedded Fibre Channel switch in an open systems blade server environment. Blade server switches almost always connect to an external switch. By applying NPIV to an embedded blade server switch, the external ports appear to another switch that supports NPIV to be N\_Ports, as if they were server HBAs. These external ports and the internal ports that connect to the blade servers are still switched in the same way as a traditional Fibre Channel switch. But an embedded switch configured as an Access Gateway device does not have external E\_Ports for uplink to another switch—instead the Access Gateway device has N\_Ports for uplinks. Figure 22 displays the interconnectivity between the embedded switch and the Edge Switch.(This is a new feature - no ISL is required and this will not be covered in this paper).





*Figure 23: Access Gateway interoperability with Edge Switch (McData -EOS or Cisco Fabric)  
(Picture (Modified) Resources Brocade Online white papers – See Annexure – A for More details)*

### Interop Mode and ISL Connectivity

Before every integration, there must be some kind of check list to refer to when it comes to the live production environment. This may be useful in implementations as interoperability is the facet of an implementation where multiple vendor products come in contact with each other. Fibre Channel standards have been put in place to guide vendors toward common external Fibre Channel interfaces

*Note:* Brocade is referred to as B Series and Cisco is referred to as C-Series in this section.

Keep the following details in mind to conduct a healthy and successful integration.

## **Check list Features and limitations:**

### **B-Series Checklist**

Before Integration, complete the following steps:

- Step B1** : verify switch firmware versions
- Steps B2, B3** : verify/configure switch domain IDs and verify switch/fabric default settings
- Step B4** : disable management server
- Step B5** : verify/configure fabric operating mode
- Step B6** : verify proper zoning configuration is in place

### **C-Series Checklist**

Before Integration, complete the following steps:

- Step C1** : verify switch firmware versions
- Step C2** : verify switch/fabric default settings
- Step C3** : verify the fabrics are in proper operating mode
- Step C4** : verify/configure switch domain IDs
- Step C5** : verify proper zoning configuration is in place

## **B-Series Features and Limitations:**

***Key bullet points/Best Practices along with the checklist must be followed strictly along with the EMC Support Matrix or Storage Vendor Support Matrix.***

When interoperability mode is set, the Brocade switch has the following limitations:

- All Brocade switches should be in Fabric OS 2.4 or later. (Refer to the EMC Support Matrix for latest FOS for Interop Mode).
- Brocade Embedded switch must be running in the Fabric Mode (No Access Gateway Mode is supported in multivendor environment).
- Interop mode affects the entire switch. All switches in the fabric must have interop mode enabled.
- Mspmgmtdeactivate must be run prior to connecting the Brocade switch to an MDS 9000 switch. This command uses Brocade proprietary frames to exchange platform information.
- If there are no zones defined in the effective configuration, the default behavior of the fabric is to allow no traffic to flow. If a device is not in a zone, it is isolated from other devices.
- Zoning can only be done with pWWNs. You cannot zone by port numbers or nWWNs.
- Domain IDs are restricted to the 97 to 127 range to accommodate Legacy McData (now Brocade) nominal restriction to this same range.
- Brocade WebTools will show an MDS 9000 switch as an anonymous switch. Only a zoning configuration of the MDS 9000 switch is possible.
- The full zone set (configuration) is distributed to all switches in the fabric.
- The following services are not valid on Brocade in Interop Mode:
  - Trunking (works within Brocade Series Switches)
    - Broadcast zones
    - Domain/port representation in zones
- The following services are not supported
  - \_Management Server
  - The Alias Server
  - \_ Secure Fabric OS

## **C-Series Features and Limitations:**

***Key bullet points/Best Practices along with the checklist must be followed strictly along with the EMC Support Matrix or Storage Vendor Support Matrix.***

The standard interoperability mode, which has been a fully functional feature since MDS SAN-OS Release 1.0(1) (refer to the latest EMC Support Matrix for the qualified IOS version), enables the MDS 9000 switch to interoperate with Brocade and McData switches when they are configured for interoperability. The standard interoperability mode allows the MDS 9000 switch to communicate over a standard set of protocols with these vendor switches.

When a VSAN is configured for the default interoperability mode, the MDS 9000 Family of switches is limited in the following areas when interoperating with non-MDS switches:

For Best Practices on the Cisco MDS Switch, keep the following bullet points in mind:

- Interop mode only affects the specified VSAN. The MDS 9000 switch can still operate with full functionality in other non-Interop mode VSANs. All switches that are part of the interoperable VSAN should have that VSAN set to interop mode, even if they do not have any end devices.
- Domain IDs are restricted to the 97 to 127 range, to accommodate Brocade/McData's nominal restriction to this same range.
- Domain IDs must be set up statically in production environments. The MDS 9000 switch will only accept one domain ID; if it does not get that domain ID, it isolates itself from the fabric.
- TE ports and PortChannels cannot be used to connect an MDS 9000 switch to a non-MDS switch.
- Only E\_ports can be used to connect an MDS 9000 switch to a non-MDS switch.
- Only the active zone set is distributed to other switches.

- If a Brocade switch issues a **cfgsave** command, the MDS 9000 switch rejects this vendor-specific command. The full zone database on the MDS 9000 switch is not updated.
- You must manually update the full zone database or copy the active zone set to the full zone database.
- The MDS 9000 switch still supports the following zoning limits per switch across all VSANs:
  - 2000 zones (as of SAN-OS 3.0, 8000 zones)

*Note:* Before configuring this number of zones in a mixed environment, determine the maximum number that can be supported by the other vendors present in the environment.

## Interoperability Modes:

### C-Series (Interop Modes)

The MDS 9000 Family of multilayer directors and fabric switches supports various interoperability types, depending on the release.

- Default or Native Mode**—This is the default mode or behavior for a VSAN that is communicating between a SAN composed entirely of MDS 9000 switches.

- Interop Mode 1**—This is the standard interoperability mode. It interoperates with Brocade and McData switches that have been configured for their own interoperability modes. Brocade and McData switches must be running in Interop mode to work with this VSAN mode.

- Interop Mode 2**—This mode, also known as legacy switch Interop mode 2, allows seamless integration with specific Brocade switches running in their own native mode of operation. Brocade switches must be configured with "core pid = 0" to work with this mode.

- Interop Mode 3**—Similar to Interop mode 2, interoperability mode 3 was introduced for Brocade switches that contained more than 16 ports. With this VSAN-based interop mode, you will not have to alter Brocade switches from their native mode (core pid = 1). They can be seamlessly added to a new or existing MDS SAN-OS VSAN. This mode is also known as legacy switch Interop mode 3.

- Interop Mode 4**—This mode, also known as legacy switch Interop mode 4, provides seamless integration between MDS VSANs and McData switches running in McData Fabric 1.0 Interop mode.

## Cisco Fibre Channel Features Affected by Interoperability

- Domain IDs

A switch may have to change its domain ID to the 97 to 127 range to accommodate the Brocade/McData 31 domain address limitation. If a domain ID is changed (which can be a disruptive event to the switch), all devices attached to the switch will need to log into the switch again. When a domain ID is changed, the switch itself will need to reregister with the principal switch in the fabric to verify domain ID uniqueness.

Disruptive: The impact of this event is switch-wide. Brocade and McData require the entire switch to be taken offline and/or rebooted when changing domain IDs.

Nondisruptive: This event is limited to the VSAN where the event is taking place. The MDS 9000 switch can perform this action, as the domain manager process for this VSAN is restarted and not the entire switch. This event still requires any devices logged into the VSAN on that switch to log in again to obtain a new FC ID.

- Fabric Shortest Path First (FSPF)
- Timers: All Fibre Channel timers must be the same on all switches as these values are exchanged by E\_ports when establishing an ISL. The timers are:
  - F\_S\_TOV (fabric stability time out value)
  - D\_S\_TOV (distributed services time out value)
  - E\_D\_TOV (error detect time out value)
  - R\_A\_TOV (resource allocation time out value)
- Trunking and PortChannels
- FC Aliases
- Default Zone Behavior
- Zoning Membership

## 1. B-Series (Interop Modes)

Brocade has two Interop Modes

- Interop Mode 0 (Native)
- Interop Mode 1

To ensure interoperability between Legacy Fabric OS v4.x-based products and Fabric OS v2.x- and v3.x-based products while maintaining compatibility with older firmware versions, a setting was created to enable the PID format to be set using either the new or old format. This is commonly known as the Core Switch PID Format setting.

### PID Format

Switches with fewer than 16 ports; the core PID will be set to 0

Switches with more than 16 ports; the core PID will be set to 1

## Interop Configuration Settings on Switches

Cisco MDS (Legacy Switch) Settings	Brocade Switch Settings
Set Interop Mode 2  Also set Interop Mode 2 if MDS switch is running IOS Release 2.1(2) and Legacy Models	Set (core pid = 0) Native Mode Interop Mode 1  On Brocade switches like 2400, 2800, 3200, and 3800 switches configured with core pid = 0 , Minimum version 3.1.1 and version 4.1.1
Set Interop Mode 3  Also set Interop Mode 2 if MDS switch is running IOS Release 2.1(2) and Legacy Models	Set (core pid = 1) (Core Mode) Interop Mode 1  On Switch 3900/12000 series with more than 16 ports.

Cisco MDS (new models) Settings	Brocade Switch Settings
Set Interop Mode 1  All new MDS switches with IOS version 3.0 and above.	Set (core pid = 1) (Default) Interop Mode 1  All Models

Please check the latest IOS version running on Cisco MDS switches in the latest EMC Support Matrix.



Please see the Example configuration with all parameters discussed in this Interop Mode section.

## **Example: Real Setup**

### **MDS 9000 Core with Brocade Edge Topology (on Embedded Blade Servers) Note:**

This is a Sample Configuration (Refer to the latest EMC Support Matrix for optimal performance).

#### **B-Series Checklist (Blade Server Embedded Switch)**

Before integration, complete the following steps:

**Step B1** : verify switch firmware versions

**Steps B2, B3** : verify/configure switch domain IDs and verify switch/fabric default settings

**Step B4** : disable management server

**Step B5** : verify/configure fabric operating mode

**Step B6** : verify proper zoning configuration is in place

#### **C-Series Checklist (Edge Switch)**

Before integration, complete the following steps:

**Step C1** : verify switch firmware versions

**Step C2** : verify switch/fabric default settings

**Step C3** : verify the fabrics are in proper operating mode

**Step C4** : verify/configure switch domain IDs

**Step C5** : verify proper zoning configuration is in place

### **B-Series configuration steps (Blade Server Switch)**

Refer to the EMC Support Matrix for the latest supported version on your OEM Blade Server - Dell, HP, IBM, etc. For simplicity, only CLI mode is covered here.

#### **Step B1: Verify switch firmware versions**

Login to the switch as an admin and verify firmware version using the telnet command "version". For example:.

Switch: login

Password: xxxxxxxx

Switch:admin>

Switch:admin> version

Kernel: 2.4.19

Fabric OS: v4.1.1

Made on: Wed Jun 18 02:59:09 2003

Flash: Thu Sep 4 19:36:06 2003

BootProm: 3.2.4

Switch:admin

#### **Steps B2, B3: Verify/configure switch domain IDs and verify switch/fabric default settings**

Ensure that all the switches between the two fabrics have unique domain IDs before they are ISL'd. List the domain ID addresses of each switch in both the fabrics and verify there are no duplicate IDs. If any duplicate addresses exist, change the IDs by assigning different domain numbers.

Here is an example on how to change domain IDs on B-series switches. Use a telnet session to perform the domain ID configurations. This requires disabling the switch temporarily, so plan accordingly.

```
Switch:admin> switchdisable
Switch:admin>
Switch:admin> configure
Configure...
    Fabric parameters (yes, y, no, n): [no] y
Domain: (1..239) [97] ← Select the domain ID in the range 97-127, if
change is
required
BB credit: (1..27) [16]
R_A_TOV: (4000..120000) [10000] ← Ensure this value is 10000,
should be
same for all switches in the fabric
E_D_TOV: (1000..5000) [2000] ← Ensure this value is 2000, should be
same for all
switches in the fabric
WAN_TOV: (1000..120000) [0]
Data field size: (256..2112) [2112]
Sequence Level Switching: (0..1) [0]
Disable Device Probing: (0..1) [0]
Suppress Class F Traffic: (0..1) [0]
SYNC IO mode: (0..1) [0]
VC Encoded Address Mode: (0..1) [0]
Core Switch PID Format: (0..1) [1] (default)
Per-frame Route Priority: (0..1) [0]
Long Distance Fabric: (0..1) [0]
Virtual Channel parameters (yes, y, no, n): [no]
Zoning Operation parameters (yes, y, no, n): [no]
RSCN Transmission Mode (yes, y, no, n): [no]
Arbitrated Loop parameters (yes, y, no, n): [no]
System services (yes, y, no, n): [no]
Portlog events enable (yes, y, no, n): [no]
No changes.
Switch:admin>
```

#### Step B4 : Disable management server

You must disable platform management services fabric-wide before enabling the interopmode and merging the fabrics. The following command will deactivate the Platform Database Management Service of each switch in the fabric.

```
Switch:admin> msPlMgmtDeactivate
This will erase all Platform entries. Are you sure?
(yes, y,
no, n): [no] y
Committing configuration...done.
Request Fabric to Deactivate Platform Management
services....
Done.
Switch:admin>
```

*Note* : You must reboot the switch after executing this command. However, you may wait to reboot the switch until completing the next step (B5), to avoid multiple rebooting.

### **Step B5: Verify/configure fabric operating mode**

For B-Series switches, the “interopmode 1” must be enabled before merging with C-series switches. This command enables interopmode on individual switches only and must be executed on all B-Series switches in the fabric. Use the following telnet command to change the operating mode.

```
login: admin
Password:xxxxxxx
Switch:admin> switchDisable
Switch:admin> 0x101a8dd0 (tThad): Jan 20 10:47:55
WARNING FW-STATUS_SWITCH, 3, Switch status changed
from HEALTHY/OK to Marginal/Warning
Switch:admin>
Switch:admin> interopmode 1
The switch effective configuration will be lost when
the operating mode is changed; do you want to
continue? (yes, y,
no, n): [no] y
Committing configuration...done.

cfgDisable: no EFFECTIVE configuration
interopMode is 1
NOTE: It is required that you boot this switch to make
this change take effect
Switch:admin> fastboot ← reboots the switch quicker,
bypassing POST
```

### **Step B6: Verify proper zoning configuration is in place**

Check for Duplicate Zone Sets (zones): To ensure proper zoning ISL and operation, verify there are no duplicate active ZoneSets or Zones across the two fabrics that need to be merged.

If duplicate zones exist, rename them, using the following telnet commands. See the B-series software manual for detailed explanation of these commands.

- **cfgShow**
- **zoneCreate**
- **cfgCreate**

Verify Proper Zone Naming

Ensure that zone names adhere to the following guidelines.

- All characters must be ASCII
- A name must be between 1 and 64 characters in length
- The first character of a name must be a letter. A letter is defined as either an uppercase [A-Z] or lower case [a-z] character
- Any character other than the first character must be a lower case character [a-z], an upper case character [A-Z], a number [0-9] or the symbol (\_).

Configure Zones using only PWWNs

In interoperability mode all zone members must be defined using port WWNs only. Defining them in any other way is not supported. For example, we can not define zones using FC port addresses or domain, port combinations etc.

### **C-Series configuration steps (EDGE Switch)**

The following steps provide information on how to verify and configure C-series switches for interoperability. While it is possible to accomplish this using either the Fabric Manager or the CLI, the following steps use only CLI (telnet).

#### **Step C1: Verify switch firmware versions (Refer to the EMC Support Matrix for the latest supported version with Blade Server OEM Models)**

```
MDS9509# show ver
Cisco Storage Area Networking Operating System (SAN-OS)
Software
TAC support: http://www.cisco.com/tac
Copyright (c) 1986-2002 by cisco Systems, Inc.
Software
kickstart: version 1.0(1) [build 1.0(0.260c)] [gdb]
system: version 1.0(1) [build 1.0(0.260c)] [gdb]
Hardware
RAM 1932864 kB
bootflash: 503808 blocks (block size 512b)
slot0: 0 blocks (block size 512b)
kickstart compile time: 11/7/2002 21:00:00
system compile time: 11/7/2002 20:00:00
```

### Step C2: Verify switch/fabric default settings

It may be necessary to change the Fibre Channel timers if they have been changed from the system defaults. The MDS 9000 and B-series FC Error Detect (ED\_TOV) and Resource Allocation (RA\_TOV) timers default to the same values. They can be changed, if needed. The RA\_TOV default is 10 seconds; the ED\_TOV default is 2 seconds. Per the FC-SW2 standard, these values must be the same on each switch within the fabric.

```
MDS9509 login: admin
Password: xxxxxxxx
MDS9509# show fctimer
F_S_TOV : 5000 milliseconds
D_S_TOV : 5000 milliseconds
E_D_TOV : 2000 milliseconds
R_A_TOV : 10000 milliseconds
```

To modify these values, use the following commands.

These changes can not be made unless all VSANs in the switch are suspended.

```
MDS9509# config t
MDS9509(config)# vsan database
MDS9509(config-vsan-db)# vsan 1 suspend
MDS9509# config t
MDS9509(config)# fctimer e_d_tov 2000
<1000-100000> E_D_TOV in milliseconds(1000-
100000)
MDS9509(config)# fctimer r_a_tov 10000
<5000-100000> R_A_TOV in milliseconds(5000-
100000)
```

### Step C3: Verify the fabrics are in proper operating mode

To enable interoperability mode on C-series switches, the first step is to place the VSAN of the E\_Ports(s) that connect to the B-series in interoperability mode.

```
MDS9509# config t
MDS9509(config)# vsan database
MDS9509(config-vsan-db)# vsan 1 interop
```

#### Step C4: Verify/configure switch domain IDs

The next step is to assign a domain ID in the range of 97 (0x61) through 127 (0x7F). While in interoperability mode, we are limited to a total of 31 switches in the fabric. In the MDS, the default is to request an ID from the principal switch. If the preferred keyword is used, the MDS will request a specific ID, but still join the fabric if the principal switch assigns a different ID. If the static keyword is used, the MDS will not join the fabric unless the principal switch agrees and assigns the requested ID.

Based on experience an under production environment static must be used to configure the domain ID.

```
MDS9509# config t  
MDS9509(config)# fcdomain domain 100 static  
                  vsan 1
```

When making changes to the domain, you may restart the MDS domain manager function for the altered VSAN. You may force a fabric reconfiguration with the *disruptive* keyword.

```
MDS9509(config)# fcdomain restart disruptive vsan 1  
Or not force a fabric reconfiguration  
MDS9509(config)# fcdomain restart vsan 1
```

#### Step C5: Verify proper zoning configuration is in place

Check for Duplicate Zone Sets (zones):

To ensure proper zoning merge and operation, verify there are no duplicate active ZoneSets or Zones across the two fabrics that need to be merged. If any duplicate zones exist, rename them. You can verify zoning information by using the “show zone” command.

For example, to verify all active zones:

```
MDS9509# show zoneset active  
zoneset name mdscore vsan 1  
zone name vz1 vsan 1  
* fcid 0x630500 [pwwn 50:06:01:60:88:02:90:cb]  
* fcid 0x610400 [pwwn 10:00:00:00:c9:24:3d:90]  
zone name vz2 vsan 1  
* fcid 0x630400 [pwwn 10:00:00:00:c9:24:3f:75]  
* fcid 0x6514e2 [pwwn 21:00:00:20:37:a7:ca:b7]  
* fcid 0x6514e4 [pwwn 21:00:00:20:37:a7:c7:e0]
```

```
* fcid 0x6514e8 [pwwn 21:00:00:20:37:a7:c7:df]
zone name vz3 vsan 1
* fcid 0x651500 [pwwn 10:00:00:e0:69:f0:43:9f]
* fcid 0x6105dc [pwwn 21:00:00:20:37:28:31:6d]
* fcid 0x6105e0 [pwwn 21:00:00:20:37:28:24:7b]
* fcid 0x6105e1 [pwwn 21:00:00:20:37:28:22:ea]
* fcid 0x6105e2 [pwwn 21:00:00:20:37:28:2e:65]
* fcid 0x6105e4 [pwwn 21:00:00:20:37:28:26:0d]
```

```
zone name $default_zone$ vsan 1
```

### **Configure Zones using only PWWNs**

As explained in the limitations section, in interoperability mode all zone members must be defined using port WWNs only. Defining them any other way is not supported. For example, we can not define zones using FC port addresses or domain, port combinations, etc. If there are any zones not defined as stated above, redefine them using proper PWWNs.



## Chapter 2

# Cisco MDS Fabric VSAN, Domain and FSPF Troubleshooting

Cisco is among the most recognized names in the networking industry. During recent years, due to their entry in the Storage Area Network market, they are now a leading competitor in the storage industry, too.

Here the focus is not only on troubleshooting but also on best practices. Implementing Best Practices proactively eliminates or reduces most troubleshooting.

As you know, VSAN is a logical way of isolating devices that are physically connected to the same storage network, but are logically considered to be part of different SAN fabrics that do not need to be aware of one another.

This chapter is divided into three parts:

Part A: Troubleshooting VSAN issues (Page 58 – Page 61).

Part B: Troubleshooting Domain ID Issues (Page 62 – Page 65).

Part C: Troubleshooting FSPF Issues (Page 66 – 75)

Let's first look at the implementation **Key Points/Best Practices of VSAN/Domain ID Assignments and FSPF and Checklist.**

***Key Points/Best Practices for VSAN Implementation:***

- Avoid using VSAN 1 (the default VSAN) for production network traffic; create at least one VSAN to carry your network traffic.
- Isolate devices in VSANs whenever practical.
- Leave fabric timers and FSPF timers at their default settings, unless changes are required because of interoperability with an existing fabric or long-haul links are being deployed.
- Use Inter-VSAN routing (IVR) only when necessary to selectively connect devices across VSANs.
- If IVR is used without NAT, ensure that domain IDs are statically configured and unique across all VSANs.
- Keep FCIP gateways in their own native VSAN to isolate disturbances when problems in the IP cloud (such as flapping links) occur. Allow only limited VSAN over FCIP, for example, Storage Systems that will participate in SRDF/Replications.
- Use VSAN-based roles to control and limit management access to your switches.
- We recommend using only the following characters in a VSAN name:
  - a-z or A-Z
  - 0 - 9
  - - (hyphen) or \_ (underscore)

### ***Key Points/Best Practices for Domain ID Assignment:***

- Use static domains in most environments.
- Disable the Domain Manager to disable the principal switch selection process.

(This is possible if all domains are statically assigned. Disabling the principal switch selection can reduce disruption when switches are rebooted or added to the fabric. This must be done on each switch that should not participate in principal switch selection. A disruptive restart of the fabric is required to apply this change.

To disable the Domain Manager, choose Fabricxx > All VSANs > Domain Manager and uncheck the Enable check box in Fabric Manager or use the `no fcdomain vsan x` CLI command.

- Keep domain ID allowed lists the same on all switches in a fabric for consistency.

If the principal switch changes, the allowed domain lists will remain the same.

- Interop Mode: Assign domain IDs between decimal 97 and 127 if the domain may be used for standards-based Interop mode.
- Do not perform frequent changes to the Domain Manager on production fabrics.
- Always save the Running Configuration to startup config if sure about the changes performed on a switch in running mode.
- Enable reconfigure fabric (RCF) rejection on every ISL port if high availability is mandatory.

### ***Key Points/Best Practices for FSPF***

- Use the default FSPF link cost, which can be configured on a per-VSAN basis for the same physical link, to provide preferred and alternate paths.  
(If you must alter the FSPF link cost, use caution to avoid asymmetric Fibre Channel routing.)
- Use the default FSPF load-balancing configuration unless you must load balance based on your unique fabric; for example, if you have FICON VSANs.
- Use the default FSPF timer configuration. If FSPF timers are misconfigured, then the switches will not reach the “two-way” state and FSPF will not operate properly.

### **Basic Checklist:**

1. Ensure the domain parameters for switches in the VSAN.
2. Ensure the physical connectivity for any problem ports or VSANs.
3. Ensure that your source (initiators) and target (storage) devices are in the same server.
4. Ensure that your source (initiators) and target (storage) devices are in the same VSAN.
5. Ensure that your source (initiators) and target (storage) devices are in the same zone.
6. Ensure that the zone is part of the active zone set.
7. Ensure the FSPF parameters for switches in the VSAN.

Most VSAN problems can be avoided by following the best practices for VSAN implementation. If needed, use the Fabric Analysis tool in Fabric Manager to verify different categories of problems such as VSANs, zoning, FCdomain, admin issues, or switch-specific or fabric-specific issues.

When suspending or deleting VSANs, do so one VSAN at a time, and wait a minimum of 60 seconds after you issue the `vsan suspend` command before you issue any other config command. Failure to do so may result in some Fibre Channel interfaces or member ports in a PortChannel becoming suspended or error-disabled.

## **Tools Available in Fabric Manager and CLI (For VSAN, FC domain, FSPF, and zone)**

### **Fabric Manager**

- Fabricxx > VSANxx to view the VSAN configuration in the Information pane.
- Fabricxx > VSANxx and select the Host or Storage tab in the Information pane to view the VSAN members.
- Fabricxx > VSANxx > Domain Manager to view the FC domain configuration in the Information pane.
- Fabricxx > VSANxx > FSPF to view the FSPF configuration in the Information pane.
- Fabricxx > VSANxx > zoneset-name to view the zone configuration for this VSAN. Zone configuration problems may appear to be a VSAN problem.

### **CLI**

- `show vsan`
- `show fcdomain`
- `show fspf`
- `show fspf internal route vsan vsan-id`
- `show fcns database vsan vsan-id`
- `show zoneset name zoneset-name vsan vsan-id`
- `show zoneset active vsan vsan-id`
- `show zone vsan vsan-id`
- `show zone status show vsan vsan-range`

*Note:* Most users are not familiar with the diverse command line capabilities available in the Cisco MDS switches, so I will be using Fabric Manager to troubleshoot. I will also show command line capability where necessary. For complete details, refer to the Cisco MDS Command Line Guide.

## Part A

### Troubleshooting VSAN

Common issues in VSAN are as follows:

- Host Cannot Communicate with Storage
- E\_Port Is Isolated in a VSAN
- Troubleshooting Interop Mode Issues

Verifying VSAN Membership Using Fabric Manager

#### Troubleshooting by Fabric Manager

- **Step 1→** Verify that both devices are in the same VSAN.  
Choose Fabricxx > VSANxx and select the Host or Storage tab in the Information pane.
- **Step 2→** If the host and storage are in different VSANs, verify which port is not in the correct VSAN and then follow these steps to change the port VSAN:
  1. Highlight the host or storage in the Information pane. You see the link to that end device highlighted in blue in the map pane.
  2. Right-click on the highlighted link and select Interface Attributes from the pop-up menu.
  3. Set the PortVSAN field to the VSAN that holds the other end device and click Apply Changes.

- **Step 3→** Right-click any ISL between the switches and select Interface Attributes. Select the Trunk Config tab and verify that the allowed VSAN list includes the VSAN found in Step 1.
- **Step 4→** If the trunk is not configured for the VSAN, set the allowed VSANs field to include the VSAN that the host and storage devices are on and click Apply Changes.

### Troubleshooting by Command Line

- **Step 1→** Use the show vsan membership command to see all the ports connected to your host and storage, and verify that both devices are in the same VSAN.
- **Step 2→** If the host and storage are in different VSANs, use the vsan database vsan vsan-id interface command to move the interface connected to the host and storage devices into the same VSAN.
- **Step 3→** Use the show interface command to verify that the trunks connecting the end switches are configured to transport the VSAN.
- **Step 4→** If the trunk is not configured for the VSAN, use the interface command and then the switchport trunk-allowed vsan command in interface mode to add the VSAN to the allowed VSAN list for the interface that connects the host and storage devices.

## Troubleshooting Isolated E\_Port Using Fabric Manager

*Cisco MDS makes E\_Port connection between Cisco ↔ Non-Cisco Switch*

### Troubleshooting by Fabric Manager

To resolve VSAN isolation on an E\_Port

- **Step 1→** Choose Switches > Interfaces > FC Physical and check the FailureCause column on the E\_Port to verify that you have a VSAN mismatch.
- **Step 2→** Choose Switches > Interfaces > FC Physical and use the PortVSAN field to correct a VSAN mismatch.

### Troubleshooting by Command Line

- **Step 1→** Use the show interface command to verify that the port is isolated because of a VSAN mismatch.  

```
switch# show interface fc2/4
```

fc2/4 is down fc2/4 is down (isolation due to port vsan mismatch)
- **Step 2→** Use the show vsan membership command to verify that the ports are in separate VSANs.  

```
switch# show vsan membership
```
- **Step 3→** Use the vsan database vsan vsan-id interface command to move the ports into the same VSAN



### **Troubleshooting Isolated TE Port Using Fabric Manager (Only)**

*Cisco MDS makes TE Port connection between Cisco ⇔ Cisco Only TE. ISL has many advantages over E\_Port which is an ISL between Cisco and Non-Cisco Switch. Refer to the Cisco MDS Fabric Manager Guide for details.*

### **Troubleshooting by Fabric Manager**

(For CLI, please refer to Cisco MDS command line guide)

- **Step 1→** Choose Switches > Interfaces > FC Physical and check the FailureCause column on the TE port to verify a trunk problem.
- **Step 2→** Choose Switches > Interfaces > FC Physical and select the Trunk Failures tab to determine the reason for the trunk problem.
- **Step 3→** Correct the problem listed in the FailureCause column  
Choose Switches > Interfaces > FC Physical and use the PortVSAN field to correct the VSAN misconfiguration problems.
- **Step 4→** Repeat this procedure for all isolated VSANs on this TE port.

### **Troubleshooting Isolated TE Port (Timers) Using Fabric Manager (only)**

If the timer has been adjusted to set an ISL between Cisco and non-Cisco switch and you are experiencing problems and must troubleshoot, first refer to the EMC support Matrix and the OEM Switch support Matrix along with the Cisco Supported Values. In production, open a case with the storage vendor to avoid delay, then follow these steps:

- **Step 1→** Use the show fctimer command to verify that the fabric timers are inconsistent across the VSANs.
- **Step 2→** Use the fctimer distribute command to enable CFS distribution for the fabric timers. Repeat this on all Switches in this VSAN.
- **Step 3→** Use the fctimer command to set each timer.
- **Step 4→** Use the fctimer commit command to save these changes and distribute them to all switches in the VSAN.

## Part B

### Troubleshooting Domain ID Issues

The full concept and details are beyond the scope of this article; please refer to the Cisco Fabric Manager Guide for complete details. Here, the Domain ID troubleshooting methods will be discussed in Fabric Mode only.

#### I. Domain ID Conflicts

The principal switch assigns domain IDs when a new switch is added to an existing fabric. However, when two fabrics merge, the principal switch selection process determines which one of the preexisting switches becomes the principal switch for the merged fabric. The election of the new principal switch is characterized by the following rules:

- A switch with a populated domain ID list takes priority over a switch which has an empty domain ID list. The principal switch becomes the one in the fabric with the populated domain ID list.
- If both fabrics have a domain ID list, the priority between the two principal switches is determined by the configured switch priority. This is a user-settable parameter. The lower the value, the higher the priority.
- If the principal switch cannot be determined by the two previous criteria, the principal switch is then determined by the WWNs of the two switches. The lower the value of the WWN, the higher the switch priority.
- Two switch fabrics might not merge. If two fabrics with two or more switches are connected, and they have at least one assigned domain ID in common, and the auto-reconfigure option is disabled (this option is disabled by default), then the E\_Ports that are used to connect the two fabrics will be isolated due to domain ID overlap.

## II. Switch is not visible to Other Switches in a VSAN

This could be because of the following reasons.

### ***FC Domain ID Overlap***

To resolve an FC domain ID overlap, you can either change the overlapping static domain ID by manually configuring a new static domain ID for the isolated switch, or disable the static domain assignment and allow the switch to request a new domain ID after a fabric reconfiguration.

### **Using Fabric Manager - Assign New Domain ID**

Note: All devices attached to the switch in the VSAN get a new FC ID when a new domain ID is assigned. Some hosts or storage devices may not function as expected if the FC ID of the host or storage device changes.

You may see the following system message in the message log when a domain ID overlap occurs:

Error Message PORT-5-IF\_DOWN\_DOMAIN\_OVERLAP\_ISOLATION: Interface [chars] is down  
(Isolation due to domain overlap)

Explanation: The interface is isolated because of a domain overlap.

Recommended Action: Use the show fcdomain domain-list to determine which domain IDs are overlapping. Use the fcdomain domain domain-id [static | preferred] vsan vsan-id CLI command or similar Fabric Manager procedure to change the domain ID for one of the overlapping domain IDs.

- **Step 1**→ Choose Switches > Interfaces > FC Logical and check the FailureCause column for an isolation or domain overlap status.
- **Step 2**→ Choose Fabricxx > VSANxx > Domain Manager to view which domains are currently in the VSAN.

- **Step 3→** Repeat Step 2 on the other switch to determine which domain IDs overlap.
- **Step 4→** Select the Configuration tab and set Config Domain and Config Type to change the domain ID for one of the overlapping domain IDs.
  - The static option tells the switch to request that particular domain ID. If it does not get that particular address, it will isolate itself from the fabric.
  - The preferred option has the switch request a specified domain ID. If that ID is unavailable, it will accept another ID.
- **Step 5→** Set the Restart drop-down menu to Disruptive and click Apply Changes to restart the Domain Manager.

### By Using CLI - Assign New Domain ID

- **Step 1→** The following example output shows the isolation error message.

**Switch# show interface fc2/14**

fc2/14 is down (Isolation due to domain overlap)

Hardware is Fibre Channel, WWN is 20:4e:00:05:30:00:63:9e

vsan is 2

Beacon is turned off

192 frames input, 3986 bytes, 0 discards

0 runts, 0 jabber, 0 too long, 0 too short

0 input errors, 0 CRC, 3 invalid transmission words

0 address id, 0 delimiter

0 EOF abort, 0 fragmented, 0 unknown class

231 frames output, 3709 bytes, 16777216 discards

Received 28 OLS, 19 LRR, 16 NOS, 48 loop inits

Transmitted 62 OLS, 22 LRR, 25 NOS, 30 loop inits

- **Step 2→** Use the `show fcdomain domain-list vsan vsan-id` command to view which domains are currently in your fabric

```
switch1# show fcdomain domain-list vsan 2
```

```
Number of domains: 2
```

```
Domain ID WWN
```

```
-----
```

```
0x4a(74) 20:01:00:05:30:00:13:9f [Local]
```

```
0x4b(75) 20:01:00:05:30:00:13:9e [Principal]
```

- **Step 3→** Repeat Step 2 on the other switch to determine which domain IDs overlap.

```
switch2# show fcdomain domain-list vsan 2
```

```
Number of domains: 1
```

```
Domain ID WWN
```

```
-----
```

```
0x4b(75) 20:01:00:05:30:00:13:9e [Local][Principal]
```

```
-----
```

In this example, switch 2 is isolated because of a domain ID 75 overlap.

- **Step 4→** Use the `fcdomain domain domain-id [static | preferred] vsan vsan-id` command to change the domain ID for one of the overlapping domain IDs. (Based of experience, use static in production environments and restart the VSAN as follows)

```
fcdomain domain 101 static vsan 1 (for setting Domain ID)
```

```
fcdomain restart vsan 1 (for refreshing vsan)
```

For complete details of command line, refer to the Cisco Command Line guide available at [cisco.com](http://cisco.com).

## Part C

### Troubleshooting Fabric Shortest Path First (FSPF)

You can achieve higher utilization with greater traffic control. FSPF is the service that can be independently configured per VSAN. Within each VSAN topology, FSPF can be configured to provide a unique routing configuration and resulting traffic flow. Here we have to use command line to determine what is happening at the protocol level.

For Fabric Manager, refer the Cisco MDS Fabric Manager Guide available at [cisco.com](http://cisco.com). A Fabric Manager screen shot is provided at the end of this section, depicting where you can modify the settings for FSPF.

*Note:* The FSPF settings are always recommended kept at default level. Refer to the Cisco MDS Fabric Manager Guide for details.

FSPF is the protocol currently standardized by the T11 committee for routing in Fibre Channel networks. The FSPF protocol has the following characteristics and features:

- Supports multipath routing.
  - Bases path status on a link state protocol.
  - Routes hop by hop, based only on the domain ID.
  - Runs only on E\_ports or TE ports and provides a loop free topology.
  - Runs on a per VSAN basis. Connectivity in a given VSAN in a fabric is guaranteed only for the switches configured in that VSAN.
  - Uses a topology database to keep track of link state on all switches in the fabric and associates a cost with each link.
  - Guarantees a fast reconvergence time in case of a topology change.
- Uses the standard Dijkstra's algorithm, but there is a static dynamic option for a more robust, efficient, and incremental Dijkstra's algorithm. The reconvergence time is fast and efficient as the route computation is done on a per VSAN basis.

## FSPF Global Configuration

By default, FSPF is enabled on switches in the Cisco MDS 9000 Family. Some FSPF features can be globally configured in each VSAN. By configuring a feature for the entire VSAN, you do not have to specify the VSAN number for every command. This global configuration feature also reduces the chance of typing errors or other minor configuration errors.

*Note:* If you are operating with other vendors using the backbone region, change this default to be compatible with those settings. FSPF is enabled by default. Generally, you do not need to configure these advanced features until it essential to do so.

The following table lists the default settings.

LSR Option	Default	Description
Acknowledgment interval (RxmtInterval)	5 seconds	The time a switch waits for an acknowledgment from the LSR before retransmission.
Refresh time (LSRefreshTime)	30 minutes	The time a switch waits before sending an LSR refresh transmission.
Maximum age (MaxAge)	60 minutes	The time a switch waits before dropping the LSR from the database.

*LSR Default Settings*

*Reference – Cisco Troubleshooting guide – See Annexure – A for more Details*

Each time a new switch enters the fabric, a link state record (LSR) is sent to the neighboring switches, and then flooded throughout the fabric.

FSPF tracks the state of links on all switches in the fabric, associates a cost with each link in its database, and then chooses the path with a minimal cost. The cost associated with an interface can be administratively changed to implement the FSPF route selection. The integer value to specify cost can range from 1 to 65,535. The default cost for 1 Gbps is 1000 and for 2 Gbps is 500.

The Hello Time intervals and the Dead Time intervals are important to watch during troubleshooting. In particular, Dead Time interval must be the same in the ports at both ends of the ISL.

Only administrators who have extensive experience in FSPF on multivendor environments should adjust FSPF parameters.

To troubleshoot FSPF using the CLI, follow these steps:

**Step 1**→ Use the **show fspf database vsan** command to verify that each path is in the FSPF database.

▪ **switch1# show fspf database**

FSPF Link State Database for VSAN 2 Domain 1 <-----1

LSR Type = 1

Advertising domain ID = 1 <-----2

LSR Age = 81 <-----3

LSR Incarnation number = 0x80000098 <-----4

LSR Checksum = 0x2cd3

Number of links = 2

NbrDomainId IfIndex NbrIfIndex Link Type Cost

-----  
237 0x00010002 0x00010001 1 1000 <-----5

238 0x00010003 0x00010002 1 1000 <-----6

FSPF Link State Database for VSAN 2 Domain 237 <-----LSR  
**for another switch**

LSR Type = 1

Advertising domain ID = 237 <-----7

LSR Age = 185

LSR Incarnation number = 0x8000000c

LSR Checksum = 0xe0a2

Number of links = 2

NbrDomainId IfIndex NbrIfIndex Link Type Cost



239 0x00010000 0x00010003 1 1000 <-----8  
 1 0x00010001 0x00010002 1 1000 <-----9  
 FSPF Link State Database for VSAN 2 Domain 238 <-----LSR  
 for another switch  
 LSR Type = 1  
 Advertising domain ID = 238  
 LSR Age = 1052

LSR Incarnation number = 0x80000013

LSR Checksum = 0xe294

Number of links = 2

NbrDomainId IfIndex NbrIfIndex Link Type Cost

-----  
 239 0x00010003 0x00010001 1 1000

1 0x00010002 0x00010003 1 1000

FSPF Link State Database for VSAN 2 Domain 239 <-----LSR

for another switch

LSR Type = 1

Advertising domain ID = 239

LSR Age = 1061

LSR Incarnation number = 0x80000086

LSR Checksum = 0x66ac

Number of links = 4

NbrDomainId IfIndex NbrIfIndex Link Type Cost

-----  
 237 0x00010003 0x00010000 1 1000

238 0x00010001 0x00010003 1 1000

1. The domain 1 view of the fabric topology.
2. Domain 1 is owner of the LSR.
3. This is a 16-bit counter starting at 0x0000, incremented by one for each switch during flooding and by one for each second held in the database. This field is used as a tie-breaker if incarnation numbers are the same.

4. This is a 32-bit value between 0x80000001 and 0x7FFFFFFF which is incremented by one each time the originating switch transmits an LSR. This is used before LSR Age.
5. The path to domainID 237, switch 1.
6. The path to domain ID 238, switch 5.
7. Switch 1, domain ID 237 is the owner.
8. The path to domain ID 239, switch 3.
9. The path to domain ID 1, switch 2.

**Step 2→** Use the show fspf vsan vsan-id interface command to verify that the FSPF parameters are correct for each interface and verify that the interface is in the FSPF active state.

**switch1# show fspf vsan 2 interface fc1/2**

FSPF interface fc1/2 in VSAN 2

FSPF routing administrative state is active <-----1

Interface cost is 1000 <-----2

Timer intervals configured, Hello 20 s, Dead 80 s, Retransmit 5 s <-----3

FSPF State is FULL <-----4

Neighbor Domain ID is 1, Neighbor Interface index is 0x00010002 <-----5

Statistics counters :

Number of packets received : LSU 46 LSA 24 Hello 103 Error packets 0

Number of packets transmitted : LSU 24 LSA 45 Hello 104 Retransmitted  
LSU 0

Number of times inactivity timer expired for the interface = 0

This displays the # of packets; Hellos should be received every 20 Sec.

1. FSPF routing is active.
2. The cost of the path out of this interface.
3. The configured FSPF timers for this interface, which must match on both sides.
4. Either Full State or Adjacent. Sent and received all database exchanges and required ACKs. Port is now ready to route frames.
5. FSPF neighbor information.

**Step 3→** Use the show fspf internal route vsan command to verify that all Fibre Channel routes are available.

### Show fspf internal route vsan 2

```
switch1# show fspf internal route vsan 2
FSPF Unicast Routes
```

VSAN	Number	Dest Domain	Route Cost	Next hops
1		0x01 (1)	1000	fc1/2
1		0xEF (239)	1000	fc1/1
1		0xED (238)	2000	fc1/1
				fc1/2

*Output Reference – Cisco Command line guide – See Annexure – A for more details*

This shows the total cost of all links.

The next hop (238) has two interfaces. This indicates that both paths will be used during load sharing. Up to sixteen paths can be used by FSPF with a Cisco MDS 9000 Family switch.

With the implementation of VSANs used with Cisco MDS 9000 Family switches, a separate instance of FSPF runs within each VSAN, and each instance is independent of the others. For this reason, FSPF issues affecting one VSAN have no effect on FSPF running in other VSANs.

### Wrong Hello Interval on an ISL Using the CLI (Loss of Two-Way Communication)

The switches will not reach the “two-way” state if FSPF is misconfigured.

- If FSPF removes the Inter-Switch Link (ISL) from the topology database.
- New link state records (LSRs) are flooded to adjacent switches to notify them that the FSPF database has changed.
- The port enters Init state and removes its neighbor’s domain ID from the Recipient Domain ID field and inserts 0xFFFFFFFF.

To resolve a wrong hello interval on an ISL using the CLI, follow these steps

**Step 1→** Use the debug fspf all command and look for wrong hello interval messages.

```
switch1# debug fspf all
```

```
Jan 5 00:28:14 fspf: Wrong hello interval for packet on interface 100f000 in VSAN 1
```

```
Jan 5 00:28:14 fspf: Error in processing hello packet , error code = 4
```

**Step 2→** Use the undebug all command to turn off debugging.

**Step 3→** Use the show fspf internal route vsan command to show FSPF information.

```
switch1# show fspf internal route vsan 1
FSPF Unicast Routes
```

VSAN Number	Dest Domain	Route Cost	Next hops
1	0xEF(239)	1000	fc1/1 <-----1
1	0xED(238)	2000	fc1/1
1	0x01(1)	3000	fc1/1 <-----2

*Output Reference – Cisco Command line guide – See Annexure – A for more details*

- 1. There is no second path to domain 238, through domain 1, switch 2.
- 2. There is no direct path to domain 1, switch 2; traffic must travel through three ISLs. This is based on the route cost column.

**Step 4→** Use the show fspf vsan vsan-id interface command to view the FSPF configuration.

```
switch1# show fspf vsan 1 interface fc1/16
```

```
FSPF interface fc1/16 in VSAN 1
```

```
FSPF routing administrative state is active
```

```
Interface cost is 500
```

```
Timer intervals configured, Hello 5 s, Dead 80 s, Retransmit 5 s <-----1
```

```
FSPF State is INIT <-----2
```

Statistics counters :

Number of packets received: LSU 0 LSA 0 Hello 2 Error packets 1

Number of packets transmitted: LSU 0 LSA 0 Hello 4 Retransmitted  
LSU 0

Number of times inactivity timer expired for the interface = 0

1. The Hello timer is not set to the default, so you should check the neighbor configuration to make sure it matches.
2. FSPF is not in FULL state, indicating a problem.

**Step 5→** Repeat Step 4 to determine the value of the Hello timer on the adjacent switch.

switch2# show fspf v 1 interface fc2/16

FSPF interface fc2/16 in VSAN 1

FSPF routing administrative state is active

Interface cost is 500

Timer intervals configured, Hello 20 s, Dead 80 s, Retransmit 5 s <---

--1

FSPF State is INIT <-----2

Statistics counters:

Number of packets received: LSU 0 LSA 0 Hello 2 Error packets 1

Number of packets transmitted: LSU 0 LSA 0 Hello 4 Retransmitted  
LSU 0

Number of times inactivity timer expired for the interface = 0

1. The neighbor FSPF Hello interval is set to the default (20 seconds).
2. FSPF is not in FULL state, indicating a problem.

**Step 6→** Use the interface command and then the fspf hello-interval command in interface mode to change the default Hello interval.

### The Final debug command:

If you want to screen every message for troubleshooting:

Use the **debug fspf all** command and look for **nonexistent region messages**.

This command must be used if you are fully aware of the output messages and have extensive knowledge of the Cisco MDS switch.

```
switch1# debug fspf all
Jan 5 00:39:31 fspf: FC2 packet received for non existent region 0 in VSAN 1 <-----1
Jan 5 00:39:33 fspf: FC2 packet received for non existent region 0 in VSAN 1
Jan 5 00:39:45 fspf: Interface fc1/1 in VSAN 1 : Event INACTIVITY , State change INIT ->
INIT
Jan 5 00:39:45 fspf: Interface fc1/2 in VSAN 1 : Event INACTIVITY , State change INIT ->
INIT <-----2
1. The neighbor switch advertising region is 0.
2. FSPF is in init state for each ISL.
```


1. The neighbor switch advertising region is 0.
2. FSPF is in init state for each ISL.

You should open a second Telnet or SSH session before entering any debug commands. If the debug output overwhelms the current session, you can use the second session to enter the **undebug all** command to stop the debug message output.

Use the **undebug all** command to turn off debugging.

### Fabric Manager - FSPF Options (For Parameters settings)

Use the following screen shot for reference purposes only. It displays the standard setting in the Cisco MDS environment. Refer to the Switch documentation when connecting MDS switch to a different vendor fabric (Brocade, McData or Qlogic, etc.), otherwise, where possible, keep it default.



Picture (Modified) Reference Online – Cisco Advanced Design Guide – See Annexure – A for more details.

1. Select Fabric

2. Expand & Click Below On FC Logical

3. Click FSPF

4. Change Values Here

Sample - Edge - Core Fabric Design

Switch	VSIAN Id, Interface	Set To Default	Cost	Admin Status	Hello Interval	Dead Interval	ReTx Interval	Neighbor State	Neighbor Domain	Neighbor PortIndex	CreateTime
100	channel16		125	up	20	80	5	Full	0x10(16)	0x4000F	2008/01/26-21:37:02
100	channel16		125	up	20	80	5	Full	0x10(16)	0x4000F	2008/01/26-21:37:02
100	channel16		125	up	20	80	5	Full	0xa(10)	0x4000F	2008/01/26-22:24:02
100	channel16		125	up	20	80	5	Full	0xa(10)	0x4000F	2008/01/26-22:24:02

Pictures (Modified) - Reference Online – Cisco Advanced Design Guide – See Annexure – A for more details.

## APPENDIX - A

This appendix lists the following documentation help that you can access on the Web and other sources.

- ❖ **Brocade Documentation help**
- ❖ **Brocade Legacy / McData Documentation help**
- ❖ **IBM Blade Center Documentation help**
- ❖ **Dell Power Edge Documentation help**
- ❖ **Cisco MDS Storage Networking Documentation help**
- ❖ **HP Blade Server help**
- ❖ **EMC Storage System / Support Matrix help**

*Note:* The IP addresses/Figures/WWN and Configuration examples are solely either designed/set and/or used/modified from the available documentation/online material mentioned in Annexure - A.

No part of the Documentation/IP Address/WWN/ Hostname/Figures/Configuration examples or other material used to write this article has any resemblance to any customers' environment.

### **Documentation help on Legacy McData / Brocade 4Gb SAN Switch Module**

***Brocade Related Figures in this paper are used from the following multiple documentation listed below along with the online resources.***

Brocade 4GB switch product manual - Included on the Brocade 4Gb SAN Switch Module for IBM Blade Center Documentation CD-ROM, the IBM Blade Center Web Site or on the Brocade Web site, through the Brocade Connect site, also can be accessed online as html.



The following are Brocade/Legacy McData documents resources for the 4GB SAN Switch Module:

- Brocade SilkWorm 4020 Hardware Reference Manual
- Brocade 4Gb SAN Switch Module for IBM Blade Center (SilkWorm 4020)
- McData 4Gb Fibre Channel Switch Module for IBM eServer Blade Center Installation Guide
- McData 4Gb Fibre Channel Switch Module for IBM eServer Blade Center Management Guide

The guides can also be obtained through the Brocade Connect Web site:  
<http://www.brocadeconnect.com> (after OEM Blade Server Product registration only)

For Legacy McData/Brocade documentation, visit the Brocade SAN Info Center and click the Resource Library location at <http://www.brocade.com> – all documents are accessible after registration only.

Additional Information can also be obtained from the following Guides.

#### Fabric OS

- Brocade Fabric Manager Administrator's Guide
- Brocade Fabric OS Command Reference Manual

#### Fabric OS Features

- Brocade Web Tools Administrator's Guide
- Brocade Fabric Watch Administrator's Guide
- Brocade Secure Fabric OS Administrator's Guide

## **Cisco Documentation**

---

Cisco-related figures used/modified/edited in this paper belong to the following documentation set for the Cisco MDS 9000 Family.

### **Cisco Fabric Manager**

- Cisco MDS 9000 Family Fabric Manager Configuration Guide
- Cisco MDS 9000 Fabric Manager Online Help

### **Command-Line Interface**

- Cisco MDS 9000 Family CLI Quick Configuration Guide
- Cisco MDS 9000 Family CLI Configuration Guide
- Cisco MDS 9000 Family Command Reference
- Cisco MDS 9000 Family Quick Command Reference
- Cisco MDS 9020 Fabric Switch Configuration Guide and Command Reference

### **Troubleshooting and Reference**

- Cisco MDS 9020 Fabric Switch System Messages Reference
- Cisco MDS 9000 Family Troubleshooting Guide
- Cisco MDS 9000 Family MIB Quick Reference
- Cisco MDS 9020 Fabric Switch MIB Quick Reference
- Cisco MDS 9000 Family SMI-S Programming Reference
- Cisco MDS 9000 Family System Messages Reference

## **Online Resources**

---

The Blade Server figures (Dell/IBM/HP), EMC figures (Clariion), Brocade figures/concept and other related details in this document also derive from the following multiple online resources.

### **HP Online resources**

<http://h71028.www7.hp.com/enterprise/cache/81454-0-0-0-121.html>

### **Dell Power Edge Server Documentation Guide**

<http://support.dell.com/support/edocs/systems/pe1855/en/index.htm>

### **IBM Blade Center**

<http://www-03.ibm.com/systems/bladecenter>

### **Brocade Online Resources**

[http://www.brocade.com/san/white\\_papers.jsp](http://www.brocade.com/san/white_papers.jsp)

### **EMC Online Resources**

CLARiiON products

<http://www.emc.com/products/family/clariion-family.htm>

EMC Topology Guide and EMC Support Matrix→ This guide is available via Powerlink. Visit <https://powerlink.emc.com/> to register and gain access to this material.