

A rectangular box with a double-line border. The background is a gradient of blue, with a white vertical bar on the right side. The title 'Fibre Channel SAN Topologies' is centered in white, bold, sans-serif font.

# Fibre Channel SAN Topologies

Version 4.1

- Fibre Channel Topology Overview
- Simple and Complex Fibre Channel SAN Topologies
- Brocade Virtual Fabrics Case Studies
- FICON Topologies

**Erik Smith**  
**Richard Hultman**  
**Dennis Klopping**

**Copyright © 2011- 2015 EMC Corporation. All rights reserved.**

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

EMC<sup>2</sup>, EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United State and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulator document for your product line, go to EMC Online Support (<https://support.emc.com>).

**Part number H8074.7**

<b>Preface</b> .....	9
<b>Chapter 1    Fibre Channel SAN Topologies</b>	
Fibre Channel topology overview .....	16
Instructions for using this TechBook.....	17
General best practices.....	19
Switch-specific best practices .....	24
Connectrix B-Series.....	24
Connectrix MDS.....	27
Host and storage layout.....	27
Switch and fabric management .....	29
Connectrix B-Series.....	29
Connectrix MDS.....	30
Security .....	31
Connectrix B-Series.....	31
Connectrix MDS.....	32
<b>Chapter 2    Simple Fibre Channel SAN Topologies</b>	
Single switch fabrics .....	34
Overview of fabric design considerations.....	34
Connectrix B example .....	35
Connectrix MDS-Series example .....	38
Two switch fabrics .....	46
Overview of fabric design considerations.....	46
Connectrix B example .....	47
Connectrix MDS example.....	59
Blade switch with direct attached storage.....	69
General steps to set up a blade server .....	70

Best practices .....	73
Host and storage layout.....	73
Switch and fabric management .....	74
Security.....	74
IBM Brocade example .....	75

### Chapter 3    **Complex Fibre Channel SAN Topologies**

Best practices .....	88
ISL subscription .....	88
Host and storage layout.....	88
Switch and fabric management .....	88
Security.....	88
Four switch full mesh.....	89
Overview of fabric design considerations .....	89
Connectrix B example .....	90
Connectrix MDS example.....	101
Compound core edge topologies.....	113
Overview of fabric design considerations .....	113
Connectrix B example .....	115
Connectrix MDS example.....	131
Heterogeneous switch interoperability .....	148
Interoperability overview.....	149
Heterogeneous SAN design.....	151
How to set up an interoperable switched fabric topology .....	151
Cisco Inter VSAN Routing (IVR) in a heterogeneous environment .....	180
Vendor-specific switch settings for interop .....	184
Heterogeneous interoperability test information .....	190

### Chapter 4    **Monitoring your SAN**

Introduction .....	202
Switch-based error types .....	204
Fabric resiliency features and recommendations.....	210
Brocade SAN resiliency features .....	210
Fabric resiliency thresholds .....	212
Quick reference for steps to address switch-based errors issues .....	213
Brocade fabric resiliency concepts.....	218
Latency and congestion bottleneck conditions .....	218
Latency severities .....	221

	Latency detection, notification, isolation, and mitigation .....	222
	Configuring FS features case study .....	225
	Case study: Brocade CLI and CMDCE .....	225
	Configuring Bottleneck Detection .....	229
	Enabling port fencing .....	233
	Configuring Edge Hold Time.....	238
	Summary .....	240
<b>Chapter 5</b>	<b>Brocade Virtual Fabrics Case Study</b>	
	Brocade Virtual Fabrics case study overview .....	242
	Objectives of Virtual Fabrics architecture .....	243
	Logical Switch capability .....	244
	Virtual Fabrics and inter-switch links .....	247
	How to configure Brocade Virtual Fabrics case study .....	250
	Brocade Virtual Fabrics versus traditional Cisco Virtual SANs .....	268
	Domain IDs must be allocated to Logical Switches while partitioning an existing fabric into logical/virtual fabrics.....	268
	Zoning information may be shared between Logical Switches, even when FIDs may not match.....	269
	Switch name must be configured on the Logical Switches.....	270
	Deleting a VF .....	270
	Brocade Base Switches concept.....	271
	FCR with Brocade VFs .....	271
<b>Chapter 6</b>	<b>FICON Topologies</b>	
	Overview .....	274
	Topology support .....	276
	Zoning practices .....	277
	Cascading .....	278
	Terminology .....	279
	IOCP considerations .....	281
	FICON and EMC ControlCenter.....	283
	CUP (fabric management server).....	284
	Connectrix B series.....	285
	Supported products.....	285
	Configuring.....	285
	OCP considerations .....	286

CUP support.....	286
Switch node identifier .....	287
EMC documentation .....	287
Connectrix MDS series.....	288
Supported products .....	288
Requirements .....	288
Configuring .....	289
OCP considerations.....	290
CUP support.....	290
FICON configuration files.....	290
Switch node identifier.....	291
FICON port numbering.....	292
References .....	293
<b>Index.....</b>	<b>295</b>

	<b>Title</b>	<b>Page</b>
1	Single switch topology example .....	34
2	Single switch SAN using DS-5100B and EZSwitchSetup with zoning ...	35
3	Connectrix MDS 9506 using VSANs .....	38
4	Two switch fabric example .....	46
5	Two switch SAN with DS-5100B and DS-5300B using CLI to configure SAN .....	48
6	Port settings .....	55
7	Two Connectrix MDS 9506s using VSANs .....	59
8	(a) Blade server direct attached to storage (b) two-tier FC SAN .....	70
9	IBM blade server directly attached to storage .....	75
10	IBM Brocade 4 GB SAN switch module (32R1812) directly attached to storage .....	76
11	Four switch full mesh fabric .....	89
12	Four switch fabric with ED-48000Bs .....	91
13	Four Connectrix MDS 9506s full mesh configuration .....	101
14	Four switch compound core edge switches .....	113
15	Four ED-48000Bs in full mesh configuration with edge switches attached.....	115
16	Four MDS 9513s in full mesh configuration with edge switches attached.....	132
17	Phase 1: Basic configuration .....	155
18	Phase 2: Adding Connectrix MDS 9513 .....	157
19	Phase 3: Moving half the host and storage ports .....	160
20	Phase 4: Completely moving host and storage ports .....	161
21	Phase 5: Adding Connectrix MDS 9216 .....	162
22	Phase 6: Moving hosts and storage to a new edge .....	164
23	Phase 7: Adding Connectrix MDS switch to core example .....	165
24	Phase 1: Basic configuration .....	168
25	Phase 2: Adding Connectrix MDS 9513 .....	171
26	Phase 3: Moving half the host and storage ports .....	174

27	Phase 4: Completely moving host and storage ports .....	175
28	Phase 5: Adding Connectrix MDS 9216 .....	176
29	Phase 6: Moving hosts and storage to a new edge .....	178
30	Phase 7: Adding Connectrix M switch to the core .....	179
31	Fabric wide effects of a latency bottleneck condition .....	219
32	Fabric wide effects of a congestion bottleneck condition .....	220
33	Uncongested environment .....	226
34	Impact of a slow drain port .....	227
35	Buffer Queue for port 7 continues to grow .....	228
36	Buffer Queue for port 7 on Switch A, port 1 .....	229
37	Port fencing dialog box .....	235
38	Port Fencing dialog box .....	236
39	Create a customized threshold .....	237
40	Apply customized threshold .....	237
41	Logical Switches .....	244
42	DISL connections between Logical Switches .....	247
43	ISL connection between Logical Switches and non-VF-capable switch .....	248
44	XISL connection between Base Switches (BS) .....	249
45	Topology A example .....	250
46	Topology B example .....	251
47	Block diagram of fabric topology .....	261
48	Cascaded IOCP .....	281
49	Cascaded FICON IOCP .....	281



---

*This EMC Engineering TechBook provides a high-level overview of Fibre Channel SAN topologies, discusses simple and complex Fibre Channel SAN topologies, shows how to monitor your SAN, and provides case studies for Brocade Virtual Fabrics. FICON connectivity is also discussed.*

*E-Lab would like to thank all the contributors to this document, including EMC engineers, EMC field personnel, and partners. Your contributions are invaluable.*

*As part of an effort to improve and enhance the performance and capabilities of its product lines, EMC periodically releases revisions of its hardware and software. Therefore, some functions described in this document may not be supported by all versions of the software or hardware currently in use. For the most up-to-date information on product features, refer to your product release notes. If a product does not function properly or does not function as described in this document, please contact your EMC representative.*

---

**Note:** This document was accurate at publication time. New versions of EMC documents might be released on EMC Online Support at <https://support.EMC.com>. Check to ensure that you are using the latest version of this document.

---

### **Audience**

This TechBook is intended for EMC field personnel, including technology consultants, and for the storage architect, administrator, and operator involved in acquiring, managing, operating, or designing a networked storage environment that contains EMC and host devices.

## EMC Support Matrix and E-Lab Interoperability Navigator

For the most up-to-date information, always consult the *EMC Support Matrix* (ESM), available through E-Lab Interoperability Navigator (ELN) at <http://elabnavigator.EMC.com>.

## Related documentation

Related documents include:

- ◆ The following documents, including this one, are available through the E-Lab Interoperability Navigator, **Documents > Topology Resource Center**, at <http://elabnavigator.EMC.com>.

These documents are also available at the following location:

<http://www.emc.com/products/interoperability/topology-resource-center.htm>

- *Backup and Recovery in a SAN TechBook*
- *Building Secure SANs TechBook*
- *Extended Distance Technologies TechBook*
- *Fibre Channel over Ethernet (FCoE): Data Center Bridging (DCB) Concepts and Protocols TechBook*
- *Fibre Channel over Ethernet (FCoE): Data Center Bridging (DCB) Case Studies TechBook*
- *iSCSI SAN Topologies TechBook*
- *Networked Storage Concepts and Protocols TechBook*
- *Networking for Storage Virtualization and RecoverPoint TechBook*
- *WAN Optimization Controller Technologies TechBook*
- *EMC Connectrix SAN Products Data Reference Manual*
- *Legacy SAN Technologies Reference Manual*
- *Non-EMC SAN Products Data Reference Manual*
- ◆ RSA security solutions documentation, which can be found at <http://RSA.com> > **Content Library**

All of the following documentation and release notes can be found at EMC Online Support at <https://support.EMC.com>.

EMC hardware documents and release notes include:

- ◆ Connectrix B series
- ◆ Connectrix MDS (release notes only)
- ◆ VNX series and CLARiiON
- ◆ Symmetrix
- ◆ VMAX

EMC software documents include those on:

- ◆ RecoverPoint
- ◆ TimeFinder
- ◆ PowerPath

The following E-Lab documentation is also available:

- ◆ Host Connectivity Guides
- ◆ HBA Guides

For Cisco and Brocade documentation, refer to the vendor's website.

- ◆ <http://cisco.com>
- ◆ <http://brocade.com>

### Authors of this TechBook

This TechBook was authored by Erik Smith, Aditya Nadkarni, Richard Hultman, and Dennis Klopping, with contributions from EMC employees, EMC engineers, EMC field personnel, and partners.

**Erik Smith** is a Consulting Technologist for the Connectrix business unit within EMC Engineering. Over the past 16 years, Erik has held various technical roles in both EMC Engineering and Technical Support. Erik has authored and coauthored several EMC TechBooks. Erik is also a member of T11.

**Richard Hultman** is a Consultant Systems Integration Engineer and has been with EMC for over 16 years. For the past 13 years, Rich has worked in the E-Lab qualifying switch firmware and hardware. Rich has over 30 years of experience including designing personal computer hardware, symmetric multi-processing systems, and disk adapters, as well as developing firmware for disk controllers.

**Dennis Klopping** is a a Principal Integration Engineer in EMC's E-Lab Product Certification and Test group. Dennis has been with EMC for over 14 years, focusing on IBM zSeries FICON interconnectivity with EMC products. He is involved with the EMC/IBM zSeries relationship, exchanging early ship features and microcode between both companies and joint customer escalations. Prior to EMC, Dennis worked at IBM for over 22 years, concentrating on IBM mainframe support.

### Conventions used in this document

EMC uses the following conventions for special notices:

#### **IMPORTANT**

**An important notice contains information essential to software or hardware operation.**

**Note:** A note presents information that is important, but not hazard-related.

## Typographical conventions

EMC uses the following type style conventions in this document.

Normal	Used in running (nonprocedural) text for: <ul style="list-style-type: none"> <li>Names of interface elements, such as names of windows, dialog boxes, buttons, fields, and menus</li> <li>Names of resources, attributes, pools, Boolean expressions, buttons, DQL statements, keywords, clauses, environment variables, functions, and utilities</li> <li>URLs, pathnames, filenames, directory names, computer names, links, groups, service keys, file systems, and notifications</li> </ul>
<b>Bold</b>	Used in running (nonprocedural) text for names of commands, daemons, options, programs, processes, services, applications, utilities, kernels, notifications, system calls, and man pages  Used in procedures for: <ul style="list-style-type: none"> <li>Names of interface elements, such as names of windows, dialog boxes, buttons, fields, and menus</li> <li>What the user specifically selects, clicks, presses, or types</li> </ul>
<i>Italic</i>	Used in all text (including procedures) for: <ul style="list-style-type: none"> <li>Full titles of publications referenced in text</li> <li>Emphasis, for example, a new term</li> <li>Variables</li> </ul>
Courier	Used for: <ul style="list-style-type: none"> <li>System output, such as an error message or script</li> <li>URLs, complete paths, filenames, prompts, and syntax when shown outside of running text</li> </ul>
<b>Courier bold</b>	Used for specific user input, such as commands
<i>Courier italic</i>	Used in procedures for: <ul style="list-style-type: none"> <li>Variables on the command line</li> <li>User input variables</li> </ul>
< >	Angle brackets enclose parameter or variable values supplied by the user
[ ]	Square brackets enclose optional values

**Where to get help** EMC support, product, and licensing information can be obtained as follows.

EMC support, product, and licensing information can be obtained on the EMC Online Support site as described next.

---

**Note:** To open a service request through the EMC Online Support site, you must have a valid support agreement. Contact your EMC sales representative for details about obtaining a valid support agreement or to answer any questions about your account.

---

### **Product information**

For documentation, release notes, software updates, or for information about EMC products, licensing, and service, go to the EMC Online Support site (registration required) at:

<https://support.EMC.com>

### **Technical support**

EMC offers a variety of support options.

**Support by Product** — EMC offers consolidated, product-specific information on the Web at:

<https://support.EMC.com/products>

The Support by Product web pages offer quick links to Documentation, White Papers, Advisories (such as frequently used Knowledgebase articles), and Downloads, as well as more dynamic content, such as presentations, discussion, relevant Customer Support Forum entries, and a link to EMC Live Chat.

**EMC Live Chat** — Open a Chat or instant message session with an EMC Support Engineer.

### **eLicensing support**

To activate your entitlements and obtain your Symmetrix license files, visit the Service Center on <https://support.EMC.com>, as directed on your License Authorization Code (LAC) letter e-mailed to you.

For help with missing or incorrect entitlements after activation (that is, expected functionality remains unavailable because it is not licensed), contact your EMC Account Representative or Authorized Reseller.

For help with any errors applying license files through Solutions Enabler, contact the EMC Customer Support Center.

If you are missing a LAC letter, or require further instructions on activating your licenses through the Online Support site, contact EMC's worldwide Licensing team at [licensing@emc.com](mailto:licensing@emc.com) or call:

- ◆ North America, Latin America, APJK, Australia, New Zealand: SVC4EMC (800-782-4362) and follow the voice prompts.

- ◆ EMEA: +353 (0) 21 4879862 and follow the voice prompts.

**We'd like to hear from you!**

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Send your opinions of this document to:

[techpubcomments@emc.com](mailto:techpubcomments@emc.com)

Your feedback on our TechBooks is important to us! We want our books to be as helpful and relevant as possible. Send us your comments, opinions, and thoughts on this or any other TechBook to:

[TechBooks@emc.com](mailto:TechBooks@emc.com)

# Fibre Channel SAN Topologies

This chapter provides an overview of Fibre Channel SAN topologies.

- ◆ Fibre Channel topology overview ..... 16
- ◆ Instructions for using this TechBook..... 17
- ◆ General best practices..... 19
- ◆ Switch-specific best practices ..... 24
- ◆ Switch and fabric management..... 29
- ◆ Security ..... 31

## Fibre Channel topology overview

This chapter provides information on Fibre Channel SAN topologies.

For valuable information that may be helpful prior to building a SAN, refer to the *Networked Storage Concepts and Protocols TechBook*, available on the [E-Lab Navigator](#), **Documents > Topology Resource Center**.

For purposes of this document, EMC® E-Lab™ uses the following definitions for simple and complex SANs:

- ◆ A *simple* Fibre Channel SAN consists of less than four Directors and switches connected by inter-switch links (ISLs) and has no more than two hops.
- ◆ A *complex* Fibre Channel SAN consists of four or more Directors and switches connected by ISLs and has any number of hops.



---

## Instructions for using this TechBook

This TechBook was written so that it could be read from front to back or used as a quick reference guide. The concept of *inheritance* was borrowed from the software development community to avoid duplicate best practices, host and storage layout, switch and fabric management, and security information.

The information in this TechBook is divided into four levels:

- ◆ The top level is the “[Fibre Channel topology overview](#)” on [page 16](#) and contains overall best practices, host and storage layout, switch and fabric management, and security information for all Fibre Channel SANs.
- ◆ The next level consists of two chapters: [Chapter 2, “Simple Fibre Channel SAN Topologies,”](#) and [Chapter 3, “Complex Fibre Channel SAN Topologies.”](#) Each of these chapters has best practices, host and storage layout, switch and fabric management, and security information specific to each configuration and it also inherits all of the information listed in “[Fibre Channel topology overview.](#)”
- ◆ Next in the hierarchy is the actual topology (that is, “[Single switch fabrics](#)” on [page 34](#), “[Two switch fabrics](#)” on [page 46](#), “[Blade switch with direct attached storage](#)” on [page 69](#), “[Four switch full mesh](#)” on [page 89](#), “[Compound core edge topologies](#)” on [page 113](#), and “[Heterogeneous switch interoperability](#)” on [page 148](#)). Each of these sections has best practices, host and storage layout, switch and fabric management, and security information specific to each topology, and also inherits all of the best practices, host and storage layout, switch and fabric management, and security information from [Chapter 2, “Simple Fibre Channel SAN Topologies,”](#) or [Chapter 3, “Complex Fibre Channel SAN Topologies.”](#)
- ◆ Finally, the actual case studies contain best practices, host and storage layout, switch and fabric management, and security information specific to the case study and each case study inherits all of the best practices above it. For easy reference, each section contains a link to the best practices section from which it inherits.

The inheritance approach is meant to enhance the readability and supportability of the document.

General information is provided in this section for the following SAN topologies:

- ◆ Chapter 2, "Simple Fibre Channel SAN Topologies"
- ◆ Chapter 3, "Complex Fibre Channel SAN Topologies"

Examples are provided using EMC Connectrix® B and Connectrix MDS switches. Information specific to these examples is detailed as needed.

This chapter also provides detailed information using case studies: Chapter 5, "Brocade Virtual Fabrics Case Study," and Chapter 6, "EMC RecoverPoint Case Study."

All of the configurations shown throughout this chapter use the same host and storage ports. This was done intentionally to better expose the differences between switch implementations and fabric topologies rather than attempt to address all possible host and storage combinations.

For a complete list of supported host/switch/storage combinations, and the most up-to-date supported configurations, refer to the *EMC Support Matrix*, available on the [E-Lab Navigator](#).

---

## General best practices

Consider the following general best practices:

- ◆ To ensure maximum data availability, EMC requires redundant physical and logical paths to prevent any single point of failure.
- ◆ Plan for failures
  - Connect the host and storage ports in such a way as to prevent a single point of failure from affecting redundant paths. For example, if you have a dual-attached host and each HBA accesses its storage through a different storage port, do *not* place both storage ports for the same server on the same Line Card or ASIC.
  - Use two power sources.
- ◆ For host and storage layout

To reduce the possibility of congestion, and maximize ease of management, connect hosts and storage port pairs to the same switch where possible.

---

**Note:** Refer to “Host and storage layout” on page 27 for information on host and storage layout.

---

- ◆ Plan cabling

Table 1 on page 20 lists the typical distances that can be supported for Fibre Channel with the different fiber types and link speeds. OM2 cable was the standard 50um cable used with Fibre Channel for many years. The higher link speeds resulted in the need for OM3 cable, which can support longer distances at the higher link speeds. OM3 cables typically have an Aqua colored jacket, as opposed to the standard orange colored jacket on OM1 and OM2 cable. Table 1 provides information on multimode media maximum distances.

Table 1 Multimode media maximum distances

Protocol	Transceiver type	Speed	Multimode media maximum distance			
			62.5µm/200MHz* km (OM1) [62.5 micron]	50µm/500 MHz* km (OM2) [50 micron]	50µm/2000 MHz* km (OM3) [50 micron]	50µm /47000 MHz* km (OM4) [50 micron]
FC	SW	2 Gb/s	150m	300m	500m	500m *
		4 Gb/s	70m	150m	380m	400m
		8 Gb/s	21m	50m	150m	190m
		10 Gb/s	33m	82m	300m	550m
		16 Gb/s	15m	35m	100m	125m
GbE	SW	1 Gb	300m	550m	1000m	1000m *
		10 Gb	33m	82m	300m	550m
		40 Gb	N/A	N/A	100m	125m

\* Denotes at least this distance. No documented distance is available at the time of this publication.

- ◆ For security

**Note:** Refer to “Security” on page 31 for information on security.

- ◆ Use single initiator zoning

**Note:** Single initiator zoning is recommended unless specified otherwise by the device vendor. For example, IBM SVC does not comply with the single initiator zoning recommendation. For specific information on how the zoning configuration differs from other device types, refer to the Zoning section in the *IBM system Storage SAN Volume Controller and Storwize V7000 Best Practices and Performance Guidelines*, available at <http://www.redbooks.ibm.com/redbooks/pdfs/sg247521.pdf>.

- Single initiator zoning with Cisco MDS Smart Zoning
  - Cisco MDS Smart Zoning supports zoning among more devices by reducing the number of zoning entries to be programmed. Smart Zoning considers device type information (initiators or targets) without increasing the size of the zone set. Smart Zoning enables you to create a single zone on every end device regardless of the type, such as a host, a target, or both, without creating multiple single initiator zones.

One zone can support multiple initiators (hosts) while still maintaining the single initiator zoning rule. This enables fewer zones to be created and simplifies single initiator zoning.

Consider the following recommendations:

- Only a host or target be used as the device type.
- For array ports that may function both as host and target device type, keep the number of hosts it is zoned to as low as possible and use a separate zone for array-to-array connections.
- Keep FC SRDF replication ports (RA), RecoverPoint Appliance ports, VPLEX ports, or other application-based FC ports in their own zone.

For Smart Zoning configuration information, refer to the *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide* at

[http://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/5\\_2/configuration/guides/fabric/nx-os/nx\\_os\\_fabric.html](http://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/5_2/configuration/guides/fabric/nx-os/nx_os_fabric.html)

- For Open Systems environments, if Cisco Smart Zoning is not being used, ideally each initiator will be in a zone with a single target. However, due to the significant management overhead that this can impose, single initiator zones can contain multiple target ports but should never contain more than 16 target ports.
- Special consideration for EMC Symmetrix® Remote Data Facility (SRDF®) ports are as follows:
  - If desired, the administrator can use *SRDF Single Zoning* when multiple sources and targets are contained in a single zone.
  - An SRDF zone should only contain RF ports.
  - If multiple zones are used, zoning must be designed so that functionality such as Dynamic SRDF meets customer requirements. For example, if Dynamic SRDF is in use, zoning requirements can change. With Dynamic SRDF, any RFs that have Dynamic SRDF connectivity established to one another must be in the same zone.
  - A maximum of 10 RFs per switch is recommended, for example when you have a two site configuration with four EMC Symmetrix DMX™ systems in each site. Each DMX contributes four SRDF ports. No single switch should have

more than ten RFs connected to it. In this example, a minimum of two switches need to be deployed at each site. The main reason for restricting a switch to ten RFs is due to Name Server traffic. Name Server traffic is an important consideration and needs to be kept to a minimum to minimize link recovery times when RSCNs occur. By distributing across multiple switches, processing of Name Server traffic is also able to scale.

- ◆ Use dual management networks whenever two or more FC switches are being used
- ◆ Before building the SAN, gather the following customer-supplied information that will be needed during the implementation
  - Information for each switch:
    - Switch names
    - Port names
    - Passwords
    - Number of HBAs and storage arrays to be connected to the switch
    - IP addresses, subnet mask, and gateway for each switch

These values are used in both manual and GUI-based setup methods.

- ◆ Use a port fencing policy

For more information on Port fencing, refer to “[Port fencing](#)” in the *Networked Storage Concepts and Protocols TechBook*, available on the [E-Lab Navigator](#), **Documents > Topology Resource Center**.
- ◆ Use the available performance monitoring tools

For more information on threshold alerts, refer to “[Threshold alerts](#)” in the *Networked Storage Concepts and Protocols TechBook*, available on the [E-Lab Navigator](#), **Documents > Topology Resource Center**.
- ◆ Use the latest supported firmware version and ensure that the same version of firmware is used throughout the fabric. In homogeneous switch vendor environments, all switch firmware versions inside each fabric should be equivalent, except during the firmware upgrade process.

---

**Note:** Refer to the *EMC Support Matrix* for the most up-to-date information.

---

- ◆ Periodically (or following any changes) back up switch configurations
- ◆ Utilize a syslog server

---

**Note:** It is also recommended to install an NTP client on the syslog server.

---

- ◆ Use persistent Domain IDs
- ◆ Estimate light budget

---

## Switch-specific best practices

This section contains switch-specific best practices for:

- ◆ “Connectrix B-Series” on page 29
- ◆ “Connectrix MDS” on page 32

---

### Connectrix B-Series

This section contains the following information:

- ◆ “Preparing for a firmware download” on page 24
- ◆ “Zoning” on page 25
- ◆ “ISL trunking” on page 26

#### Preparing for a firmware download

Before executing a firmware download, performing the following tasks is recommended. This information will not only help validate that the firmware download did not disrupt the existing configuration, but also provide the support team with sufficient information in an unlikely event that something goes wrong after a firmware download.

1. Establish a Telnet session and log in to the Connectrix B switch.
2. It is advisable to have session logging enabled so that the output of your commands is saved to a file.
3. Enter the **switchShow** command to display switch information.
4. Enter the **portCfgShow** command to display current switch port settings.
5. Enter the **nsShow** command to display the total number of devices in the switch local Name Server database.
6. Enter the **nsAllShow** command to display the total number of devices in the switch global Name Server database.
7. Enter the **fabricShow** command to display the total number of domains in the fabric. If you are changing a dual domain Connectrix B director to a single domain Connectrix B director, this value will be one domain less after the operation.
8. Display the MIB configuration information using the **snmpMibCabShow** or **agtCfgShow** command.



9. Upload the switch configuration settings to an FTP server using the **configUpload** command.
10. Enter the **supportShow** and **supportSave** commands (these commands are version dependent) to provide baseline information for advanced support.

---

**Note:** If you are upgrading a Connectrix B director that is configured with two logical domains, perform these steps for *both* logical switches.

---

Before you enter the **firmwareDownload** command, read the release notes for the new firmware to find out if there are any firmware download issues.

After the **firmwareDownload** on a Connectrix B director switch, it is recommended you validate that the firmware versions have been synchronized on both CPs.

---

**Note:** A best practice for performing firmware upgrades is through the CLI.

---

## Zoning

Note the following recommendations and best practices:

- ◆ To make zone edits, use either the Brocade CLI or Connectrix Manager Converged Network Edition (CMCNE), which is the Brocade fabric management GUI.
- ◆ Use only one interface, CLI or CMCNE, for making zone edits at a given time. Using both interfaces concurrently can result in severe issues. For example, zone changes made through one of the interfaces do not get processed and, in fact, are lost. To address these types of issues, Brocade defects have been opened and resolved. SAN administrators should use any one application for all zoning changes. With FOS 7.1.x, there is appropriate user notification embedded in the Brocade firmware to notify users about any concurrent zoning changes.
- ◆ The recommended default zone setting on all switches in a FOS fabric is *no access*, which implies that the access level is closed with all nodes isolated and without access to each other. With this setting, only the user-created zonesets will dictate what nodes can access each other. To prevent the zones from being distributed across the fabric, verify that there are no two switches in the fabric with different default zone settings, (i.e., *no access* vs. *all access*).

- ◆ A zoneset can be managed and activated from any switch in the fabric, but should be managed from a single entry switch within a fabric. This avoids complications arising from multiple users accessing different switches in a fabric to make concurrent zone changes.
- ◆ The system administrators should coordinate zoning configuration activity to avoid running into a situation where two administrators are making changes simultaneously.
- ◆ To avoid any lengthy outages due to errors in Connectrix B SAN configurations, back up the existing configuration before making any changes.
- ◆ To avoid the high risk involved when adding a new unauthorized switch to a Connectrix B fabric, limit the creation of switch-to-switch ports. This can be done by locking the already connected switch-to-switch ports in the SAN using the **portCfgEport** command. Such locking down of E\_Ports is persistent across reboots. A **portCfgEport <port number>,0 <disable>** must be run on ports that are not connected to other switches in the fabric to block them from forming ISLs between switches.

### ISL trunking

More than a best practice, the administrator configuring a Connectrix B SAN must be aware that the frame-level trunking for Connectrix B switches requires all ports in a given ISL trunk to reside within an ASIC group on each end of the link.

- ◆ On 2 Gb/s switches, port groups are built on contiguous 4-port groups, called *quads*. For example, on a Connectrix DS-8B2, there are two quads: ports 0-3 and ports 4-7.
- ◆ On 4 Gb/s switches like the Connectrix DS-4100B, trunking port groups are built on contiguous 8-port groups, called *octets*. In this product, there are four octets: ports 0-7, 8-15, 16-23, and 24-31.
- ◆ On 8 Gb/s switches like the Connectrix DS-5100B, trunking port groups are built on *octets*. On this product, there are five octets: ports 0-7, 8-15, 9-23, 24-31 and 32-39.
- ◆ On 16 Gb/s switches like the Connectrix DS-6510B, trunking port groups are built on *octets*. On this product, there are six octets: ports 0-7, 8-15, 9-23, 24-31, 32-39 and 40-47.

The administrator must use the ports within a group specified above to form an ISL trunk. It is also possible to configure multiple trunks within a port group.

---

## Connectrix MDS

The following are requirements and guidelines for using IVR NAT:

IVR NAT port login (PLOGI) requests received from hosts are delayed for a few seconds to perform the rewrite on the FC ID address. If the host's PLOGI timeout value is set to a value less than five seconds, it may result in the PLOGI being unnecessarily aborted and the host being unable to access the target. EMC® recommends that you configure the host bus adapter for a timeout of at least ten seconds (most HBAs default to a value of 10 or 20 seconds).

---

**Note:** IVR NAT requires Cisco MDS SAN-OS Release 2.1(1a) or later on all switches in the fabric performing IVR. If you have isolated switches with an earlier release that are involved in IVR, you must remove any isolated fabrics from being monitored by Fabric Manager server and then re-open the fabric to use IVR NAT.

---

---

## Host and storage layout

The correct way to attach hosts and storage to a SAN is completely dependent upon customer environments. Historically, the best practice placed hosts on edge switches and high-use storage ports on core switches. This was recommended because high-use storage ports are sometimes accessed by many different hosts on different parts of the fabric. If this is the case in your environment, this configuration would still be the best option. However, if you have high-use storage ports that are only accessed by a couple of hosts and it is possible to locate them all on the same switch, this is the preferred configuration instead of forcing the use of inter-switch links (ISLs). ISLs are a valuable and limited resource and should be reserved for providing connectivity between ports that are unable to be placed on the same switch.

With this in mind, the following information provides helpful general guidelines:

- ◆ Whenever practical, locate HBAs and the storage ports they will access on the same switch. If it is not practical to do this, minimize the number of ISLs the host and storage need to traverse.
- ◆ Some of the switch class products being produced today only contain a single ASIC. If this is the case, then the positioning of the host and storage ports is strictly a matter of personal

preference. However, if the switch being used contains multiple ASICs, try to connect host and storage pairs to the same ASIC. This prevents using the shared internal data transfer bus and reduces switch latency. In addition to performance concerns, consider fault tolerance as well. For example, if a host has two HBAs, each one accessing its own storage port, do not attach both HBAs, both storage ports, or all of the HBA and storage ports to the same ASIC.

- ◆ When working with hosts that have more than one connection to more than one storage port, always connect the HBAs and, if possible, the storage ports that it accesses to different FC switches. If a completely separate fabric is available, connect each HBA and storage port pair to different fabrics. Refer to the “Methodology 1: Balanced fabrics” located in the *Networked Storage Concepts and Protocols TechBook*, available on the [E-Lab Navigator](#), **Documents > Topology Resource Center**.

---

## Switch and fabric management

The following are management interfaces common to all platforms:

- ◆ CLI  
The command line interface (CLI) can be accessed through Telnet and Secure Shell (SSH, which authenticates remote computers). On some platforms, it is also possible to access the CLI using a serial cable.
- ◆ SNMP  
Simple Network Management Protocol (SNMP) is a TCP/IP protocol that generally uses the User Datagram Protocol (UDP) to exchange messages between a management device and a network management system.

Specific switch and fabric management information follows.

---

### Connectrix B-Series

All switches in the EMC Connectrix B family can be managed by CLI, Web Tools, and CMCNE.

- ◆ Web Tools  
Web Tools provides switch/fabric management through a web browser pointed to the IP address of the switch. Brocade uses a built-in web server for this function.
- ◆ Connectrix Manager Converged Network Edition (CMCNE)  
CMCNE is a SAN management software tool for managing, monitoring, configuring, and reporting on Connectrix switches and directors. CMCNE provides a single pane of glass view into the management of traditional FC SANs and converged FCoE networks.

It not only provides a GUI for the core SAN management functions such as zoning and switch configuration, but also provides for switch and event monitoring with dashboards and automated call home to EMC. Also provided are real-time and historical performance monitoring, reporting, security, and RBAC management.

CMCNE can also manage Brocade IP switch products within the same management interface. It is available in three basic tiers that can scale to your environment size:

- CMCNE Professional Edition. Intended for the management of smaller fabrics with Connectrix B Series departmental switches.
- CMCNE Professional Plus Edition. Increases fabric scalability and adds support for the ED-DCX-4S-B and ED-DCX8510-4B.
- CMCNE Enterprise Edition. A full-featured Enterprise class management solution for the largest of data center SAN infrastructures with the high density Director class products, ED-DCX-B and ED-DCX8510-8B.

---

### Connectrix MDS

- ◆ Device Manager
- ◆ Fabric Manager
- ◆ Fabric Manager server

---

## Security

It is important to secure your fabric. General security best practices for an FC SAN are as follows:

- ◆ Implement some form of zoning
- ◆ Change the default password
- ◆ Disable unused or infrequently used Management Interfaces
- ◆ Use SSL or SSH if available
- ◆ Limit physical access to FC switches

Specific switch and fabric management information follows.

---

### Connectrix B-Series

Fabric Operating Software (FOS) versions 5.2.x and above now include an access control list (ACL) feature which gives the SAN administrator the ability to restrict both device and switch login throughout the fabric.

The following two ACL policies offer the base Fabric Operating System (FOS):

- ◆ Device Connection Control (DCC) policy
- ◆ Switch Connection Control (SCC) policy

FOS versions 5.2.x and above introduced Role Based Access Control (RBAC).

Also introduced in FOS versions 5.2.x and above was the concept of Admin Domains (AD).

To maintain a secure network, you should avoid using Telnet (you can use secTelnet, or any other unprotected application when you are working on the switch. For example, if you use Telnet to connect to a machine, and then start an SSH or secure Telnet session from that machine to the switch, the communication to the switch is in clear text and therefore *not* secure.

The FTP protocol is also not secure. When you use FTP to copy files to or from the switch, the contents are in clear text. This includes the remote FTP server login and password. This limitation affects the following commands: **saveCore**, **configUpload**, **configDownload**, and **firmwareDownload**.

---

## Connectrix MDS

The MDS supports the following protocols:

- ◆ SSH
- ◆ SFTP
- ◆ FCSP
- ◆ DHCHAP



## Simple Fibre Channel SAN Topologies

This chapter provides the following information on simple Fibre Channel SAN topologies.

- ◆ Single switch fabrics ..... 34
- ◆ Two switch fabrics..... 46
- ◆ Blade switch with direct attached storage..... 69

## Single switch fabrics

This section provides examples of single switch fabrics.

### Overview of fabric design considerations

#### General layout

A single switch fabric consists of only a single switch (Figure 1). The switch is also connected to a single management LAN through IP.




Figure 1 Single switch topology example

#### Best practices

For general best practices in single switch fabrics, refer to “General best practices” on page 19.

#### Host and storage layout

**Note:** The correct way to attach hosts and storage to a SAN is completely dependent upon the customers’ environment, but the following information may be helpful.

In this example, both the host and storage ports are located on the same switch. There is also the opportunity to locate the host and storage pairs so that they are in the same quad or octet of ports which are controlled by the same switch ASIC. This connection scenario helps eliminate frames traveling over the backplane of the switch and increases the speed of frame routing. Many switch class products being sold today are switch-on-a-chip architecture, and do not contain many discrete ASICs. For these architectures, port placement for performance or HA concerns does not need to be considered.

For general information on host and storage layout in single switch fabrics, refer to “[Host and storage layout](#)” on page 27.

### Switch and fabric management

All switch and fabric management tools can be used to configure this environment. For general information on switch and fabric management, refer to “[Switch and fabric management](#)” on page 29.

### Security

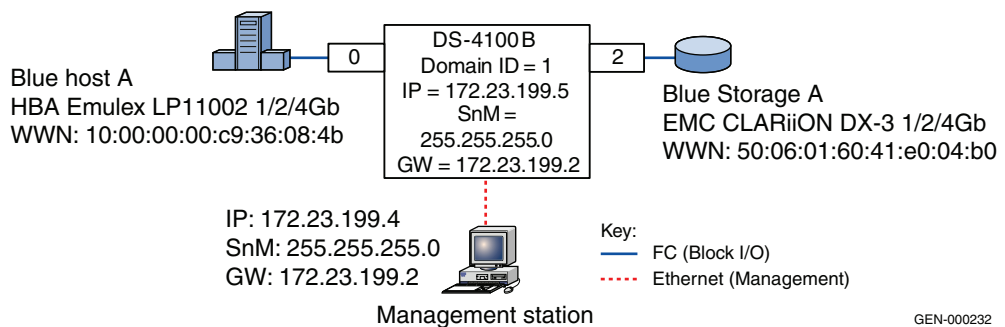
Even when dealing with single switch fabrics, it is important to think about security. Typically, in a single switch fabric, if physical access to the switch can be controlled, default passwords are changed, and unused interfaces are disabled, zoning will be enough to secure the fabric. It is also recommended to implement some form of Port Binding. This can be as simple as enabling Port Binding if the feature exists, or as complicated as disabling all unused ports and hard-setting port type to F\_Port.

Refer to “[Security](#)” on page 31 for general security information.

## Connectrix B example

### General layout

[Figure 2](#) shows a single switch SAN using DS-5100B, and EZSwitchSetup with typical zoning.



**Figure 2** Single switch SAN using DS-5100B and EZSwitchSetup with zoning

### Supported switches

Connectrix DS-220B and DS-5100B are supported with EZSwitchSetup.

### Best practices

For general best practices in single switch fabrics, refer to “[General best practices](#)” on page 19.

For Connectrix B specific best practices, refer to “[Connectrix B-Series](#)” on page 24.

**Host and storage layout**

When using Typical Zoning in EZSwitchSetup, the GUI guides the SAN administrator on recommended placement of both host and storage.

In a single switch SAN design, both hosts and storage can be placed anywhere on the switch. Connectrix B switches group ASIC-controlled port groups, either by quad or octet. In some circumstances, you should connect both host and storage to these groups which eliminates frames from traveling through the backplane of the switch to reach their destination.

For general information on host and storage layout in single switch fabrics, refer to [“Host and storage layout” on page 27](#).

**Switch and fabric management**

For this case study, the EZSwitchSetup tool is used to setup a Connectrix DS-5100B switch with one host and one storage array. EZSwitchSetup offers three zoning options. The Typical Zoning option is used in this case study. For more information on EZSwitchSetup, refer to [“Connectrix MDS” on page 30](#).

For general information on host and storage layout in single switch fabrics, refer to [“Switch and fabric management” on page 29](#).

**Security**

For general information on security in single switch fabrics, refer to [“Security” on page 31](#).

**Setting up this topology****Assumptions specific to this case study:**

- ◆ The switch is not powered on.
- ◆ Network drop, IP addresses, subnet mask, and gateway are provided by the customer.
- ◆ License keys have been obtained.
  - Go to the URL listed on the transaction code certificate that shipped with the product to get the license key.
- ◆ The customer has provided a server or laptop with CD drive, serial DB-9 connector, and a NIC card to be used to configure the switch.
- ◆ Configuration will be done using the *EZSwitchsetup CD*.

**Configure a single switch SAN**

To configure a single switch SAN, follow these steps:

**Note:** Host and storage ports do not have to be connected until [Step 8](#).

1. Power-on the switch.
2. Connect to the switch management port using RS-232 serial cable. Set the workstation serial port to use 9600 baud rate, 8 databits, no parity, 1 stop bit, and no flow control.
3. Connect RJ-45 network cable from the workstation to the network management port of switch.
4. Launch the EZSwitchSetup CD.
5. Issue the switch IP address (172.23.199.5), subnet mask (255.255.255.0), and default gateway (172.23.199.2), and then click **Next**.
6. Supply a password for the Admin account, issue a switch name, change the date, and then click **Next**.
7. Select **Typical Zoning**, on the **Select Zoning** screen, and then click **Next**.
8. Select the ports as host or storage on the **Configure Typical Zoning** screen, and then click **Next**.  
Each time you click, you toggle the port between a desired host port, designated by the color blue, the letter **H**, and a desired storage port, designated by the color green, and the letter **S**.
9. Connect host and storage using the **Configure Typical Zoning** GUI as a guide.
10. Use the pull-down menu selections on the **Specify Devices** screen to select the number of HBA connections and the number of storage connections planned for the switch, and then click **Next**.
11. The **Connect Devices** screen provides a graphical view of recommended connections for both host and storage on the switch. Connection status is represented by colored lines.
  - A green line indicates a good connection
  - A red line indicates an invalid connection
  - A blue dashed line indicates a missing connectionClick **Next**.
12. The **Finish** screen supplies a summary of the switch configuration and allows you to print the configuration if needed. Click **Finish**.

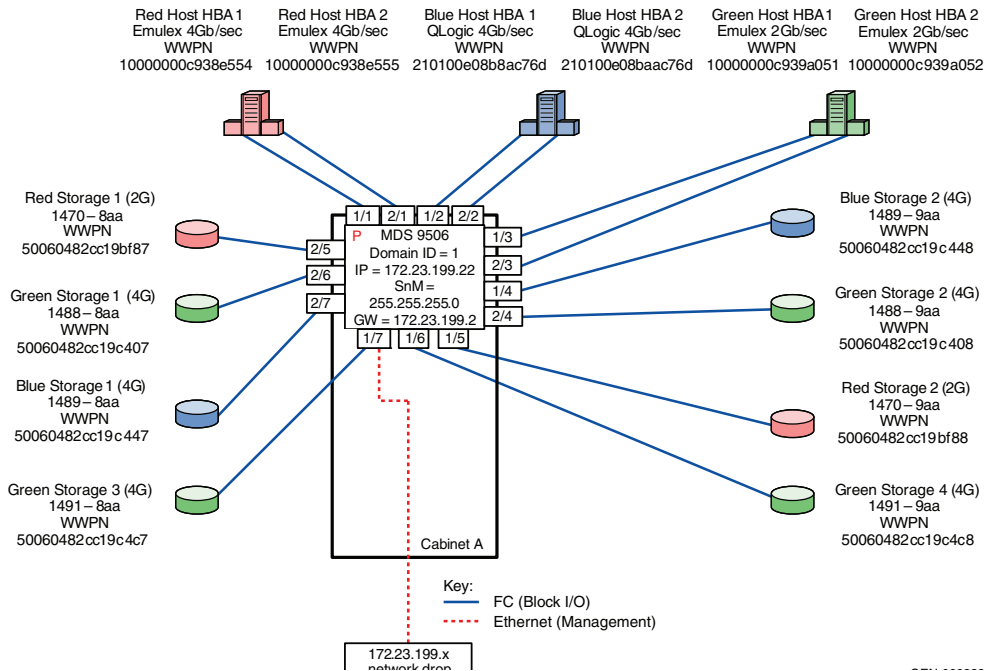
- Validate zoning configuration by ensuring that each HBA has access to the storage device(s).

Each HBA should see every storage device when using **Typical Zoning**.

The **Validate** link in the task panel of the **Switch Manager** page checks for devices that are not zoned properly, allows you to delete the devices from the zoning database, and displays a matrix of which HBA can see which storage device.

### Connectrix MDS-Series example

**General layout** Figure 3 illustrates an Connectrix MDS 9506 using VSANs.



**Figure 3** Connectrix MDS 9506 using VSANs

#### Best practices

For general information on best practices in single switch fabrics, refer to “Best practices” on page 34. Specific information for this example follows.

Port fencing is on by default.

<b>Host and storage layout</b>	<p>For general information on host and storage layout in single switch fabrics, refer to “<a href="#">Host and storage layout</a>” on page 34. Specific information for this example follows.</p> <p>There are no host or storage restrictions for Line Rate Mode cards. Oversubscribed cards should be used for hosts only.</p>
<b>Switch and fabric management</b>	<p>For general information on switch and fabric management in single switch fabrics, refer to “<a href="#">Switch and fabric management</a>” on page 35. Specific information for this example follows.</p> <p>Cisco Fabric Manager can be used for this example.</p>
<b>Security</b>	<p>For general information on security in single switch fabrics, refer to “<a href="#">Security</a>” on page 35. Specific information for this example follows.</p> <p>Port Binding can be used for security.</p>
<b>Setting up this topology</b>	<p><b>Assumptions specific to this case study:</b></p> <ul style="list-style-type: none"> <li>◆ The switches are installed in an EMC-supplied cabinet. <ul style="list-style-type: none"> <li>• For installation instructions, see <i>Connectrix EC-1500 Cabinet Installation and Setup Manual</i>, which can be accessed from EMC Online Support at <a href="https://support.emc.com">https://support.emc.com</a>.</li> </ul> </li> <li>◆ The proper power receptacles have been provided by the customer. <ul style="list-style-type: none"> <li>• For Cabinet power requirements, refer to <i>Connectrix EC-1500 Cabinet Installation and Setup Manual</i>, which can be accessed from EMC Online Support at <a href="https://support.emc.com">https://support.emc.com</a>.</li> </ul> </li> <li>◆ The switches have <i>not</i> been connected to the power source and are <i>not</i> powered on.</li> <li>◆ Network drops, IP addresses, subnet mask, and gateway have been provided by the customer. <ul style="list-style-type: none"> <li>• For switch or cabinet network requirements, refer to the <i>EMC Connectrix SAN Products Data Reference Manual</i>, on the <a href="#">E-Lab Navigator, Documents &gt; Topology Resource Center</a>.</li> </ul> </li> </ul> <hr/> <p><b>Note:</b> Connectrix MDS switches can be placed on either a public or private network. There are advantages to both configurations. For more information, refer to “Public versus private” in the <i>Networked Storage Concepts and Protocols TechBook</i>, available through the on the <a href="#">E-Lab Navigator, Documents &gt; Topology Resource Center</a>.</p>

In this example it is assumed that the customer has provided one Ethernet cable and an IP of 172.23.199.22.

- ◆ The proper number of line cards have been installed in each chassis. In this case, two line cards in each chassis are required and installed in slots 1 and 2.
  - For help in determining how many ports are required, see "Determining customer requirements" in the *Networked Storage Concepts and Protocols TechBook*, available on the [E-Lab Navigator, Documents > Topology Resource Center](#).
- ◆ License keys have been obtained.
  - Go to the URL listed on the transaction code certificate that shipped with the product to get the license key.
- ◆ A laptop, supplied by the installer, will be used to configure the IP addresses of the switches, and this laptop has a serial DB-9 connector.
- ◆ The customer has provided a temporary password that will be used as the default password when configuring the IP address.
- ◆ Fabric Manager will be used for VSAN setup.

### Configure the IP address

To configure the IP address:

1. Power up the cabinet by connecting power cords to the power receptacles provided by the customer.
2. Select one of the switches to configure and set the IP to 172.23.199.22.
3. Supply a network connection to the appropriate subnet.
4. Connect to the serial port of the switch, using an RS232 serial cable, with a baud rate of 9600, 8 data bits, no parity, 1 stop bit and no flow control.

The **login** prompt appears.

5. Log in the first time with a username of *admin* and a password of *admin*.

You are prompted to supply a new strong password for CLI user *admin*.
6. For this example, select **no** when asked if you want to run setup.



**Note:** This example starts with the switch that has a Domain ID of 1 and an IP address of 172.23.199.22.

7. Repeat the above steps for each switch, supplying the appropriate IP.

### CLI commands to configure the IP and gateway

- ◆ Switch# *config terminal*

Enter configuration commands, one per line.

```
Switch(config)# interface mgmt 0
Switch(config-if)# IP address 172.23.199.22 255.255.255.0
```

End with CNTL/Z.

- ◆ Switch# *config terminal*

Enter configuration commands, one per line.

```
Switch(config)# ip default-gateway 172.23.199.2
```

End with CNTL/Z.

To authorize access on a switch for Device and Fabric Manager, run this command on every switch while supplying a username (nnn) and password (ppp):

- ◆ Switch#*conf t*

```
Switch(config)# snmp-server user nnn network-admin auth md5
ppp
Switch(config)#end
Switch# copy running-config startup-config
Switch# exit
```

### Install Fabric Manager and Device Manager

To install Fabric Manager and Device Manager:

1. Open your web browser.
2. Enter the IP address of the switch into the address bar.
3. Follow the prompts and accept all defaults to install both Fabric Manager and Device Manager.

Fabric Manager and Device Manager can be started using the configured SNMP server username and password in “CLI commands to configure the IP and gateway” on page 41.

### Configure a VSAN

To configure a VSAN:

1. Open the Device Manager for the switch with an IP address of **172.23.199.22**.
2. Open the **VSAN** dialog box by selecting **VSAN**.
3. Click **Create**.
4. Enter the value of **100** into the **VSAN ID** field.
5. Set the **VSAN Name** to be **“Red\_VSAN\_100”**.
6. Use the default interop mode.
7. Click **Create**.

### Configure the other VSANs in this physical switch

To configure the other VSANs, complete the following steps:

1. Repeat [Step 2](#) through [Step 7](#), above, for VSAN 200 and 300 noting that:
  - For Virtual switch 200, use a VSAN name of **“Green\_VSAN\_200”**.
  - For Virtual switch 300, use a VSAN name of **“Blue\_VSAN\_300”**.
2. Using the following table, assign and enable the ports to the proper VSAN using Device Manager for the switch with the IP address **172.23.199.22**.

Slot #	Port #	Name	VSAN ID
1	1	Red Host HBA 1	100
1	2	Blue Host HBA 1	300
1	3	Green Host HBA 1	200
1	4	Blue Storage 2	300
1	5	Red Storage 2	100
1	6	Green Storage 4	200
1	7	Green Storage 3	200
1	8		
2	1	Red Host HBA 2	100

Slot #	Port #	Name	VSAN ID
2	2	Blue Host HBA 2	300
2	3	Green Host HBA 2	200
2	4	Green Storage 2	200
2	5	Red Storage 1	100
2	6	Green Storage 1	200
2	7	BlueStorage 1	300
2	8		

### Connect cables

To connect the cables:

1. Connect host and storage ports.
2. Attach fiber cable between switches and N\_Ports.

### Configure domains

To configure domains:

1. Open **Fabric Manager**. A topology of two switches appears.
2. From **Fabric Manager** open the “Red\_VSAN\_100” folder.
3. Select **Domain Manager**.
4. Select the **Configuration** menu.
5. Set a Domain ID, 1 for switch 172.23.199.22 and 2 for 172.23.200.22.
6. To set a principal switch, set the priority to 1 in the domain menu.
7. Repeat [Step 2](#) through [Step 6](#) for “Green\_VSAN\_200” and “Blue\_VSAN\_300”.

### Zone hosts and storage

To zone hosts and storage:

1. From **Fabric Manager**, select “Red\_VSAN\_100”.
2. Select **Edit Full Zone Database**.
3. Create a zone by selecting the **Zone** menu and clicking the **Add** button.

4. Provide a descriptive name for the zone. This example zones "Red host HBA 1" and "Red Storage 1", so "**RedHBA1\_1470\_8aa**" will be used. Click **OK**.
5. Locate, then click, "**Red Host HBA 1**" (WWPN 10000000c938e554) in the **Potential zone members list**.
6. Click the right-pointing arrow on the divider between the **Potential members list** and the zones list to add the HBA to the zone. Select **Add to zone or alias**.
7. Locate, then click, "**Red Storage 1**" (WWPN 50060482cc19bf87) in the **Potential zone members list**.
8. Click the right-pointing arrow on the divider between the **Potential members list** and the **zones list** to add the Storage port to the zone.
9. Repeat [Step 2](#) through [Step 8](#) for all host and storage pairs in the environment.
10. Create a zone set by selecting the **Zonesets** menu, and then click the **Add** button.
11. Provide a descriptive name for the zone set. This example uses the name "**RED Fabric 1**".

Add only those zones that are necessary on Red\_Fabric\_1. In this case, the two red zones listed below, "**RedHBA1\_1470\_8aa**" and "**RedHBA2\_1470\_9aa**" are used. Repeat for other Fabrics, when completed, you should have three zone sets as shown below.

```
Zone set name = "Red_Fabric_1"
```

```
Zone name = "RedHBA1_1470_8aa"
Zone Member = "10000000c938e554"
Zone Member = "50060482cc19bf87"
```

```
"Red_Fabric_2"
```

```
Zone name = "RedHBA2_1470_9aa"
Zone Member = "10000000c938e555"
Zone Member = "50060482cc19bf88"
```

```
Zone set name = "Blue_Fabric_1"
```

```
Zone name = "BlueHBA1_1489_8aa"
Zone Member = "210100e08b8ac76d"
Zone Member = "50060482cc19c447"
```

```
"Blue_Fabric_1"
```

```
Zone name = "BlueHBA2_1489_9aa"  
Zone Member = "210100e08baac76d"  
Zone Member = "50060482cc19c448"
```

```
Zone set name = "Green_Fabric"
```

```
Zone name = "GreenHBA1_AllGreenStorage"  
Zone Member = "10000000c939a051"  
Zone Member = "50060482cc19c407"  
Zone Member = "50060482cc19c408"  
Zone Member = "50060482cc19c4c7"  
Zone Member = "50060482cc19c4c8"
```

```
Zone name = "GreenHBA2_AllGreenStorage"  
Zone Member = "10000000c939a052"  
Zone Member = "50060482cc19c407"  
Zone Member = "50060482cc19c408"  
Zone Member = "50060482cc19c4c7"  
Zone Member = "50060482cc19c4c8"
```

### Complete the SAN setup

At this time, the SAN is ready to pass I/O from host to storage. Other steps, such as configuring LUN masking and modification of host configuration files, are required before the SAN setup is complete. Refer to the *Cisco MDS 9000 Family Fabric Manager Configuration Guide* for more details.

### Configure security

- ◆ Enable Switch Binding.
- ◆ Enable Port Binding.

### Configure proactive monitoring and countermeasures

ISL thresholds are 80% by default.

### Configure port fencing

Port fencing is on by default.


## Two switch fabrics

This section contains information on two switches in a Fibre Channel SAN topology.

### Overview of fabric design considerations

#### General layout

Figure 4 shows an example of a two switch fabric. The switches are connected using two ISLs. Both switches are connected to Management Network A.



GEN-000229

Figure 4 Two switch fabric example

Every switch type can be used in any position in this configuration.

#### Best practices

The following best practices are specific for two switch fabrics.

- ◆ ISL subscription best practice — While planning the SAN, keep track of how many host and storage pairs utilize the ISLs between domains. As a general best practice, if two switches are connected by ISLs, ensure that there is a minimum of two ISLs between them and that there are no more than six initiator and target pairs per ISL. For example, if 14 initiators access a total of 14 targets between two domains, a total of three ISLs would be necessary. This best practice should not be applied blindly when setting up a configuration. Consider the applications that will use the ISLs.
- ◆ One of the use cases for a two switch fabric is distance extension. In these configurations, it is essential to monitor the ISLs for oversubscription conditions (utilization > 80%) which may lead to back pressure and any errors that are incrementing, especially bit errors or invalid transmission words, as these can lead to credit starvation. For more information on credit starvation, refer

to “Buffer-to-buffer credit information” in the *Extended Distance Technologies TechBook*, available through the on the [E-Lab Navigator, Documents > Topology Resource Center](#). See the individual case studies below for information on how to configure in each environment.

For general information on best practices in all SANs, refer to “[General best practices](#)” on page 19.

### Host and storage layout

Specific information for two switch fabrics follows.

- ◆ In the two switch fabric examples used in this section, hosts and storage can be connected to either switch, but should be connected to the *same* switch when possible. A notable exception to this is in a distance extension environment when the two switches are used to aggregate many different connections over an ISL and provide additional BB\_Credit (buffer-to-buffer credit). In this configuration, the whole point of having two switches is to use the ISL.

For general information on host and storage layout in all SANs, refer to “[Host and storage layout](#)” on page 27.

### Switch and fabric management

Specific information for two switch fabrics follows.

All management applications can be used to monitor this environment.

For general information on switch and fabric management in all SANs, refer to “[Switch and fabric management](#)” on page 29.

### Security

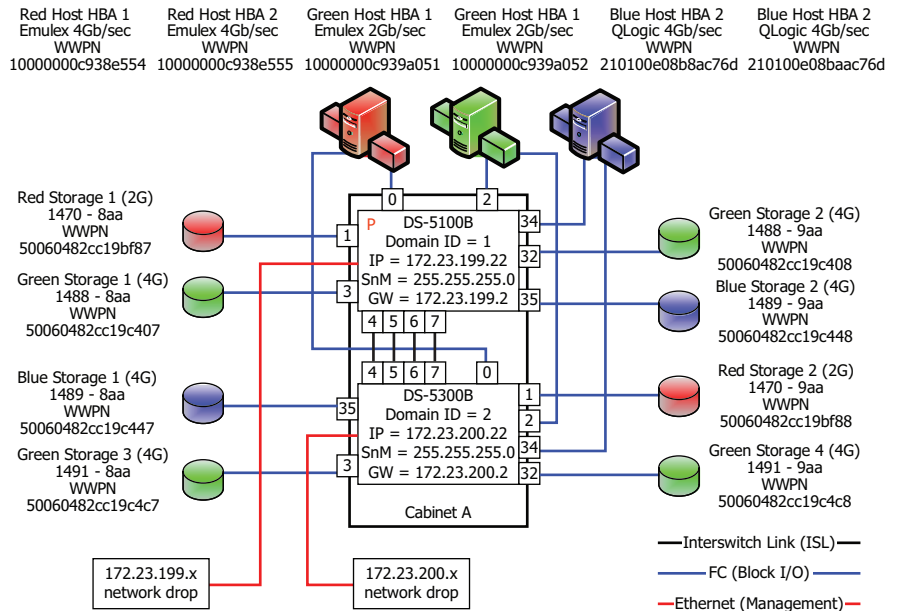
For general information on security in all SANs, refer to “[Security](#)” on page 31.

---

## Connectrix B example

This section contains following information for the Connectrix B example.

**General layout** Figure 5 illustrates a two switch SAN with a DS-5100B and a DS-5300B using CLI to configure the SAN.



GEN-002144

**Figure 5** Two switch SAN with DS-5100B and DS-5300B using CLI to configure SAN

**Best practices**

For general information on best practices for two switch fabrics, refer to “Best practices” on page 46.

Specific information for this example follows.

Even when trunking is not being deployed immediately, it is recommended that you locate ISLs on the same quad or octet to facilitate a smooth transition to a trunking configuration at a later time.

For more Connectrix B specific best practices, refer to “Connectrix B-Series” on page 24.

**Host and storage layout**

Both hosts and storage can be placed anywhere on the SAN. Connectrix B groups ASIC controlled port groups by octet on the DS-5300B. In some circumstances, it may be recommended to connect both host and storage to these groupings. This eliminates frames traveling through the backplane of the switch to reach their destination.



For general information on host and storage layout for two switch fabrics, refer to “[Host and storage layout](#)” on page 27.

### Switch and fabric management

In this example, the CLI (Command Line Interface) is used to create a two switch fabric consisting of a DS-5300B and a DS-5100B. Once the two switch fabric has been connected, three hosts and their associated storage will be attached and properly configured.

For general information on switch and fabric management for two switch fabrics, refer to “[Switch and fabric management](#)” on page 29.

### Security

For general information on security for two switch fabrics, refer to “[Security](#)” on page 31. For more information on security, refer to the *Building Secure SANs TechBook*, on the [E-Lab Navigator](#), **Documents > Topology Resource Center**.

### Setting up this topology

#### Assumptions specific to this case study:

- ◆ The Fibre Channel switches are installed in an EMC-supplied cabinet.
  - For installation instructions, see *Connectrix EC-1500 Cabinet Installation and Setup Manual*, which can be accessed from EMC Online Support at <https://support.emc.com>.
- ◆ Redundant power sources are available.
  - For switch power requirements, refer to the *EMC Connectrix SAN Products Data Reference Manual*, available on the [E-Lab Navigator](#), **Documents > Topology Resource Center**.
- ◆ A laptop is available with the following specifications:
  - Running some version of Windows
  - HyperTerminal is installed
  - Laptop serial ports are DB-9 connections and COM1 will be used to configure the switch IP addresses
- ◆ A serial cable (straight through) and an Ethernet cable (crossover) are available.
- ◆ Network drops, IP addresses, subnet mask, and gateway have been provided by the customer.
- ◆ SFP transceivers and compatible fiber cables are available as required.
- ◆ Access to an FTP server, for backing up (uploading) or downloading the switch configuration is available.

- ◆ License keys have been obtained.
  - Go to the URL listed on the transaction code certificate that shipped with the product to get the license key.
- ◆ The customer has provided a temporary password that will be used as the default password when configuring the IP address.
- ◆ Trunking licenses have been purchased by the Customer for both switches and are available for installation.

### Configure the IP address

To configure the IP address:

---

**Note:** Both the DS-5300B and DS-5100B use an RJ-45 connector for the serial port. Because of this, an RJ-45 serial cable (10-ft (3 m) long) is shipped with switch. Use these to configure the IP address on the DS-5300B and DS-5100B.

---

1. Attach the RJ-45 serial cable between the serial port on the DS-5100B and the RS-232 serial port on the management PC (COM1).
2. Power up the switch by connecting the power cords to the power receptacles provided by the customer. Make sure that all power cords are connected to the switch for maximum redundancy.

---

**Note:** Both power supplies must be connected in order to bring the switch online.

---

3. Configure HyperTerminal:
  - a. Open HyperTerminal by clicking on the **Start / Programs / Accessories / Communications / HyperTerminal**.  
The **Connection Description** dialog box displays.
  - b. Type a descriptive director name in the **Name** field and click **OK**.  
The **Connect To** dialog box displays.
  - c. Ensure the **Connect using** field displays **COM1** and click **OK**.  
The **COM1 Properties** dialog box displays.
  - d. Ensure the following port settings parameters have been properly selected:

- Bits per second: 9600
  - Databits: 8
  - Parity: None
  - Stop bits: 1
  - Flow control: None
- e. Press **Return** to get a prompt.
4. Log in to the switch using the default values: Username: *admin* and Password: *password*.

### IMPORTANT

**It is strongly recommended that when prompted to change the password, you change it to the password that was provided by the customer. This can also be done using the `passwd` command from the command prompt at any time.**

5. At the prompt, enter **ipaddrset** and press **Return**.
6. When prompted, supply the IP address (**172.23.199.22**), subnet mask (**255.255.255.0**) and gateway address (**172.23.199.2**). Enter **Y** for all other values to leave them at their defaults.

**Note:** The Fibre Channel addresses and DHCP are not used for this example.

```
switch:admin> ipaddrset
Ethernet IP address [10.77.77.77]:10.32.53.47
Ethernet Subnetmask [255.0.0.0]:255.255.240.0
Fibre Channel IP address [0.0.0.0]:
Fibre Channel Subnetmask [0.0.0.0]:
Gateway IP address [0.0.0.0]:10.32.48.1
IP address is being changed...Done.
Committing configuration...Done.
```

7. Verify IP address change using **ipaddrshow**.

```
switch:admin> ipaddrshow
SWITCH
Ethernet IP address: 10.32.53.47
Ethernet Subnetmask: 255.255.240.0
Fibre Channel IP address: none
Fibre Channel Subnetmask: none
Gateway IP address: 10.32.48.1
DHCP: Off
```

8. Power down the switch and disconnect the serial cable.

9. Connect the switch to the 10/100BaseT Ethernet connection for the 172.23.199.x network which was provided by the customer.
10. Power up the switch.  
The switch can now be accessed via an IP-based management tool.
11. Repeat [Step 1](#) through [Step 10](#) on the DS-5300B using an IP address of **172.23.200.22**, Subnet mask of **255.255.255.0**, and a Gateway of **172.23.200.2**.

---

**Note:** The RJ-45 serial cable and DB-9 adapter will need to be used to configure the switch IP address. Connect the DB-9/RJ-45 adapter to COM1 and the RJ-45 serial cable between the adapter and the switch.

---

### Configure FC switches

To configure FC switches:

1. Set the switch name for the DS-5100B.
  - a. From the switch prompt, enter **switchname DS-5100B**.  

---

**Note:** The following configurations need to be done with the switch *disabled*.
2. Configure the fabric parameters.
  - a. From the switch prompt, enter **switchdisable** to disable the switch.
  - b. From the switch prompt, enter **configure** to enter the configuration parameter menu.
  - c. Enter **Y** at the **Fabric Parameters** prompt.
  - d. Enter **1** for desired Domain ID at the Domain prompt and press **Enter**.
  - e. The R\_A\_TOV should be automatically set to **10000**. If it is not, enter **10000** at the prompt and press **Enter**.
  - f. The E\_D\_TOV should be automatically set to **2000**. If it is not, enter **2000** at the prompt and press **Enter**.
  - g. Accept the following defaults for the rest of the fields under the **Fabric Parameters** menu by pressing **Enter** after each prompt:

- WAN\_TOV = 0
- MAX\_HOPS = 7
- Data field size = 2112
- Sequence Level Switching = 0
- Disable Device Probing = 0
- Suppress Class F Traffic = 0
- Switch PID Format = 1
- Per-frame Route Priority = 0
- Long Distance Fabric = 0
- BB\_Credit = 16

---

**Note:** For this case study, there is no long distance between the DS-5300B switches. The ISLs connecting the two are less than 10 km.

---

- h. At the **Insistent Domain ID Mode** prompt, enter **y** to accept the **Insistent Domain ID** setting.

---

**Note:** When this mode is set, the switch attempts to acquire the domain number programmed in its **Switch Fabric Settings** from the fabric.

---

- i. Accept the default values from the remaining **Fabric Parameter Menu** items by pressing **Enter** after each prompt:

- Virtual Channel parameters (yes, y, no, n): **[no]**
- F\_Port login 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]**
- ssl attributes (yes, y, no, n): **[no]**
- http attributes (yes, y, no, n): **[no]**
- snmp attributes (yes, y, no, n): **[no]**
- rpcd attributes (yes, y, no, n): **[no]**
- cfgload attributes (yes, y, no, n): **[no]**
- web tools attributes (yes, y, no, n): **[no]**

3. Install the necessary licenses.

---

**Note:** In this case, the only license to be installed will be the trunking license.

---

- a. From the switch prompt, enter **licenseadd** *<license for this switch>*. Example:

```
switch:admin> licenseadd byeSRbdzyQkzfTS0
```

---

**Note:** To enable trunking on the switch after unlocking the license, you need to re-initialize the ports. To re-initialize the ports, you can either disable, and then re-enable, the switch or disable, and then re-enable, the affected ISL ports. Since the switch is already disabled and will be re-enabled in the next step, the explicit disabling and re-enabling of the switch and/or switch ports will not be necessary.

---

4. From the switch prompt, enter **switchenable** to enable the switch.
5. Repeat [Step 1](#) through [Step 4](#) for the DS-5300B using the switch name of DS-5300B and **2** for the desired Domain ID.

### Verify the firmware version loaded on the switches

Run the **firmwareshow** command to verify the firmware version.

```
switch:admin> firmwareshow  
Primary partition: v5.2.1a  
Secondary Partition: v5.2.1a  
switch:admin>
```

## Verify port settings

Use command **PortCfgShow** to show current configuration of a port. All ports should be set to Auto-negotiate speed and the port type should not be locked to either L\_Port or G\_Port. See [Figure 6](#).

Ports of Slot	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Speed	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN	AN
Trunk Port	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON
Long Distance	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--
UC Link Init	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--
Locked L_Port	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--
Locked G_Port	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--
Disabled E_Port	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--
ISL R_RDY Mode	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--
RSCN Suppressed	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--
Persistent Disable	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--
NPIU capability	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON	ON
Mirror Port	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--

Figure 6 Port settings

## Set the switch date and time

**Note:** To assist with the review of support logs should the need arise, it is recommended that you sync up switch time to real time, ideally via an NTP server. You can synchronize the local time of the principal or primary fabric configuration server (FCS) switch to that of an external Network Time Protocol (NTP) server. In this example, the date and time will be set manually.

To set the date and time of a switch manually:

1. Using Telnet, log in to the switch as admin.
2. Enter the date command at the command line using the following syntax:

```
date "MMDDhhmm[CC]YY"
```

where:

- MM is the month (01-12)
- DD is the date (01-31)
- hh is the hour (00-23)
- mm is minutes (00-59)
- CC is the century (19-20)
- YY is the year (00-99)

---

**Note:** Year values greater than 69 are interpreted as 1970-1999; year values less than 70 are interpreted as 2000-2069. The date function does not support Daylight Savings Time or time zones, so changes will have to be reset manually.

---

Example:

```
switch:admin> date
Fri May 5 21:50:00 UTC 1989
switch:admin>
switch:admin> date "0624165203"
Tue Jun 24 16:52:30 UTC 2003
switch:admin>
```

### Install SFP transceivers and connect cables

To install SFP transceivers and connect the cables:

1. Install the SFP transceivers in the Fibre Channel ports.

---

**Note:** The transceivers are keyed to ensure correct orientation. If a transceiver does not install easily, ensure that it is correctly oriented.

---

2. Starting with the ISLs, connect the fiber cables one at a time, verifying the login status of each as they are attached. If a connectivity issue is encountered during this phase, it is easier to troubleshoot it now rather than after all cables have been attached.
  - a. Connect ISLs between ports 4-7 as shown in [Figure 5 on page 48](#).
3. Connect host and storage ports.
  - a. Attach fiber cable between switches and N\_Ports, as shown in [Figure 5 on page 48](#).
  - b. Verify port login status using the **switchshow** and **nsshow** commands.

### Zone hosts and storage

To zone hosts and storage:

1. Create zones using the **zonecreate** commands below:

```
zonecreate "RedHBA1_1470_8aa", "10:00:00:00:c9:38:e5:54;
50:06:04:82:cc:19:bf:87"
zonecreate "RedHBA2_1470_9aa", "10:00:00:00:c9:38:e5:55;
50:06:04:82:cc:19:bf:88"
```



```

zonecreate "BlueHBA1_1489_8aa", "21:01:00:e0:8b:8a:c7:6d;
50:06:04:82:cc:19:c4:47"
zonecreate "BlueHBA2_1489_9aa", "21:01:00:e0:8b:aa:c7:6d;
50:06:04:82:cc:19:c4:48"
zonecreate "GreenHBA1_AllGreenStorage", "10:00:00:00:c9:39:e5:51;
50:06:04:82:cc:19:c4:07; 50:06:04:82:cc:19:c4:08; 50:06:04:82:cc:19:c4:c7;
50:06:04:82:cc:19:c4:c8"
zonecreate "GreenHBA2_AllGreenStorage", "10:00:00:00:c9:39:e5:52;
50:06:04:82:cc:19:c4:07; 50:06:04:82:cc:19:c4:08; 50:06:04:82:cc:19:c4:c7;
50:06:04:82:cc:19:c4:c8"

```

2. Create the configuration by using the **cfgcreate** command.

```

cfgcreate "Oct_31_06_1140" , "RedHBA1_1470_8aa; RedHBA2_1470_9aa;
BlueHBA1_1489_8aa; BlueHBA2_1489_9aa; GreenHBA1_AllGreenStorage;
GreenHBA2_AllGreenStorage"

```

3. Enable the configuration by using the **cfgenable** command.

```

cfgenable "Oct_31_06_1140"

```

4. Enter **Y** at the confirmation prompt.
5. Enter **cfgshow** to display zoning info.

When completed, the zone information should be similar to what is shown below.

#### Defined configuration:

```

cfg: Oct_31_06_1140
  RedHBA1_1470_8aa; RedHBA2_1470_9aa; BlueHBA1_1489_8aa; BlueHBA2_1489_9aa;
  GreenHBA1_AllGreenStorage; GreenHBA2_AllGreenStorage"
zone: RedHBA1_1470_8aa
  10000000c938e554; 50060482cc19bf87
zone: RedHBA2_1470_9aa
  10000000c938e555; 50060482cc19bf88
zone: BlueHBA1_1489_8aa
  210100e08b8ac76d; 50060482cc19c447
zone: BlueHBA2_1489_9aa
  210100e08baac76d; 50060482cc19c448
zone: GreenHBA1_AllGreenStorage
  10000000c939a051; 50060482cc19c407;
  50060482cc19c408; 50060482cc19c4c7;
  50060482cc19c4c8
zone: GreenHBA2_AllGreenStorage
  10000000c939a052; 50060482cc19c407;
  50060482cc19c408; 50060482cc19c4c7;
  50060482cc19c4c8

```

#### Effective configuration:

```

CFG: Oct_31_06_1140
Zone: RedHBA1_1470_8aa

```

```

10000000c938e554
50060482cc19bf87
Zone: RedHBA2_1470_9aa
10000000c938e555
50060482cc19bf88
Zone: BlueHBA1_1489_8aa
210100e08b8ac76d
50060482cc19c447
Zone: BlueHBA2_1489_9aa
210100e08baac76d
50060482cc19c448
Zone: GreenHBA1_AllGreenStorage
10000000c939a051
50060482cc19c407
50060482cc19c408
50060482cc19c4c7
50060482cc19c4c8
Zone name = "GreenHBA2_AllGreenStorage
10000000c939a052
50060482cc19c407
50060482cc19c408
50060482cc19c4c7
50060482cc19c4c8

```

### Save configuration

In case the configuration is lost, or unintentional changes are made, keep a backup copy of the configuration file on a host computer.

To upload a configuration file:

1. Verify that the FTP service is running on the host computer. The host must have an FTP server application running.
2. Connect to the switch through the Telnet and log in as admin.
3. Enter the **configUpload** command.

The command becomes interactive and you are prompted for the required information.

Example:

```

switch:admin> configupload
Protocol (scp or ftp) [ftp]: ftp
Server Name or IP address [host]: 192.1.2.3
User Name [user]: JohnDoe
File Name [config.txt]: /pub/configurations/config.txt
Password: xxxxxx
configUpload complete: All config parameters are uploaded.
switch:admin>

```

## Complete the SAN setup

At this point the SAN is ready to pass I/O from host to storage. Other steps, such as configuring LUN Masking and modification of host configuration files, are required before the SAN setup is complete. Refer to the OS configuration guide for more details.

## Connectrix MDS example

**Note:** VSANs will be configured and used in this example. Refer to *EMC Connectrix SAN Products Data Reference Manual*, available on the [E-Lab Navigator, Documents > Topology Resource Center](#), for more information on the Connectrix MDS VSAN feature.

**General layout** Figure 7 shows two MDS 9506s using VSANs.




Figure 7 Two Connectrix MDS 9506s using VSANs

**Best practices** For general information on best practices for two switch fabrics, refer to “[Best practices](#)” on page 46. Specific information for this example follows.

By default thresholds are set to 80% utilization.

**Host and storage layout** For general information on host and storage layout for two switch fabrics, refer to “[Host and storage layout](#)” on page 47. Specific information for this example follows.

Line Rate Mode cards have no special restrictions. Over-subscribed cards should be used for hosts only.

**Switch and fabric management** For general information on switch and fabric management for two switch fabrics, refer to “[Switch and fabric management](#)” on page 47. Specific information for this example follows.

Cisco Fabric Manager may be used for management.

**Security** For general information on security for two switch fabrics, refer to “[Security](#)” on page 47. Specific information for this example follows.

Use Switch Binding and Port Binding for security.

### Setting up this topology

#### Assumptions specific to this case study:

- ◆ The switches are installed in an EMC-supplied cabinet.
  - For installation instructions, see *Connectrix EC-1500 Cabinet Installation and Setup Manual*, which can be accessed from EMC Online Support at <https://support.emc.com>.
- ◆ The proper power receptacles have been provided by the customer.
  - For switch power requirements, refer to the *EMC Connectrix SAN Products Data Reference Manual*, available on the [E-Lab Navigator, Documents > Topology Resource Center](#).
  - For Cabinet power requirements, refer to *Connectrix EC-1500 Cabinet Installation and Setup Manual* which can be accessed from EMC Online Support at <https://support.emc.com>.
- ◆ The switches have *not* been connected to the power source and are *not* powered on.
- ◆ Network drops, IP addresses, subnet mask, and gateway have been provided by the customer.

For switch or cabinet network requirements, refer to the *EMC Connectrix SAN Products Data Reference Manual*, available through the on the [E-Lab Navigator, Documents > Topology Resource Center](#).

---

**Note:** Connectrix MDS switches can be placed on either a public or private network. There are advantages to both configurations. For more information, refer to "Public versus private" in the *Networked Storage Concepts and Protocols TechBook*, available on the [E-Lab Navigator, Documents > Topology Resource Center](#).

In this example, it is assumed that the customer has provided two Ethernet cables and that one of them is on the 172.23.199.x network and that the other is connected to the 172.23.200.x network.

---

- ◆ The correct number of line cards have been installed into each chassis. In this case, two line cards in each chassis are required and installed in slots 1 and 2.
  - For help in determining how many ports are required, refer to "Determining customer requirements" in the *Networked Storage Concepts and Protocols TechBook*, available on the [E-Lab Navigator, Documents > Topology Resource Center](#).
- ◆ License keys have been obtained.
  - Go to the URL listed on the transaction code certificate that shipped with the product to get the license key.
- ◆ Use the laptop supplied by the installer to configure the IP addresses of the switches; this laptop has a serial DB-9 connector.
- ◆ Use the temporary password provided by the customer as the default password when configuring the IP address.
- ◆ Use Fabric Manager for VSAN setup.

### Configure the IP address

To configure the IP address:

1. Power up the cabinet by connecting the power cords to the power receptacles provided by the customer.
2. Select one of the switches to configure and set the IP to 172.23.199.22.
3. Supply a network connection to the appropriate subnet.

4. Connect to the serial port of the switch using an RS232 serial cable, with a baud rate of 9600, 8 data bits, no parity, 1 stop bit and no flow control.

The **login** prompt should display.

5. Log in the first time with username *admin* and password *admin*.  
You should be prompted to supply a new strong password for CLI user *admin*.
6. For this example, select **no** when asked if you want to run setup.

---

**Note:** This example will start with the switch that will have a Domain ID of **1** and an IP address of **172.23.199.22**.

---

7. Repeat above steps for each switch, supplying the appropriate IP.

### CLI commands to configure the IP and gateway

- ◆ Switch# *config terminal*

Enter configuration commands, one per line.

```
Switch(config)# interface mgmt 0
Switch(config-if)# IP address 172.23.199.22 255.255.255.0
End with CNTL/Z.
```

- ◆ Switch# *config terminal*

Enter configuration commands, one per line.

```
Switch(config)# ip default-gateway 172.23.199.2
End with CNTL/Z.
```

To authorize access on a switch for Device and Fabric Manager, run this command on every switch while supplying a username (nnn) and password (ppp):

- ◆ Switch#*conf t*

```
Switch(config)# snmp-server user nnn network-admin auth md5
ppp
Switch(config)#end
Switch# copy running-config startup-config
Switch# exit
```

### Installing Fabric Manager and Device Manager

To install Fabric Manager and Device Manager:

1. Open your web browser.

2. Enter the IP address of the switch in the address bar.
3. Follow the prompts and accept all defaults to install both Fabric Manager and Device Manager.

Fabric Manager and Device Manager can be started using the configured snmp-server username and password provided in “[CLI commands to configure the IP and gateway](#)” on page 41.

### Creating a VSAN

To create a VSAN:

1. Open the Device Manager for the switch with an IP address of **172.23.199.22**.
2. Select **FC** from top toolbar.
3. Select **VSAN**.
4. Select **Create VSAN**.
5. Enter the value of **100** into the **VSAN ID** field.
6. Set the **VSAN Name** to be “**Red\_VSAN\_100**”.
7. Use the default interop mode.
8. Click **Create**.
9. Configure the other VSANs in this physical switch.

*Example 1: IP address  
172.23.199.22*

For the switch with the IP address **172.23.199.22**:

- a. Repeat [Step 2](#) (beginning on [page 63](#)) through [Step 8](#) for VSAN 200 and 300.
  - For Virtual switch 200, use a VSAN name of “**Green\_VSAN\_200**”.
  - For Virtual switch 300, use a VSAN name of “**Blue\_VSAN\_300**”.
- b. Following the tables below, assign and enable the ports to the proper VSAN using Device Manager.

Slot #	Port #	Name	VSAN ID
1	1	TE ISL to SW 2	1
1	2	Red Host HBA 1	100
1	3	Red Storage 1	100
1	4	Green Host HBA 1	200
1	5	Green Storage 1	200

Slot #	Port #	Name	VSAN ID
1	6	Blue Host HBA 1	300
1	7		
1	8		
2	1	TE ISL to SW 2	1
2	2	Green Storage 2	200
2	3	Blue Host HBA 1	300
2	4	Blue Storage 2	300
2	5		
2	6		
2	7		
2	8		

*Example 2: IP address  
172.23.200.22*

For the switch with the IP address **172.23.200.22**:

- a. Repeat [Step 2](#) (beginning on [page 63](#)) through [Step 8](#) for VSAN 100, 200 and 300 on this switch.
  - For VSAN 100 use a Name “**Red\_VSAN\_100**”.
  - For VSAN 200 use a VS Name “**Green\_VSAN\_200**”.
  - For VSAN 300, use a VS Name “**Blue\_VSAN\_300**”.
- b. Assign the ports to the proper VSAN using Device Manager, using the following table:

Slot #	Port #	Name	VSAN ID
1	1	TE ISL to SW 1	1
1	2	Red Host HBA 2	100
1	3	Blue Storage 1	300
1	4	Green Storage 1	200
1	5		
1	6		
1	7		
1	8		
2	1	TE ISL to SW 1	1
2	2	Green HBA 2	200
2	3	Red Storage 2	100
2	4	Blue Host HBA 2	300
2	5	Green Storage 4	200
2	6		
2	7		
2	8		



## Connecting cables

To connect the cables:

1. Connect ISLs.
  - a. Attach fiber cable between switches as shown in [Figure 7 on page 59](#).
  - b. After all cables are connected, use Fabric Manager to verify that all ISL connections are up.
  - c. Re-arrange icons to accurately reflect the switch configuration.
2. Connect host and storage ports.
  - a. Attach fiber cable between switches and N\_Ports.

## Configure domains

To configure domains:

1. Open **Fabric Manager**. It should show a topology of two switches.
2. Open the “**Red\_VSAN\_100**” folder from **Fabric Manager**.
3. Select **Domain Manager**.
4. Select the **Configuration** menu.
5. Set a Domain ID. **1** for switch **172.23.199.22** and **2** for **172.23.200.22**.
6. Set the priority to **1** in the domain menu, to set a principal switch.
7. Repeat [Step 2](#) through [Step 6](#) for “**Green\_VSAN\_200**” and “**Blue\_VSAN\_300**”.

## Zone hosts and storage

To zone hosts and storage:

1. From **Fabric Manager** select “**Red Vsan 100**”.
2. Select **Edit Full Zone Database**.
3. Create a zone by clicking **Zone** button under the **Zones Tree**.
4. Provide a descriptive name for the zone. This example will zone “Red host HBA 1” and “Red Storage 1”, so “**RedHBA1\_1470\_8aa**” will be used. Press **Enter**.

5. Locate, then click, "**Red Host HBA 1**" (WWPN 10000000c938e554) in the **Potential zone members list**.
6. Click the right-pointing arrow on the divider between the **Potential members list** and the **zones list** to add the HBA to the zone.
7. Locate, then click, "**Red Storage 1**" (WWPN 50060482cc19bf87) in the **Potential zone members list**.
8. Click the right-pointing arrow on the divider between the **Potential members list** and the **zones list** to add the Storage port to the zone.
9. Repeat [Step 2](#) through [Step 7](#) for all host and storage pairs in the environment.
10. Create a zone set by clicking **New Set** under the **Zone sets Tree**.
11. Provide a descriptive name for the zone set. This example uses the name "**RED Fabric 1**".

Add only those zones that will be necessary on Red\_Fabric\_1. In this case only the zone named "RedHBA1\_1470\_8aa" should be added to the Red\_Fabric\_1 zone set. Repeat for other Fabrics. When completed, you should have 5 zone sets, as shown below.

```
Zone set name = "Red_Fabric_1"
```

```
Zone name = "RedHBA1_1470_8aa"
Zone Member = "10000000c938e554"
Zone Member = "50060482cc19bf87"
```

```
Zone name = "RedHBA2_1470_9aa"
Zone Member = "10000000c938e555"
Zone Member = "50060482cc19bf88"
```

```
Zone set name = "Blue_Fabric_1"
```

```
Zone name = "BlueHBA1_1489_8aa"
Zone Member = "210100e08b8ac76d"
Zone Member = "50060482cc19c447"
```

```
Zone name = "BlueHBA2_1489_9aa"
Zone Member = "210100e08baac76d"
Zone Member = "50060482cc19c448"
```

```
Zone set name = "Green_Fabric"
```

```
Zone name = "GreenHBA1_AllGreenStorage"
Zone Member = "10000000c939a051"
Zone Member = "50060482cc19c407"
Zone Member = "50060482cc19c408"
Zone Member = "50060482cc19c4c7"
Zone Member = "50060482cc19c4c8"
```

```
Zone name = "GreenHBA2_AllGreenStorage"
Zone Member = "10000000c939a052"
Zone Member = "50060482cc19c407"
Zone Member = "50060482cc19c408"
Zone Member = "50060482cc19c4c7"
Zone Member = "50060482cc19c4c8"
```

### Optional: Configure IVR (Inter-VSAN Routing)

The configuration of VSANs on a fabric allows for security, scalability, and availability. However, this isolation of traffic between VSANs prevents users from accessing resources, such as tape libraries, located in other VSANs. The solution to this limitation is Cisco's Inter-VSAN Routing feature, which allows initiators in one VSAN to access targets in other VSANs without merging the VSANs. Perhaps a host in the Red VSAN needs to access storage in the Blue VSAN. Configuring IVR zones and IVR zone sets containing the allowed initiators and targets allows communication between these resources.

1. In Fabric Manager
  - a. Click the **Zone** tab in upper tool bar.
  - b. Click the **IVR** tab.
  - c. Select **Wizard**.
2. Select the VSANs that will participate in IVR in the fabric.  
Select **VSAN 100, 200** and **300**.
3. Select the end devices that you want to communicate over IVR.  
Select the following:
  - VSAN 100: 10000000c938e554**
  - VSAN 200: 50060482cc19c407**
  - VSAN300: 50060482cc19c447**
  - VSAN 300: 210100e08b8ac76d**
4. Enter the VSAN ID of the VSAN you want to use as the transit VSAN between the VSANs selected for the IVR zone.

Select **VSAN 1** as the transit VSAN.

---

**Note:** VSAN 1 connects both switches and all trunking traffic will pass over this link to communicate with VSANs in other switches.

---

5. Set the IVR zone and IVR zone set.

IVR NAME = **IVRZONE1**

IVR ZONENET NAME = **IVRZONESET1**

6. Verify all steps that Fabric Manager will take to configure IVR in the fabric.
7. Click **Finish** if you want to enable IVR NAT and IVR topology and to create the associated IVR zones and IVR zone set.

or

Click **Cancel** to exit the IVR Wizard without saving any changes.

8. The **Save Configuration** dialog box displays. You can save the configuration of the master to be copied to other IVR-enabled switches. Click either **Continue Activation** or **Cancel**.
9. Click **Finish**.

### **Complete the SAN setup**

At this point the SAN is ready to pass I/O from host to storage. Other steps, such as configuring LUN masking and modification of host configuration files, are required before the SAN setup is complete. Refer to the *Cisco MDS 9000 Family Fabric Manager Configuration Guide* for more details.

## Blade switch with direct attached storage

A blade server with two embedded FC switch modules directly attached to storage complies with the definition of a simple two-switch, one-hop SAN model. Before discussing the general layout of this topology type, it is interesting to note how the use of a blade server can simplify a desired two switch-one hop SAN design.

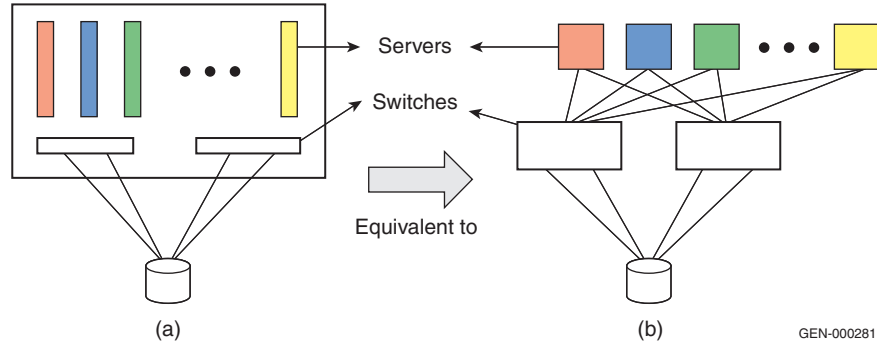
### Consider the following scenario:

An end-user needs to hook up ten independent servers to storage using two departmental switches. Each of the ten servers needs its own power supply, cooling mechanism, and cabling. These aspects must also be considered for the departmental switches. This whole two-tier set up can be replaced by a 7U blade server chassis which can house at least ten independent servers, depending on the vendor type (IBM, Dell, HP, and so on). The chassis contains the power, cooling, and cabling components, and often incorporates a network switch. The blades within a single chassis can still run different applications and play independent roles.

The chassis also provides the ability to embed a pair of FC switch modules with a maximum of eight external ports, depending on the vendor (Brocade or Cisco), which can be attached to an external fabric, or in this case, the storage. Each switch module is internally connected to each of the blade servers in the blade server chassis. The behavior, features, and management of the switch modules are similar to those of departmental stand-alone edge switches.

The detailed blade server concept, its value to the FC SAN world, the basic architecture, and the various EMC-supported blade switch modules are further discussed in the *Fibre Channel over Ethernet (FCoE) Data Center Bridging (DCB) Concepts and Protocols TechBook*, available on the [E-Lab Navigator](#), **Documents > Topology Resource Center**.

With this in mind, this example examines the general layout of a blade server switch module directly attached to storage. [Figure 8 on page 70](#) is divided into two parts. The left side (a) represents the blade server chassis directly attached to storage through the switch modules, while the right side (b) represents the two-tier FC SAN to which it is equivalent.



**Figure 8 (a) Blade server direct attached to storage (b) two-tier FC SAN**

This design occupies about 7U of cabinet space, and offers fully redundant fabrics. Most switch modules today have 4 G ports. Considering a switch module has at the most 6 ports (depending on the vendor type), a blade server is capable of up to 48 GB/s throughput.

## General steps to set up a blade server

To set up a blade server:

1. Configure the management module.
  - a. Physically install the blade server chassis and all its components in a customer-supplied rack as per the vendor-specific blade server hardware setup documentation.
  - b. Power up the switch by connecting the power cords to the power receptacles provided by the customer.
  - c. After powering it up, it is essential to configure the management module. The management module is generally located on the back face of the blade server chassis and has a serial console port and an Ethernet port.

**Note:** At this stage, refer to the “Best practices” on page 73 to study the two options available to configure the management modules and then select the preferred option between the two.

It is also advisable to refer to the specific blade server vendor documentation to obtain the configuration details for assigning an IP to the management module.

There are also different options to consider while doing this at data centers (for example, whether it is advisable to have the payload LAN and management LAN on the same LAN segment, and so on).

d. In general, if configuring the management module through the *serial console* using the *Hyperterminal* application, it is essential to gather the following information:

- Connection type: COM 1 or COM 3
- Other settings:
  - Bits per second
  - Data bits
  - Parity bits
  - Stop bits

If configuring the management through the *Ethernet port*, that is, through Telnet, the following information is required:

- Default IP of management module
- Default username
- Default password

e. Once configured, the management module can be used to configure the other modules including the I/O modules.

2. Configure the other modules and the server blades.

a. To configure other modules:

Once the management module is configured with an IP address, it can be accessed through CLI or a Web interface in most cases, and can be used to configure, or rather assign, IPs to the other I/O modules, such as fibre channel switch modules (FCSM), and the Ethernet module if required.

b. To configure the server blades:

Generally, when supplied from the vendor, the server blades are configured for a network internal to the blade server chassis. Most blade servers today have VGA and USB connectors on the server blade which can be used to access and configure the server blades. The desired OS can be installed on the blade using these means of access, and on completing this installation, the other features, such as IP address, etc., can be set up for the server blades.

3. Check the following switch module configuration settings:
  - a. The mode of operation: The switch modules must be in their respective native mode for this kind of a configuration.
    - Desired mode on Brocade switch module: Brocade native mode (interopmode 0)
  - b. Domain ID settings:
    - Domain ID range for Brocade switch modules: 1 – 239
  - c. Ensure that the switch firmware revision and the switch management application are supported versions.
  - d. If default zoning enabled does not allow the host ports to communicate with the storage ports, then the appropriate WWNN or WWPN zoning needs to be performed to establish connectivity. As a result, the server blades can log into the storage and read from or write to the respective storage devices it is mapped to.
    - Default zoning on the Brocade module restricts the initiator (host) ports to see or communicate with the target (storage) ports. Therefore, zoning must be configured on the Brocade modules.

4. Ensure switch module connectivity (physical).

The switch modules must be connected directly to the storage using FC cables. EMC recommends at least two connections between any two FC components in a SAN. In this case, at least four cables must be connected in total: two from each switch module to the respective storage ports.

The switch modules, as previously stated, are internally connected to each of the blade servers. Thus, a specific number of ports on each switch, depending on the number of server blades the blade server chassis can house, are dedicated F\_ports, while the other external ports are G\_ports. For this topology, these ports can be configured as F\_ports since they will be hooked up to storage.

5. Use CLI/Web GUI to verify connectivity.

The switch module CLI or switch management application can be used to verify that the physical connectivity was successful and that the switch can see the storage ports. The name server



information can be obtained at this time so as to confirm whether all the N\_Ports: HBA ports from the server blades and the storage ports, show up.

---

## Best practices

For general best practices, refer to [“General best practices” on page 19](#). The following best practices are recommended for configuring the blade server and setting up the direct attach to storage:

- ◆ There are two options available to access the management functions of the FC switch modules through the link to the management module. Option 1 is recommended for easy access and management of the switch module.
  - Option 1: If the IP address assigned to the switch is within the subnet of the management module, then the switch management functions can be accessed without launching the Management Module GUI.
  - Option 2: If the IP address assigned to the switch is not within the subnet of the management module, then the switch management functions can be accessed only by launching the Management Module GUI and then navigating to the specific functions.
- ◆ Connect each switch module to storage ports that are mapped to the same set of storage devices. Thus, each server blade can access the storage ports it is zoned with using two paths, each path passing through a different switch module. If either of the switch modules goes down, the server continues to have access to the storage port it is zoned with and there is no disruption in traffic between the server and storage.
- ◆ Always check supported switch module firmware release notes for caveats on direct attach to storage, if any.

---

## Host and storage layout

For general host and storage layout information, refer to [“Host and storage layout” on page 27](#). The following information is specific to blade servers.

There is little an end-user can do with the host and storage layout in the case of a blade server since the architecture allows every server

blade in the blade server chassis to access both the switch modules. As a result, even if a storage port is hooked up to one of the switch modules, all the server blades can access the storage port through the switch module unless, and until, any of the internal or external ports on the switch modules are blocked to prevent access between the server and storage in this case. However, zoning can take care of the access settings.

---

### Switch and fabric management

For general switch and fabric information, refer to [“Switch and fabric management” on page 29](#). The following information is specific to blade switches.

It is recommended that end-users use the supported vendor specific web application and CLI for managing the switch.

- ◆ The Brocade modules can be managed using Brocade CLI or Web Tools.
- ◆ The Brocade M series modules can be managed using Brocade M series SAN browser.

---

### Security

For general information on security, refer to [“Security” on page 31](#). Specific information for this example follows. For further information on security, refer to the *Building Secure SANs TechBook*, available on the [E-Lab Navigator](#), **Documents > Topology Resource Center**.

- ◆ Blade servers usually exist in the “access layer” of a data center architecture and have specific requirements for security. These mainly deal with establishing secure access to switch, protection against attacks, and identifying users and other servers that are accessing the network. For easy configuration, management, and control, many capabilities, such as firewalls, are consolidated in the embedded switches in this example.
- ◆ The access to the embedded switch must be secured. This requires specific features, such as SSHv2 or RADIUS. Both protocols are means of securing connectivity to access the management interface of the switch. Without these capabilities, it would be considerably easier for an unauthorized user to change the configuration of the switch thus affecting the blade server’s performance.

- ◆ The switch must also provide capabilities to prevent, as well as defend against network attacks. This is enabled by the following features: Port security (limited number of MAC addresses per port), MAC address notification (whether MAC address has moved), DARP inspection (ties a port to an ARP request, helping to ensure that a default gateway cannot be spoofed), IP source guard (Protects against IP being spoofed).
- ◆ An access layer switch must make sure that the network allows only authorized devices to connect to it.

Security features on blade server FC switch modules include:

- ◆ ISL and ELS authentication, as defined in FC-SP, provide a means to authenticate the identity of a connected switch, host or target and/or authorize a list of devices to join a fabric.
- ◆ ISL security is supported on E\_Ports. ELS security is supported on F\_Ports.
- ◆ Fabric Binding, Port Binding, and Switch Binding are introduced as a means to control the switch composition of a fabric. They may be enabled on the respective vendor switch modules using any of the switch management applications, just as they are for equivalent standalone switches.
- ◆ Security configuration management is similar to zoning configuration management.

## IBM Brocade example

### General layout

This example, shown in [Figure 9](#), will show an IBM Brocade 4 GB SAN switch module (32R1812) directly attached to storage.

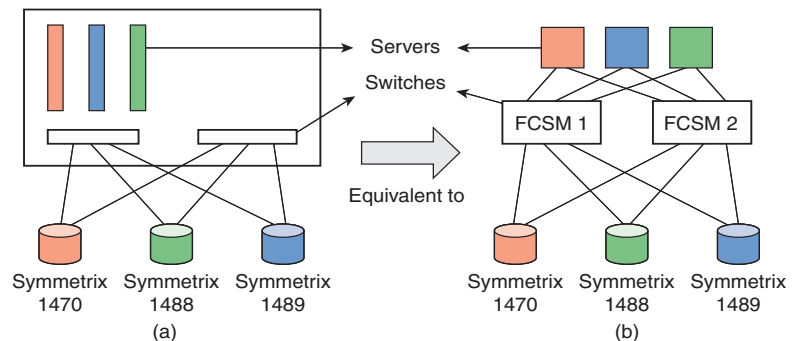
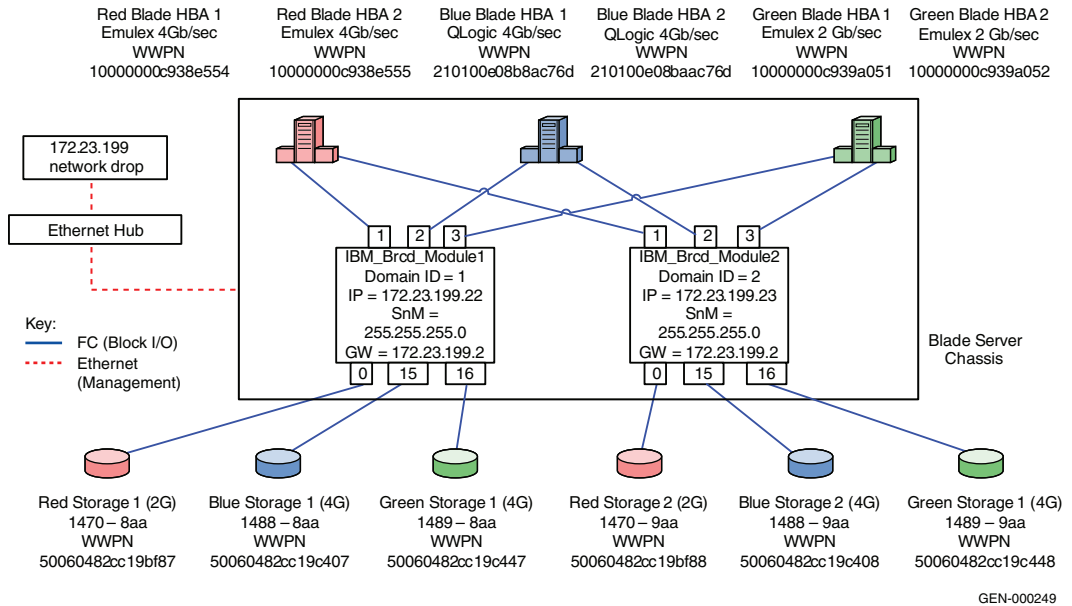


Figure 9

IBM blade server directly attached to storage

Figure 10 represents an IBM blade server chassis with three server blades placed in slots 1, 2, and 3 of the ten slot chassis. Each of the server blades has a dual port HBA. One of the HBA ports on each server blade is internally connected to one of the Brocade switch modules, while the other HBA port on each server blade is connected to the other Brocade switch module.



**Figure 10 IBM Brocade 4 GB SAN switch module (32R1812) directly attached to storage**

Each switch module has 14 internal ports allocated for the 14 server blades the chassis can house. There are 6 external ports on the switch module. In this case, 3 of the 6 ports on each module are hooked up to a port on three different EMC Symmetrix® systems. For this example, assume that Posrt 0, 15, and 16 on Brocade Switch Module 1 (FCSM 1) are hooked to Symmetrix systems 1470 8AA, 1488 8A,A and 1489 8AA, while the Ports 0, 1,5 and 16 on FCSM 2 are hooked to Symmetrix systems 1470 9AA, 1488 9AA, and 1489 9AA, respectively.

**Setting up this topology**

**Assumptions before configuring the components of this topology:**

- ◆ The IBM Blade server chassis and its components have been installed in the cabinet and powered up as per the IBM hardware installation guide for blade servers. (<http://www.redbooks.ibm.com/redbooks/pdfs/sg246342.pdf>)

- ◆ The following components are used for this example:
  - Two Ethernet modules
  - Brocade switch I/O modules
  - Management module
  - Three server blades
- ◆ The IP addresses and operating systems on the server blades can be configured by referring to the following document from IBM: <http://www.redbooks.ibm.com/redbooks/pdfs/sg247313.pdf>.

### Set up an IBM Brocade 4 GB SAN

To set up an IBM Brocade 4 GB SAN switch module direct attached to a Symmetrix system:

1. Configure the Management Module for the IBM Blade server:

The primary setup task for the Management Module would be assigning IP addresses, which are necessary to communicate with the Management Module Web GUI and command line interfaces.

The IBM Management Module has two network interfaces:

- An external interface (eth 0), which is accessible using the 10/100/1000Base T connector on the Management Module.
- An internal interface (eth 1), which is connected to the management interfaces of all the installed I/O modules including the IBM Brocade 4 GB SAN switch module in this case.

---

**Note:** The default static IP address of the external interface of the Management module is 192.168.170.125 with a default subnet mask of 255.255.255.0. The default IP address for the internal interface is statically assigned to be 192.168.70.126.

---

The steps for configuring the IP address on the Management Module are as follows:

- a. Prepare a checklist with configuration details:

Decide on the IP addresses, subnet masks, and gateway addresses to be assigned to the external and internal interfaces of the Management Module. They all must belong to the same subnet. In this example, the IP address on the external interface is changed to 172.23.199.60 and the IP address on the internal interface is changed to 172.23.199.61, which are both on the 199 subnet.

- b. Connect the Management Module to a workstation using a cross over network cable.
- c. Configure a static IP address for the workstation that is in the same subnet as the Management Module default IP addresses. In this case, a static address of 192.168.70.100 with a subnet mask of 255.255.255.0 was used for the workstation. IBM recommends not using addresses in the range of 192.168.70.125 through 192.168.70.130 since they conflict with the default addresses assigned by the management module.
- d. Connect to the Management Module GUI or Web Interface by pointing the Web browser on the workstation to:  
`http://192.168.70.125`.
- e. Enter a valid user ID and password to log in to the Module management interface. The factory default configuration of a Management Module defines a user ID named USERID with a password of PASSWORD.

---

**Note:** The number 0 is between the W and the R in PASSWORD.

---

- f. Select the **Network interfaces** option under the **Management Module (MM) Control** menu.
- g. Enter the desired external and internal IP addresses, subnet masks, and default gateway for the Management Module. In this example:  
  
External IP address: 172.23.199.60  
Subnet mask: 255.255.255.0  
Gateway address: 172.23.199.2  
  
Internal IP address: 172.23.199.61  
Subnet mask: 255.255.255.0  
Gateway address: 172.23.199.2  
  
Click **Save** to store these new IP addresses.
- h. Restart the Management Module.
- i. Pull out the cross over cable from the Ethernet port of the Management Module and connect it with an Ethernet cable to a hub on the 199 subnet.

One can now connect to the Management Module Web interface and Telnet into its Command Line Interface using the 172.23.199.60 IP address assigned to its external network interface.

## 2. Configure the Brocade SAN switch modules.

The primary set up tasks for the Brocade modules are to:

- Assign IP addresses to the Brocade switch (I/O module) Management interfaces.
- Enable the external Ethernet port on the Brocade modules.

---

**Note:** When a new switch I/O module is first installed, the Management Module assigns a default IP address to the management interface of the I/O module. The default IP address is chosen based on the I/O module bay on the back plane of the chassis where the I/O module is installed. Those I/O modules installed in I/O module bays 1, 2, 3, and 4 are assigned IP addresses 192.168.70.127, 192.168.70.128, 192.168.70.129, and 192.168.70.130, respectively.

---

The steps to configure the IP addresses on the Brocade modules are as follows:

### a. Prepare a checklist with configuration details:

Decide on the IP addresses, subnet masks, and gateway addresses to be assigned to the external interfaces of the two Brocade modules. They must belong to the same subnet as the Management module. This example configures the IP address of one Brocade module as 172.23.199.22 and the other as 172.23.199.23.

### b. Connect to the Management Module by pointing the web browser to: <http://172.23.199.60>.

### c. From the Management Module interface select **I/O Module Tasks > Management**.

### d. Select the specific I/O module based on the I/O bay it is pushed into. It is a general practice to place the Brocade modules in I/O bays 3 and 4. The IP address fields are updated as follows:

- For Bay 3: Brocade module 1  
IP address: 172.23.199.22  
Subnet mask: 255.255.255.0  
Gateway address: 172.23.199.2
- For Bay 4: Brocade module 2  
IP address: 172.23.199.23  
Subnet mask: 255.255.255.0  
Gateway address: 172.23.199.2



- e. Click **Save** to activate the new IP address.
- f. Select the **Advance Management** link for each of the Brocade modules, set the **External Ports** field to **Enable**, and click **Save**. Leave everything else at default settings.

At this point both the Brocade switch modules have an IP address. A Brocade Web Tools browser can now be pointed to this IP or one can Telnet into this IP address to use the Brocade CLI to manage the switch.

3. Configure the switches.
  - a. Telnet into one of the Brocade switch modules by issuing a **Telnet 172.23.199.22** command.
  - b. Enter the default username and password of **USERID** and **PASSWORD**.

#### **IMPORTANT**

**It is strongly recommended that when prompted to change the password, you change it to a password that was provided by the customer.**

- c. Verify mode of operation. On issuing an **interopmode** (CLI) command, the switch must return **interopmode: Off** which implies that the switch is running in its native Brocade mode. If not, disable the switch by issuing a **switchdisable** command followed by the **interopmode 0** command. A reboot is required to restore normal configuration on the switch.
- d. Assign a Domain ID: The switch modules automatically take a Domain ID in the range of 1– 239. If a specific Domain ID is desired, as in this case,
  - The Brocade module needs to be disabled by issuing **switchdisable**.
  - A new Domain ID can be assigned by running **configure**.
  - When prompted to choose *yes* or *no* for **Configure.Fabric parameters**, type **y**.
  - For the Domain ID setting, type **1**.
  - To have Domain ID fixed to **1**, press **Enter** until you reach the **Insistent Domain ID Mode** field and type **y**.
  - Press **Enter** for the rest of the values to accept default settings.

- e. Assign a switch name to the Brocade module by issuing the **switchname** *IBM\_brcd\_module1* command.
- f. Ensure that a supported firmware version is running on the switch module by running **version** on the switch CLI. If not, download the latest firmware version supported by EMC and IBM.
- g. Verify that no zoning is enabled or is active on the switch modules by running a **cfgshow** command.
- h. Configure the ports:

To configure port speed to auto negotiate issue the **portcfgspeed** *<portnumber>* **0** command.

**Note:** Outside of a hardware requirement or if a known problem with auto-negotiation exists between two port types. We recommend that you leave ports at auto-negotiate for both port type and speed.

To name the port, issue a **portName** *<portnumber>* [**name**] command. For example, in this case the internal ports on FCSM1 connected to the blade server HBAs can be named as "Red Host HBA 1", "Green Host HBA 1", and "Blue Host HBA 1" as specified in the following tables:

#### Switch 1, FCSM 1:

Port #	Symbolic port name	Port type	Port speed
0	Red Storage 1	F_port	AutoNeg.
1	Red Blade HBA 1	F_port	AutoNeg.
2	Blue Blade HBA 1	F_port	AutoNeg.
3	Green Blade HBA 1	F_port	AutoNeg.
15	Blue Storage 1	F_port	AutoNeg.
16	Green Storage 1	F_port	AutoNeg.

#### Switch 2, FCSM 2:

Port #	Symbolic port name	Port type	Port speed
0	Red Storage 2	F_port	AutoNeg.
1	Red Blade HBA 2	F_port	AutoNeg.
2	Blue Blade HBA 2	F_port	AutoNeg.

Port #	Symbolic port name	Port type	Port speed
3	Green Blade HBA 2	F_port	AutoNeg.
15	Blue Storage 2	F_port	AutoNeg.
16	Green Storage 2	F_port	AutoNeg.

Repeat steps [Step a](#) through [Step f](#) for the other switch module (172.23.199.23, switch Name: IBM\_brcd\_module2”).

4. Connect the FC switch modules and verify connectivity.

Connect FC cables from the first two or left-most external port on the switch modules to the two storage devices. To verify connectivity:

- a. A **switchshow** (CLI) on any of the above configured Brocade switch modules with this configuration will show ports 1, 2, and 3 as F\_Ports with the respective Server Blade HBA port WWNs. Ports 0, 11 on both the Brocade modules will again show up as F\_Ports with the respective Symmetrix port WWNs.

The red and green colors indicate what initiator is mapped to which target. The red server blade must log in to Symmetrix dir 3A and needs to be zoned accordingly while the green server blades are meant to log in to Symmetrix dir 3C and needs to be zoned accordingly.

- b. A **fabricshow** (CLI) on any of the modules must show two domains, each switch module representing an independent domain. A Name Server query must list ten ports: six ports coming from the three servers with dual port HBAs and four storage ports.

As shown in [Figure 10 on page 76](#), each server blade has two paths to access the respective Symmetrix director it is zoned with. The two paths go through the two switch modules respectively.

5. Configure zoning on the Brocade switch modules:

Zoning is configured on the Brocade modules using either the Brocade CLI or Web Tools. This example uses Brocade Web Tools to do the zoning.

- a. Point the web browser to the Brocade module (172.23.199.22), which opens Web Tools.

- b. Select the **Zoning** icon on the bottom left of the interface.
- c. Enter username and password as *USERID* and *PASSWORD*.
- d. On the **Zoning** interface screen, select the **Zone** tab and then select **Create**.
- e. On the **Create New Zone** dialog box, provide a descriptive zone name. This example will zone "Red Host HBA 1" and "Red Storage 1", so "**RedHBA1\_1470\_8aa**" will be entered.
- f. Verify that the **Zone Name** shows in the drop-down menu. Enter the respective zone members from the **Member Selection list** of N\_Ports: Host and storage WWPNs from the left side of the zoning dialog box into the "Zone Members" area on the right.
- g. Repeat [Step d](#) through [Step f](#) for all host and storage pairs in the environment.
- h. Create a zone set by clicking on the **Config** tab and then clicking **New**.
- i. In the **Create New Config** dialog box, provide a descriptive name for the configuration (zone set). In this case, the date of "Oct\_31\_06\_1140" will be used.
- j. Add the zones created in [Step d](#) through [Step g](#) into the configuration created in [Step i](#).
- k. Select **Actions** from the menu bar, and then select **Enable Config** to activate and save the desired zoning configuration; in this case "Oct\_31\_06\_1140".

When completed, the active or effective configuration displayed should be similar to what is shown below:

```
Zone set name = "Oct_31_06_1140"
```

```
Zone name = "RedHBA1_1470_8aa"
Zone Member = "10000000c938e554"
Zone Member = "50060482cc19bf87"
```

```
Zone name = "RedHBA2_1470_9aa"
Zone Member = "10000000c938e555"
Zone Member = "50060482cc19bf88"
```

```
Zone name = "GreenHBA1_1489_8aa"
Zone Member = "10000000c939a051"
Zone Member = "50060482cc19c447"
```

```
Zone name = "GreenHBA2_1489_9aa"  
Zone Member = "10000000c939a052"  
Zone Member = "50060482cc19c448"  
  
Zone name = "BlueHBA1_AllBlueStorage"  
Zone Member = "210100e08b8ac76d"  
Zone Member = "50060482cc19c407"  
Zone Member = "50060482cc19c408"  
  
Zone name = "BlueHBA2_AllBlueStorage"  
Zone Member = "210100e08baac76d"  
Zone Member = "50060482cc19c407"  
Zone Member = "50060482cc19c408"
```

### Complete the SAN setup

At this point, the SAN is ready to pass I/O from the host to storage. Other steps, such as configuring LUN Masking and modification of host configuration files, are required before the SAN setup is complete. Refer to the OS configuration guide for more details.



## Complex Fibre Channel SAN Topologies

This chapter provides the following information on complex Fibre Channel SAN topologies.

- ◆ Best practices..... 88
- ◆ Four switch full mesh ..... 89
- ◆ Compound core edge topologies ..... 113
- ◆ Heterogeneous switch interoperability..... 148

---

## Best practices

General best practices for simple and complex Fibre Channel SAN topologies are described in [“General best practices” on page 19](#). The information in this section is specific to complex Fibre Channel SAN topologies only.

---

### ISL subscription

While planning the SAN, keep track of how many host and storage pairs will be utilizing the ISLs between domains. As a general best practice, if two switches will be connected by ISLs, ensure that there is a minimum of two ISLs between them and that there are no more than six initiator and target pairs per ISL. For example, if 14 initiators access a total of 14 targets between two domains, a total of three ISLs would be necessary. Consider the applications that will use the ISLs before applying this best practice when setting up a configuration.

---

### Host and storage layout

For host and storage layout information for both simple and complex Fibre Channel SAN topologies, refer to [“Host and storage layout” on page 27](#).

---

### Switch and fabric management

For switch and fabric management information for both simple and complex Fibre Channel SAN topologies, refer to [“Switch and fabric management” on page 29](#).

---

### Security

It is important to secure your fabric. For general information on security, refer to [“Security” on page 31](#). For more information on security, refer to the *Building Secure SANs TechBook*, available on the [E-Lab Navigator, Documents > Topology Resource Center](#).



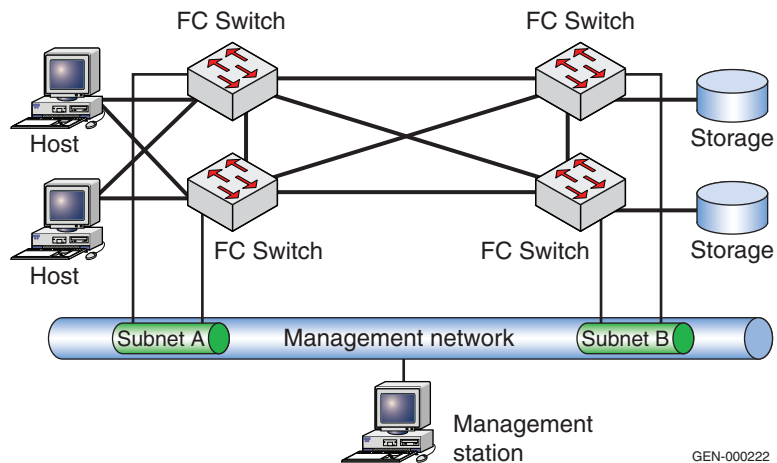
## Four switch full mesh

This section contains information on four switch full mesh topologies.

### Overview of fabric design considerations

#### General layout

In the four switch full mesh fabric shown in [Figure 11](#), each switch is connected to every other switch with a minimum of two ISLs. This prevents any ISL from becoming a single point of failure. Each switch is also connected to a management LAN through IP.



**Figure 11** Four switch full mesh fabric

Each switch type can be used in any position.

#### Best practices

For general best practices for all SANs, refer to [“General best practices” on page 19](#). For specific best practices to complex Fibre Channel SAN topologies, refer to [“Best practices” on page 88](#).

Specific information for four switch full mesh fabrics follows:

One of the use cases for a four switch full mesh fabric is distance extension. In these configurations, it is essential to monitor the ISLs for oversubscription conditions (utilization > 80%) which may lead to back pressure and any errors that are incrementing, especially bit errors or invalid transmission words, as these may lead to credit starvation.

For more information on credit starvation, refer to “BB-Credit Loss” section in the *Networked Storage Concepts and Protocols TechBook*, available on the [E-Lab Navigator, Documents > Topology Resource Center](#). See the individual case studies below for information on how to configure in each environment.

### Host and storage layout

Specific information for four switch full mesh fabrics follows:

In the four switch fabric examples used in this section, hosts and storage can be connected to any switch, but should be connected to the same switch when possible. A notable exception to this is in a distance extension environment when the switches are used to aggregate many different connections over an ISL and provide additional BB\_Credit. In this configuration, the whole point of having multiple switches is to use the ISLs.

For general information on host and storage layout or all SANs, refer to “[Host and storage layout](#)” on page 27.

### Switch and fabric management

For general information on switch and fabric management or all SANs, refer to “[Switch and fabric management](#)” on page 29.

### Security

For general information regarding security or all SANs, refer to “[Security](#)” on page 31. For more information on security, refer to the *Building Secure SANs TechBook*, available on the [E-Lab Navigator, Documents > Topology Resource Center](#).

---

## Connectrix B example

This section contains information for this Connectrix B example.

### General layout

As shown in [Figure 12 on page 91](#), a four switch full mesh fabric has been created by connecting each switch to every other switch in the fabric with two ISLs. Each switch is also connected to a Management LAN via Ethernet cables. There are three hosts (Red, Blue, and Green) as well as three sets of storage ports (also Red, Blue, or Green). The colors are intended to indicate which hosts access which storage ports.

In addition to the CTPs, each ED-48000B contains the following three blades:

- ◆ Slot 1 — FC4-16
- ◆ Slot 2 — FC4-32
- ◆ Slot 3 — FR4-18i

To enhance the continuity of this document, this configuration has been created with the intention of reusing it in other sections. As a result, not all of the capabilities of the FR4-18i will be utilized in this section.

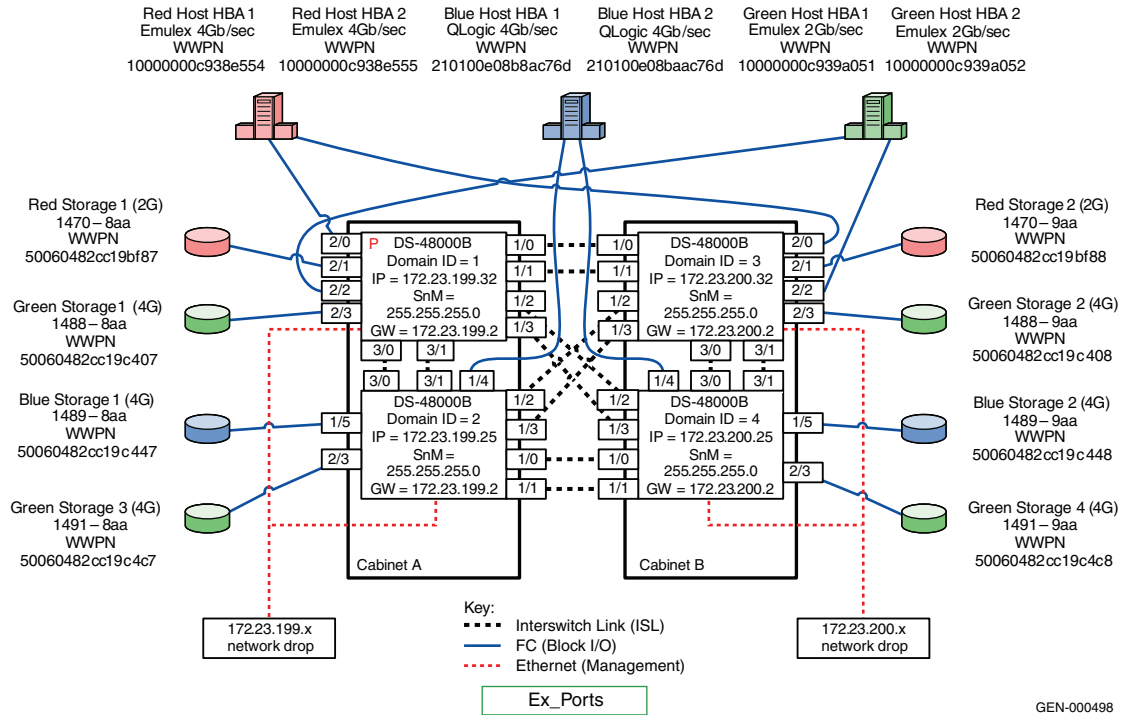


Figure 12 Four switch fabric with ED-48000Bs

All Connectrix® B series switches supporting v5.2.x Fabric Operating software and higher are supported using this topology. In this example, the Connectrix ED-48000B director is used.

**Best practices**

For general information on best practices for four switch fabrics, refer to “General best practices” on page 19.

For Connectrix B specific best practices, refer to “Connectrix B-Series” on page 24.

**ISL layout**

In this example, the ISLs are connected to ports on Slot 1 (FC4-16) and on Slot 3 (FR4-18i) and not to Slot 2 (FC4-32). The reasoning behind this layout is that the FC4-32 is oversubscribed (16:8), and because of this, utilizing these ports for ISLs may make the environment more

susceptible to congestion and backpressure. The ISLs are spread out over two different blades to enhance fault tolerance and serviceability. ISLs destined to the same Domain are kept within the same port octet to take advantage of trunking.

### Host and storage layout

Both hosts and storage can be placed anywhere on the SAN. However in this example, we will try to conserve the ports on the FC4-16 and only use them for ISLs or host and storage pairs which could potentially utilize the full 4 Gb/s bandwidth. Host and storage pairs that cannot fully utilize the full 4 Gb/s bandwidth are placed onto the FC4-32 blade. For example, while the Red host HBAs are 4 Gb/sec, the storage ports they access are only 2 Gb/s. Because of this, they are placed onto the FC4-32. Similar reasoning is applied to the Green host since the HBAs are 2 Gb/s even though the storage ports are 4 Gb/sec. The Blue host and storage ports are all 4 Gb/s so they are placed on the FC4-16.

---

**Note:** It is not essential to configure the environment this way, but it does ensure that all resources are being efficiently utilized.

---

For general information on host and storage layout for four switch fabrics, refer to [“Host and storage layout” on page 88](#).

### Switch and fabric management

In this example, CLI will be used to configure the environment, but Fabric Manager or Web Tools could also have been used.

For general information on switch and fabric management for four switch fabrics, refer to [“Switch and fabric management” on page 88](#).

### Security

The connectivity and device discovery for a large Connectrix B switch fabric may be secured by appointing the following binding techniques. All these are Connectrix B specific features.

- ◆ Fabric Binding is a security method for restricting switches within a multiple-switch fabric. The SCC policy prevents unauthorized switches from joining a fabric. Switches are authenticated using digital certificates and unique private keys provided to the Switch Link Authentication Protocol (SLAP).
- ◆ Switch Binding is a security method for restricting devices that connect to a particular switch. If the device is another switch, this is handled by the SCC policy. If the device is a host or storage device, the Device Connection Control (DCC) policy binds those devices to a particular switch. Policies range from completely restrictive to reasonably flexible, based upon customer needs.

- ◆ Port Binding is a security method for restricting host or storage devices that connect to particular switch ports. The DCC policy also binds device ports to switch ports. Policies range from completely restrictive to reasonably flexible, based on customer needs. For switches running Fabric OS v5.2.0 and later, the SCC ACL with strict fabric-wide consistency can also be used for Switch Binding in addition to the Secure Fabric OS mechanism.

The method to enable SCC and DCC policies has been provided in “[Enabling the Switch Connection Policy \(SCC\)](#)” on page 130.

For general information on security for four switch fabrics, refer to “[Security](#)” on page 88. Specific information for this example follows.

### Setting up this topology

#### Assumptions specific to this case study:

- ◆ The switches are installed in an EMC-supplied cabinet.
  - For installation instructions, see *Connectrix EC-1500 Cabinet Installation and Setup Manual*, which can be accessed from EMC Online Support at <https://support.emc.com>.
- ◆ The proper power receptacles have been provided by the customer.
  - For switch power requirements, refer to the *EMC Connectrix SAN Products Data Reference Manual*, available on the [E-Lab Navigator](#), **Documents > Topology Resource Center**.
- ◆ The switches have *not* been connected to the power source and are *not* powered on.
- ◆ Network drops, IP addresses, subnet mask, and gateway have been provided by the customer.
  - For switch or cabinet network requirements, refer to the *EMC Connectrix SAN Products Data Reference Manual*, available on the [E-Lab Navigator](#), **Documents > Topology Resource Center**.

---

**Note:** In this example, we are going to assume that the customer provided us with 8 Ethernet cables and that four of them are on the 172.23.199.x network and that the other four are connected to the 172.23.200.x network.

---

- ◆ The correct number of line cards have been installed into the ED-4800Bs.

- For help in determining how many ports are required, refer to “Determining customer requirements” in the *Networked Storage Concepts and Protocols TechBook*, available on the [E-Lab Navigator](#), **Documents > Topology Resource Center**.
- ◆ License keys have been obtained.
  - Go to the URL listed on the transaction code certificate that shipped with the product to get the license key.
- ◆ The customer has provided a temporary password that will be used as the default password when configuring the IP address.
- ◆ FOS v5.2.0a or greater is installed on all switches in the fabric.

### Configure the IP address

To configure the IP address:

---

**Note:** Connectrix B switches may ship with a default IP address not on the 172.23.199.x subnet. The ED-4800B director uses a maximum of three IPs per unit: one IP for the switch and one IP for each control processor.

---

1. Attach the provided serial cable between the serial port on Domain 1 and an RS-232 serial port on the management PC. The serial cable is wired with only pins 2, 3, and 5 wired straight through.
2. Power up the switch by connecting the power cords to the power receptacles provided by the customer.
3. Run a terminal emulation program, such as Hyperterm, on Windows hosts, or TERM in a UNIX environment.
4. Configure the terminal for 9600 Baud, 8 Data Bits, No Parity, 1 stop bit, and no flow control.
5. Press **Return** to get a prompt.

### **IMPORTANT**

**It is strongly recommended that when prompted to change the password, you change it to a password that was provided by the customer. This can also be done using the `passwd` command from the prompt at any time.**

---

6. Log in using the default values: Username: *admin* Password: *password*.
7. At the prompt, enter **ipaddrset -sw1** and press **Return**.

8. When prompted, supply IP address (172.23.199.22), subnet mask (255.255.255.0), and gateway address (172.23.199.2).

---

**Note:** The Fibre Channel addresses will not be used for this example.

---

9. At the prompt, enter **ipaddrset -cp 0** and press **Return**.
10. When prompted, supply hostname (**48K\_1\_23**), IP address (**172.23.199.23**), subnet mask (**255.255.255.0**), and gateway address (**172.23.199.2**).
11. At the prompt, enter **ipaddrset -cp 1** and press **Return**.
12. When prompted, supply hostname (**48K\_1\_24**), IP address (**172.23.199.24**), subnet mask (**255.255.255.0**), and gateway address (**172.23.199.2**).
13. Power down the switch and disconnect the serial cable.
14. Connect the switch to a 10/100BaseT Ethernet connection.
15. Power up the switch.

The switch can now be accessed with IP-based management.

16. Repeat these steps for Domain 2, 3, and 4 using the following information:

	Domain ID	Switch IP	CP0 IP	CP0 Name	CP1 IP	CP1 name
<b>Domain 1</b>	4	172.23.199.22	172.23.199.23	48K_1_23	172.23.199.24	48K_1_24
<b>Domain 2</b>	7	172.23.199.25	172.23.199.26	48K_2_26	172.23.199.27	48K_2_27
<b>Domain 3</b>	10	172.23.200.22	172.23.200.23	48K_3_23	172.23.200.24	48K_3_24
<b>Domain 4</b>	13	172.23.200.25	172.23.200.26	48K_4_26	172.23.200.27	48K_4_27

---

**Note:** For Domains 3 and 4, use the gateway address of **172.23.200.2**.

---

### Configure FC switches

To configure FC switches:

1. Set the switch name and fabric parameters for the switch with the Domain ID of 1.

---

**Note:** The following configurations need to be done with the switch *disabled*.

---

2. Configure the fabric parameters.
  - a. From the switch prompt, enter **switchdisable** to disable the switch.
  - b. From the switch prompt, enter **configure** to enter the configuration parameter menu.
  - c. Enter **Y** at the **Fabric Parameters** prompt.
  - d. Enter **1** for desired Domain ID at the Domain prompt and press **Enter**.
  - e. The R\_A\_TOV should be automatically set to 10000. If it is not, enter **10000** at the prompt and press **Enter**.
  - f. The E\_D\_TOV should be automatically set to 2000. If it is not, enter **2000** at the prompt and press **Enter**.
  - g. Accept the following defaults for the rest of the fields under the **Fabric Parameters** menu by pressing **Enter** after each prompt:
    - WAN\_TOV = 0
    - MAX\_HOPS = 7
    - Data field size = 2112
    - Sequence Level Switching = 0
    - Disable Device Probing = 0
    - Suppress Class F Traffic = 0
    - Switch PID Format = 1
    - Per-frame Route Priority = 0
    - Long Distance Fabric = 0
    - BB\_Credit = 16

---

**Note:** For this case study, there is no long distance between the DS-48000B switches. The ISLs connecting the two are less than 10 km.

---

- h. At the **Insistent Domain ID Mode** prompt, enter **y** to accept the **Insistent Domain ID** setting. When this mode is set, the switch attempts to acquire the domain number programmed in its **Switch Fabric Settings** from the fabric.



- i. Accept the default values from the remaining **Fabric Parameter Menu** items by pressing **Enter** after each prompt:

- Virtual Channel parameters (yes, y, no, n): **[no]**
- F\_Port login 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]**
- ssl attributes (yes, y, no, n): **[no]**
- http attributes (yes, y, no, n): **[no]**
- snmp attributes (yes, y, no, n): **[no]**
- rpcd attributes (yes, y, no, n): **[no]**
- cfgload attributes (yes, y, no, n): **[no]**
- web tools attributes (yes, y, no, n): **[no]**

---

**Note:** You may also press **CNTRL D** after making the last change in the menu to exit and save the changes. This will eliminate the need to accept the default values for the rest of the menu items.

---

- j. Repeat from [Step a](#) for switches with the Domain IDs 2, 3, and 4.

### Connect cables

To connect the cables:

1. Connect ISLs.
  - a. Attach Fiber cable between switches as shown in [Figure 12 on page 91](#).
  - b. After all cables are connected, use **switchshow** and **topologyshow** commands to ensure all links.
2. Connect host and storage ports.
  - a. Attach fiber cable between switches and N\_Ports.
3. Verify port login status.
  - a. After all cables are connected, use the **switchshow** CLI command to verify the all of the ports logged into the switch.

## Zone hosts and storage

To zone hosts and storage, Telnet into one of the switches in the fabric and using the following zoning commands:

1. Create zones using the **zonecreate** commands below:

```
zonecreate "RedHBA1_1470_8aa", "10:00:00:00:c9:38:e5:54;
50:06:04:82:cc:19:bf:87"
zonecreate "RedHBA2_1470_9aa", "10:00:00:00:c9:38:e5:55;
50:06:04:82:cc:19:bf:88"
zonecreate "BlueHBA1_1489_8aa", "21:01:00:e0:8b:8a:c7:6d;
50:06:04:82:cc:19:c4:47"
zonecreate "BlueHBA2_1489_9aa", "21:01:00:e0:8b:aa:c7:6d;
50:06:04:82:cc:19:c4:48"
zonecreate "GreenHBA1_AllGreenStorage", "10:00:00:00:c9:39:e5:51;
50:06:04:82:cc:19:c4:07; 50:06:04:82:cc:19:c4:08; 50:06:04:82:cc:19:c4:c7;
50:06:04:82:cc:19:c4:c8"
zonecreate "GreenHBA2_AllGreenStorage", "10:00:00:00:c9:39:e5:52;
50:06:04:82:cc:19:c4:07; 50:06:04:82:cc:19:c4:08; 50:06:04:82:cc:19:c4:c7;
50:06:04:82:cc:19:c4:c8"
```

2. Create the configuration by using the **cfgcreate** command.

```
cfgcreate "Oct_31_06_1140" , "RedHBA1_1470_8aa; RedHBA2_1470_9aa;
BlueHBA1_1489_8aa; BlueHBA2_1489_9aa; GreenHBA1_AllGreenStorage;
GreenHBA2_AllGreenStorage"
```

3. Enable the configuration by using the **cfgenable** command.

```
cfgenable "Oct_31_06_1140"
```

4. Enter **Y** at the confirmation prompt.

5. Enter **cfgshow** to display zoning info.

When completed, the zone information should be similar to what is shown below.

### Defined configuration:

```
cfg: Oct_31_06_1140
  RedHBA1_1470_8aa; RedHBA2_1470_9aa; BlueHBA1_1489_8aa; BlueHBA2_1489_9aa;
  GreenHBA1_AllGreenStorage; GreenHBA2_AllGreenStorage"
zone: RedHBA1_1470_8aa
10000000c938e554; 50060482cc19bf87
zone: RedHBA2_1470_9aa
10000000c938e555; 50060482cc19bf88
zone: BlueHBA1_1489_8aa
210100e08b8ac76d; 50060482cc19c447
zone: BlueHBA2_1489_9aa
210100e08baac76d; 50060482cc19c448
zone: GreenHBA1_AllGreenStorage
```

```

10000000c939a051; 50060482cc19c407;
50060482cc19c408; 50060482cc19c4c7;
50060482cc19c4c8
zone: GreenHBA2_AllGreenStorage
10000000c939a052; 50060482cc19c407;
50060482cc19c408; 50060482cc19c4c7;
50060482cc19c4c8

```

### Effective configuration:

```

CFG: Oct_31_06_1140
Zone: RedHBA1_1470_8aa
10000000c938e554
50060482cc19bf87
Zone: RedHBA2_1470_9aa
10000000c938e555
50060482cc19bf88
Zone: BlueHBA1_1489_8aa
210100e08b8ac76d
50060482cc19c447
Zone: BlueHBA2_1489_9aa
210100e08baac76d
50060482cc19c448
Zone: GreenHBA1_AllGreenStorage
10000000c939a051
50060482cc19c407
50060482cc19c408
50060482cc19c4c7
50060482cc19c4c8
Zone name = "GreenHBA2_AllGreenStorage
10000000c939a052
50060482cc19c407
50060482cc19c408
50060482cc19c4c7
50060482cc19c4c8

```

### Save configuration

In case the configuration is lost, or unintentional changes are made, keep a backup copy of the configuration file on a host computer.

To upload a configuration file:

1. Verify that the FTP service is running on the host computer. The host must have an FTP server application running.
2. Connect to the switch through the Telnet and log in as admin.
3. Enter the **configUpload** command.

The command becomes interactive and you are prompted for the required information. For example:

```
switch:admin> configupload
```

```

Protocol (scp or ftp) [ftp]: ftp
Server Name or IP address [host]: 192.1.2.3
User Name [user]: JohnDoe
File Name [config.txt]: /pub/configurations/config.txt
Password: xxxxxx
configUpload complete: All config parameters are uploaded.
switch:admin>

```

### Enable the Switch Connection Policy (SCC)

To enable the SCC:

1. At the switch prompt, enter `fdccfg -fabwideset "SCC:S;DCC;"`
2. Press **Enter**.

This command will set a strict SCC and tolerant DCC fabric-wide consistency policy.

---

**Note:** When a switch is joined to a fabric with a strict Switch Connection Control (SCC) or Device Connection Control (DCC) fabric-wide consistency policy, the joining switch must have a matching fabric-wide consistency policy. If the strict SCC or DCC fabric-wide consistency policies do not match, the switch cannot join the fabric and the neighboring E\_Ports will be disabled. If the strict SCC and DCC fabric-wide consistency policies match, the corresponding SCC and DCC access control list (ACL) policies are compared.

---

3. To verify that the policy has been set, the `fdccfg --showall` command can be run on any switch in the fabric. Any switch on the fabric should show output similar to:

```

switch:admin> fdccfg --showall
Local Switch Configuration for all Databases:-
DATABASE - Accept/Reject
-----
SCC - accept
DCC - accept
PWD - accept
Fabric Wide Consistency Policy:- "SCC:S;DCC"

```

### Complete the SAN setup

At this point the SAN is ready to pass I/O from host to storage. Other steps, such as configuring LUN Masking and modification of host configuration files, are required before the SAN setup is complete. Refer to the OS configuration guide for more details.

## Connectrix MDS example

**General layout** Figure 13 illustrates four Connectrix MDS 9506s full mesh configuration.




Figure 13 Four Connectrix MDS 9506s full mesh configuration

### Best practices

For general information on best practices for four switch fabrics, refer to “Best practices” on page 88. Specific information for this example follows.

By default thresholds are set to 80% utilization.

### Host and storage layout

For general information on host and storage layout for four switch fabrics, refer to “Host and storage layout” on page 88. Specific information for this example follows.

Line Rate Mode cards have no special restrictions. Over-subscribed cards should be used for hosts only.

**Switch and fabric management**

For general information on host and storage layout for four switch fabrics, refer to “[Switch and fabric management](#)” on page 88. Specific information for this example follows.

For this topology Fabric Manager can be used.

**Security**

For general information on security, for four switch fabrics, refer to “[Security](#)” on page 88. Specific information for this example follows.

Use switch and Port Binding for security.

**Setting up this topology****Assumptions specific to this case study:**

- ◆ The switches are installed in an EMC-supplied cabinet.
  - For installation instructions, see *Connectrix EC-1500 Cabinet Installation and Setup Manual*, which can be accessed from EMC Online Support at <https://support.emc.com>.
- ◆ The proper power receptacles have been provided by the customer.
  - For switch power requirements, refer to the *EMC Connectrix SAN Products Data Reference Manual*, available on the [E-Lab Navigator](#), **Documents > Topology Resource Center**.
  - For Cabinet power requirements, refer to *Connectrix EC-1500 Cabinet Installation and Setup Manual*, which can be accessed from EMC Online Support at <https://support.emc.com>.
- ◆ The switches have *not* been connected to the power source and are *not* powered on.
- ◆ Network drops, IP addresses, subnet mask, and gateway have been provided by the customer.

For switch or cabinet network requirements, refer to the *EMC Connectrix SAN Products Data Reference Manual*, available on the [E-Lab Navigator](#), **Documents > Topology Resource Center**.

---

**Note:** Connectrix MDS switches can be directly connected to the customer's LAN. The switches can be placed on either a public or private network. There are advantages to both configurations. For more information, refer to "Public versus private" in the *Networked Storage Concepts and Protocols TechBook*, available on the [E-Lab Navigator](#), **Documents > Topology Resource Center**.

This example assumes that the customer has provided us with two Ethernet cables and that one of them is on the 172.23.199.x network and that the other is connected to the 172.23.200.x network.

---

- ◆ The proper number of line cards have been installed into the Connectrix MDS 9513s.
  - For help in determining how many ports are required, refer to "Determining customer requirements" in the *Networked Storage Concepts and Protocols TechBook*, available on the [E-Lab Navigator](#), **Documents > Topology Resource Center**.
- ◆ License keys have been obtained.
  - Go to the URL listed on the transaction code certificate that shipped with the product to get the license key.
- ◆ A laptop, connected to a Connectrix MDS serial port, will be used to configure the IP addresses of the switches.
- ◆ The customer has provided a temporary password that will be used as the default password when configuring the IP address.
- ◆ Cisco CLI, Fabric Manager, and Device Manager will be used.

### Configure the IP address

To configure the IP address:

1. Power up the cabinet by connecting the power cords to the power receptacles provided by the customer.
2. Select one of the switches to configure and set the IP to 172.23.199.22.
3. Supply a network connection to the appropriate subnet.
4. Using an RS232 serial cable, connect to the serial port of the switch with a baud rate of 9600, 8 data bits, no parity, 1 stop bit and no flow control.

The **login** prompt should display.

5. Log in the first time with username *admin* and password *admin*.

You should be prompted to supply a new strong password for CLI user admin.

6. For this example, select **no** when asked if you want to run setup.

**Note:** This example will start with the switch that will have a Domain ID of 1 and an IP address of 172.23.199.22.

### CLI commands to configure the IP and gateway

- ◆ Switch# *config terminal*

Enter configuration commands, one per line.

Switch(config)# *interface mgmt 0*

Switch(config-if)# *IP address 172.23.199.22 255.255.255.0*

End with CNTL/Z.

- ◆ Switch# *config terminal*

Enter configuration commands, one per line.

Switch(config)# *ip default-gateway 172.23.199.2*

End with CNTL/Z.

To authorize access on a switch for Device and Fabric Manager, run this command on every switch while supplying a username (nnn) and password (ppp):

- ◆ Switch# *conf t*

Switch(config)# *snmp-server user nnn network-admin auth md5 ppp*

Switch(config)# *end*

Switch# *copy running-config startup-config*

Switch# *exit*

### Install Fabric Manager and Device Manager

To install Fabric Manager and Device Manager:

1. Open your web browser.
2. Enter the IP address of the switch into the address bar.
3. Follow the prompts and accept all defaults to install both Fabric Manager and Device Manager.



Fabric Manager and Device Manager can be started using the configured snmp-server username and password in “CLI commands to configure the IP and gateway” on page 104.

### Configure a VSAN

To configure a VSAN:

1. Open the Device Manager for the switch with an IP address of **172.23.199.22**.
2. Open the **VSAN** dialog box by selecting the **VSAN** menu item.
3. Click **Create**.
4. Enter the value of **100** into the **VSAN ID** field.
5. Set the **VSAN Name** to be “**Red\_VSAN\_100**”.
6. Use the default interop mode.
7. Click **Create**.

### Configure the other VSANs in this physical switch

To configure the other VSANs:

1. Repeat [Step 2](#) through [Step 7](#) above, in “Configure a VSAN,” for VSAN 200 and VSAN 300 nothing that:
  - a. For virtual switch 200, use **VSAN Name “Green\_VSAN\_200”**.
  - b. For virtual switch 300, use **VSAN Name “Blue\_VSAN\_300”**.
2. Assign and enable the ports to the proper VSAN using Device Manager for the switch.
  - a. Following [Step 1](#) and [Step 2](#), configure the ports of the switch with an IP address of 172.23.199.22 as shown in the tables below.

Slot #	Port #	Name	VSAN ID
1	1	TE ISL to SW 3	1
1	2	TE ISL to SW 4	1
1	3		
1	4		
1	5		

Slot #	Port #	Name	VSAN ID
1	6		
1	7	TE ISL to SW 2	1
1	8		

Slot #	Port #	Name	VSAN ID
2	1	TE ISL to SW 3	1
2	2	TE ISL to SW 4	1
2	3	Red Storage 1	100
2	4	Green Storage 1	200
2	5	Red Host HBA 1	100
2	6	Green Host HBA 1	200
2	7	TE ISL to SW 2	1
2	8		

- b. Following [Step 1](#) and [Step 2](#), configure the ports of the switch with an IP address of 172.23.199.23 as shown in the tables below.

Slot #	Port #	Name	VSAN ID
1	1	TE ISL to SW 3	1
1	2	TE ISL to SW 4	1
1	3	Blue Storage 1	300
1	4	Green Storage 3	200
1	5	Blue Host HBA 2	300
1	6		
1	7	TE ISL to SW 1	1
1	8		

Slot #	Port #	Name	VSAN ID
2	1	TE ISL to SW 3	1
2	2	TE ISL to SW 4	1
2	3	Blue Storage 1	300

Slot #	Port #	Name	VSAN ID
2	4	Green Storage 3	200
2	5	Green Storage 4	200
2	6		
2	7	TE ISL to SW 1	1
2	8		

- c. Following [Step 1](#) and [Step 2](#), configure the ports of the switch with an IP address of 172.23.200.22 as shown in the tables below.

Slot #	Port #	Name	VSAN ID
1	1	TE ISL to SW 1	1
1	2	TE ISL to SW 2	1
1	3		
1	4		
1	5		
1	6		
1	7	TE ISL to SW 4	1
1	8		

Slot #	Port #	Name	VSAN ID
2	2	TE ISL to SW 2	1
2	3	TE ISL to SW 4	1
2	4		
2	5	Red Host HBA 2	100
2	6	Green Host HBA 2	200
2	7	Red Storage 2	100
2	8	Green Storage 2	200

- d. Following [Step 1](#) and [Step 2](#), configure the ports of the switch with an IP address of 172.23.200.23 as shown in the table below.

Slot #	Port #	Name	VSAN ID
1	1	TE ISL to SW 1	1
1	2	TE ISL to SW 2	1
1	3		
1	4		
1	5	Blue Host HBA 2	300
1	6		
1	7	TE ISL to SW 3	1
1	8		

Slot #	Port #	Name	VSAN ID
2	1	TE ISL to SW 1	1
2	2	TE ISL to SW 2	1
2	3	TE ISL to SW 3	1
2	4		
2	5		
2	6		
2	7	Blue Storage 2	300
2	8	Green Storage 4	200

### Connect cables

To connect the cables:

1. Connect ISLs.
  - a. Attach fiber cable between switches as shown in [Figure 13 on page 101](#).
  - b. After all cables are connected, use Fabric Manager to verify that all ISL connections are up.
  - c. Re-arrange icons to accurately reflect the switch configuration.

**Note:** When looking at the topology view after persisting the fabric, you can immediately detect if something has changed in the environment. For example, if an ISL or device disappeared, yellow alert icons display. Because of this feature, it is recommended to *always* persist the fabric *after* changes have been made.

2. Connect host and storage ports.
  - a. Attach fibre cable between the switches and N\_Ports.

### Zone hosts and storage

To zone hosts and storage:

1. Open the **Zoning** menu in Fabric Manager for the desired VSAN.
2. Create a zone by clicking **New Zone** under **Zones**.
3. Provide a descriptive name for the zone. This example will zone "Red host HBA 1" and "Red Storage 1". Type "**RedHBA1\_1470\_8aa**" and press **Enter**.
4. In zone "**RedHBA1\_1470\_8aa**" select (**WWPN 10000000c938e554**) select **add to zone**.
5. Select "**Red Storage 1**" (**WWPN 50060482cc19bf87**) in the **potential zone members list**.
6. Create the VSAN zoneset, add the zones, then activate the zoneset.
7. Repeat [Step 2](#) through [Step 6](#) for all host and storage pairs in the environment.

```
Zone set name = "VSAN Red 100"
```

```
Zone name = "RedHBA1_1470_8aa"
Zone Member = "10000000c938e554"
Zone Member = "50060482cc19bf87"
```

```
Zone name = "RedHBA2_1470_9aa"
Zone Member = "10000000c938e555"
Zone Member = "50060482cc19bf88"
```

```
Zone set name = "VSAN Green 200"
```

```
Zone name = "GreenHBA1_1489_8aa"
Zone Member = "210100e08b8ac76d"
Zone Member = "50060482cc19c447"
```

```

Zone name = "GreenHBA2_1489_9aa"
Zone Member = "210100e08baac76d"
Zone Member = "50060482cc19c448"

Zone set name = "VSAN Blue 300"

Zone name = "BlueHBA1_AllBlueStorage"
Zone Member = "10000000c939a051"
Zone Member = "50060482cc19c407"
Zone Member = "50060482cc19c408"
Zone Member = "50060482cc19c4c7"
Zone Member = "50060482cc19c4c8"

Zone name = "BlueHBA2_AllBlueStorage"

Zone Member = "10000000c939a052"
Zone Member = "50060482cc19c407"
Zone Member = "50060482cc19c408"
Zone Member = "50060482cc19c4c7"
Zone Member = "50060482cc19c4c8"

```

### Complete the SAN setup

At this point the SAN is ready to pass I/O from host to storage. Other steps, such as configuring LUN masking and modification of host configuration files, are required before the SAN setup is complete. Refer to the *Cisco MDS 9000 Family Fabric Manager Configuration Guide* for more details.

### Configure IVR with Network Address Translation (NAT)

The configuration of VSANs on a fabric allows for security, scalability and availability. However, this isolation of traffic between VSANs prevents users from accessing resources, such as tape libraries, located in other VSANs. The solution to this limitation is Cisco's Inter-VSAN Routing feature, which allows initiators in one VSAN to access targets in other VSANs without merging the VSANs. Perhaps a host in the Red VSAN needs to access storage in the Blue VSAN. Configuring IVR zone and an IVR zone set containing the allowed initiators and targets allows communication between these resources.

This procedure also includes the configuration of IVR NAT, which allows duplicate Domain IDs to exist in the same fabric.

Without Network Address Translation (NAT), IVR requires unique Domain IDs for all switches in the fabric. You can enable IVR NAT to allow non-unique Domain IDs. This feature simplifies the deployment of IVR in an existing fabric where non-unique Domain IDs may be present.

---

**Note:** To use IVR NAT, it must be enabled in all IVR-enabled switches in the fabric IVR configuration distribution. By default, IVR NAT and IVR configuration distribution are disabled in all switches in the Cisco MDS 9000 Family.

---

1. In Fabric Manager:
  - a. Click the **Zone** tab in upper tool bar.
  - b. Click the **IVR** tab.
  - c. Select **Wizard**.

To migrate to IVR NAT mode click **Yes**; otherwise click **No**. You see the **IVR Zone Wizard**.

---

**Note:** If you are not using IVR NAT, Fabric Manager may display an error message if all the switches participating in IVR do not have unique Domain IDs. You must reconfigure those switches before configuring IVR.

---

2. Select the VSANs that will participate in IVR in the fabric.  
Select **VSAN 100, 200 and 300**.
3. Select the end devices that you want to communicate over IVR.

Select the following:

**VSAN 100: 10000000c938e554**

**VSAN 200: 50060482cc19c407**

**VSAN300: 50060482cc19c447**

**VSAN 300: 210100e08b8ac76d**

4. Enter the VSAN ID of the VSAN you want to use as the transit VSAN between the VSANs selected for the IVR zone.

Select **VSAN 1** as the transit VSAN.

---

**Note:** VSAN 1 connects both switches and all trunking traffic will pass over this link to communicate with VSANs in other switches.

---

5. Set the IVR zone and IVR zone set.

IVR NAME = **IVRZONE1**

IVR ZONENET NAME = **IVRZONESET1**

6. Verify all steps that Fabric Manager will take to configure IVR in the fabric.
7. Click **Finish** if you want to enable IVR NAT and IVR topology and to create the associated IVR zones and IVR zone set.  
  
or  
  
Click **Cancel** to exit the IVR Wizard without saving any changes.
8. The **Save Configuration** dialog box displays. You can save the configuration of the master switch to be copied to other IVR-enabled switches.  
  
Click either **Continue Activation** or **Cancel**.
9. Click **Finish**.

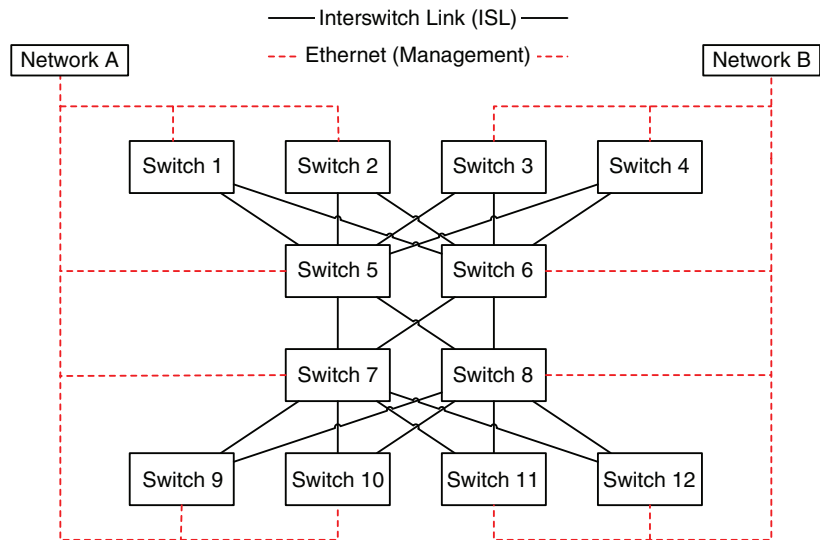


## Compound core edge topologies

This section provides examples of compound core edge switch topologies.

### Overview of fabric design considerations

**General layout** Figure 14 shows an example of a four switch compound core edge switches.



GEN-000252

**Figure 14** Four switch compound core edge switches

In the fabric shown in Figure 14, every core switch is connected to every other core switch while edge switches are only connected to two of the four cores. This is a classic core/edge design. Half of the switches are connected to Management Network A, and the other half are connected to Management Network B.

Only director class products should be used in the core of the fabric. This is due to the high number of ports that are consumed by ISLs and not due to other resource limitations (such as, CPU).

Both director class products and departmental class switches can be used in the edge position.

**Best practices**

Specific information on a compound core edge switch follows.

- ◆ Layout the host and storage connectivity such that if a switch fails, not all of a particular hosts storage becomes inaccessible.
- ◆ The use of two separate management networks is more common with balanced fabrics, but it can still be employed when only one fabric is used.
- ◆ ISL subscription best practice — While planning the SAN, keep track of the number of host and storage pairs that would be utilizing the ISLs between domains. As a general best practice, if two switches are connected by ISLs, ensure that there is a minimum of two ISLs between them, and that there are no more than six initiator and target pairs per ISL. For example, if 14 initiators access a total of 14 targets between two domains, a total of three ISLs are necessary. This best practice should not be applied blindly when setting up a configuration. Consider the applications that will use the ISLs.

For general information on best practices for all SANs, refer to [“General best practices” on page 19](#).

**Host and storage layout**

Specific information a compound core edge switch follows.

In the examples that follow, host and storage pairs are located in the following locations:

- ◆ Host on the edge and storage on the edge (Red)
- ◆ Host on the core and storage on the core (Blue)
- ◆ Host on the core and storage on the edge (Green)

The decision on where to place these host and storage pairs was not arbitrary. These were deliberately placed in areas that may not typically be thought of as "good" places to attach host and storage ports. This was done to stress the point that the only good place to attach host and storage ports is where it makes the most sense, given the customer's environment.

For general information on host and storage layout for all SANs, refer to [“Host and storage layout” on page 27](#).

**Switch and fabric management**

For general information on switch and fabric management for all SANs, refer to [“Switch and fabric management” on page 29](#).

**Security**

For general information on security for all SANs, refer to [“Security” on page 31](#).

## Connectrix B example

Figure 15 shows four ED-48000Bs in a full mesh configuration with edge switches attached.

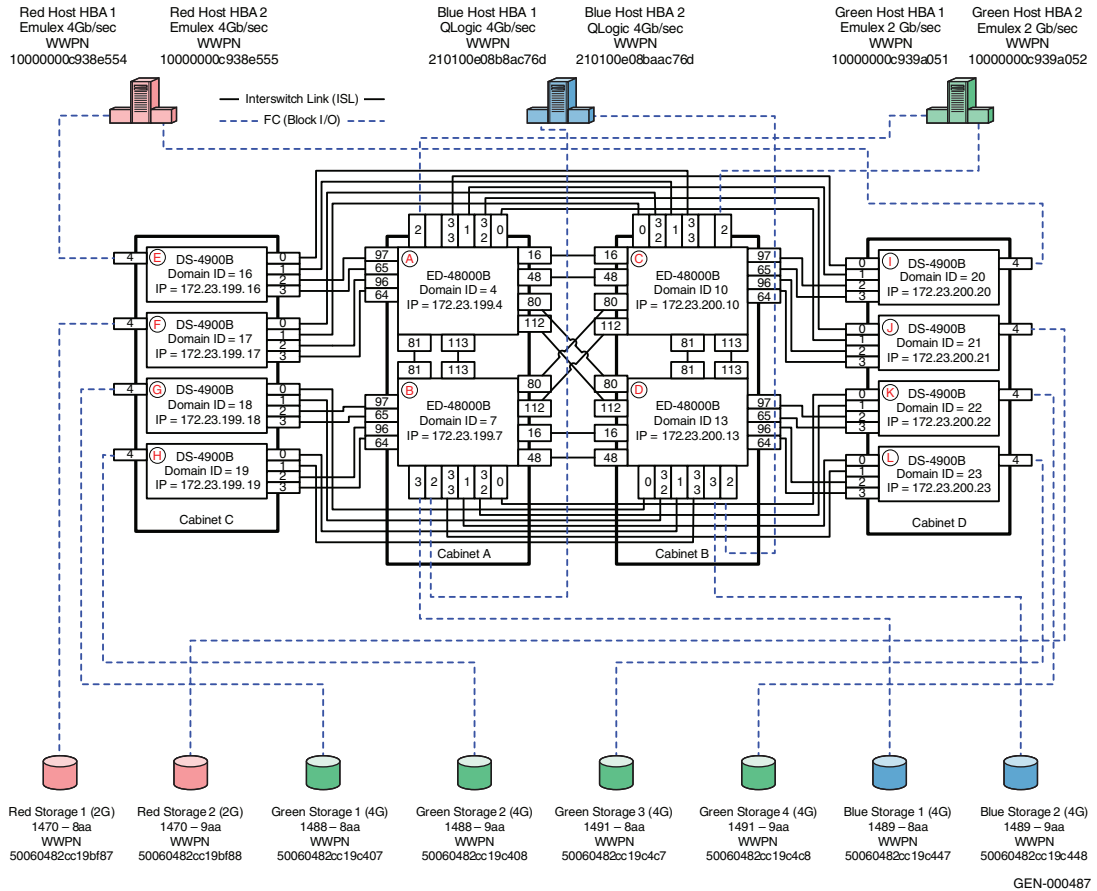


Figure 15 Four ED-48000Bs in full mesh configuration with edge switches attached

**Note:** Any director class Connectrix B product, such as the ED-24000B or ED-48000B, can be used in the core of this fabric. All EMC-supported Connectrix B director and switch class products can be used as edge switches.

**Best practices**

Specific information for this example is as follows:

While connecting the edge switches to the ED-48000B director switch it is recommended to be conversant with the architecture of the ED-48000B, its local switching capability, and the over-subscription management of the different switch port blades (i.e., the supported 16-port, 32-port, and 48-port blades) while designing scalable fabrics. An understanding of the product design will help utilize the ports on the switch blades for switch or host/target connectivity, enabling the simultaneous uncongested ports on all ports as long as simple best practices are followed.

Refer to the following link for further information on this subject:  
<http://www.brocade.com/products/competitive/directors.jsp>

For general information on best practices for a compound core edge switch, refer to “[Best practices](#)” on page 114. For Connectrix B specific best practices, refer to “[Connectrix B-Series](#)” on page 24.

**Host and storage layout**

To ensure fairness and increase the reliability of the SAN by eliminating any single point of overall failure and minimizing the impact that any single hardware failure would have, the N\_Port and E\_Port connections have been spread out on the ED-48000B over as many switch blade ports as possible.

The N\_Port and E\_Port connections on the DS-5300B are attached without any special considerations since there are no benefits to performance or reliability by spreading out the connections on the DS-4900B. In case the user wants to configure a trunk on the Connectrix B series switches, it is necessary that all ports in a given ISL trunk reside within an ASIC group on each end of the link. On 2 Gb/s switches, port groups are built on contiguous 4-port groups, called *quads*. On 4 Gb/s switches, like the Connectrix DS-48000B and the DS-4900B in this example, trunking port groups are built on contiguous 8-port groups, called *octets*. In these products, there are four octets: ports 0-7, 8-15, 16-23, and 24-31. The user must use the ports within a group as specified above to form an ISL trunk. It is also possible to configure multiple trunks within a port group.

For this case study, an attempt is made to connect ISLs to the same number ports on both switches to help assist with troubleshooting should the need arise.

For general information on host and storage layout for a compound core edge switch, refer to “[Host and storage layout](#)” on page 114.

**Switch and fabric management**

In this example, both CLI and Web Tools will be used to set up the compound core edge SAN topology, with the four Connectrix ED-48000Bs at the core and the eight DS-4900Bs at the edge.

For general information on switch and fabric management for a compound core edge switch, refer to [“Switch and fabric management” on page 114](#). Specific information for this example follows.

**Security**

The connectivity and device discovery for a large Connectrix B switch fabric may be secured by appointing the following binding techniques. All these are Connectrix B specific features.

- ◆ Fabric Binding is a security method for restricting switches within a multiple-switch fabric. The SCC policy prevents unauthorized switches from joining a fabric. Switches are authenticated using digital certificates and unique private keys provided to the Switch Link Authentication Protocol (SLAP).
- ◆ Switch Binding is a security method for restricting devices that connect to a particular switch. If the device is another switch, this is handled by the SCC policy. If the device is a host or storage device, the Device Connection Control (DCC) policy binds those devices to a particular switch. Policies range from completely restrictive to reasonably flexible, based upon customer needs.
- ◆ Port Binding is a security method for restricting host or storage devices that connect to particular switch ports. The DCC policy also binds device ports to switch ports. Policies range from completely restrictive to reasonably flexible, based on customer needs. For switches running Fabric OS v5.2.0 and later, the SCC ACL with strict fabric-wide consistency can also be used for Switch Binding in addition to the Secure Fabric OS mechanism.

The method to enable SCC and DCC policies has been provided at the end of the fabric configuration steps for this case study (refer to [“Enabling the Switch Connection Policy \(SCC\)” on page 130](#)).

For general information on security for a compound core edge switch, refer to [“Security” on page 114](#). Specific information for this example follows.

**Setting up this topology**

Assumption specific to this case study:

- ◆ The ED-48000B director class and the DS-4900B edge switches are installed in an EMC-supplied cabinet.

For installation instructions, refer to the *Connectrix EC-1500 Cabinet Installation and Setup Manual*, accessible from EMC Online Support at <https://support.emc.com>.

- ◆ The proper power receptacles have been provided by the customer.

For switch power requirements, refer to “Connectrix B series directors and switches” in the *EMC Connectrix SAN Products Data Reference Manual*, available on the [E-Lab Navigator](#), **Documents > Topology Resource Center**.

- ◆ The switches have not been connected to the power source and are not powered on.
- ◆ Network drops, IP addresses, subnet mask, and gateway have been provided by the customer.

For switch or cabinet network requirements, refer to “Connectrix B series directors and switches” in the *EMC Connectrix SAN Products Data Reference Manual*, available on the [E-Lab Navigator](#), **Documents > Topology Resource Center**.

- ◆ The proper number of line cards have been installed into the ED-4800Bs.

For help in determining how many ports are required, refer to “Determining customer requirements” in the *Networked Storage Concepts and Protocols TechBook*, available on the [E-Lab Navigator](#), **Documents > Topology Resource Center**.

- ◆ License keys have been obtained.

Go to the URL listed on the transaction code certificate that shipped with the product to get the license key.

- ◆ The customer has provided a temporary password that will be used as the default password when configuring the IP address.
- ◆ An EMC-supported Brocade FOS is installed on all switches in the fabric.

Refer to the *EMC Support Matrix* for the most current support information.

- ◆ A 32-port blade is installed in the first slot, a 16-port blade is installed in the second slot, a 48-port blade is installed in the fourth slot and the 32-port blade is installed in the fifth slot of all the ED-4800B director switches.

## Configure the IP address

To configure the IP address:

---

**Note:** Connectrix B switches may ship with a default IP address that is not on the desired subnet. The ED-48000B director uses a maximum of three IPs per unit: one IP for the switch and one IP for each control processor. The DS-4900B uses only one IP.

---

1. Attach the provided serial cable between the serial port on one of the ED-48000B switches (switch A) and an RS-232 serial port on the management PC. The serial cable is wired with only pins 2, 3, and 5 wired straight through.
2. Power up the switch by connecting the power cords to the power receptacles provided by the customer.
3. Run a terminal emulation program, such as Hyperterm on Windows hosts or TERM in a UNIX environment.
4. Configure the terminal for 9600 Baud, 8 Data Bits, No Parity, 1 stop bit, and no flow control.
5. Press **Return** to get a prompt.

---

**Note:** It is strongly recommended that when prompted to change the password, you change it to a password that was provided by the customer. This can also be done using the **passwd** command from the prompt at any time.

---

6. Log in using the default values: Username: *admin*; Password: *password*.
7. At the prompt, enter **ipaddrset -sw1** and press **Return**.
8. When prompted, supply IP address (172.23.199.4), subnet mask (255.255.255.0), and gateway address (172.23.199.2).

---

**Note:** The Fibre Channel addresses will not be used for this example.

---

9. At the prompt, enter **ipaddrset -cp 0** and press **Return**.
10. When prompted, supply hostname (**48K\_1\_5**), IP address (**172.23.199.5**), subnet mask (**255.255.255.0**), and gateway address (**172.23.199.2**).
11. At the prompt, enter **ipaddrset -cp 1** and press **Return**.

12. When prompted, supply hostname (48K\_1\_6), IP address (172.23.199.6), subnet mask (255.255.255.0), and gateway address (172.23.199.2).
13. Power down the switch and disconnect the serial cable.
14. Connect the switch to a 10/100BaseT Ethernet connection.
15. Power up the switch. The switch can now be accessed with IP-based management.
16. Repeat these steps for the other three switches B, C, and D using the following information:

	Domain ID	Switch IP	CP0 IP	CP0 Name	CP1 IP	CP1 Name
Switch A	4	172.23.199.4	172.23.199.5	48K_1_5	172.23.199.6	48K_1_6
Switch B	7	172.23.199.7	172.23.199.8	48K_1_8	172.23.199.9	48K_1_9
Switch C	10	172.23.200.10	172.23.200.11	48K_1_11	172.23.200.12	48K_1_12
Switch D	13	172.23.200.13	172.23.200.14	48K_1_14	172.23.200.15	48K_1_15

17. Repeat [Step 1](#) through [Step 6](#) for all of the DS-4900Bs: Switch E, F, G, H, I, J, K and L one-by-one, and then execute the following steps on switch E:
18. At the prompt, enter **ipaddrset** and press **Return**.
19. When prompted, supply IP address (172.23.199.16), subnet mask (255.255.255.0), and gateway address (172.23.199.2).
20. Power down the switch and disconnect the serial cable.
21. Connect the switch to a 10/100BaseT Ethernet connection.
22. Power up the switch. The switch can now be accessed with IP-based management.
23. Repeat steps [Step 17](#) through [Step 21](#) for the other seven switches F, G, H, I, J, K, and L using the following information:

	Domain ID	Switch IP	Switch Name
Switch E	16	172.23.199.16	4900_16
Switch F	17	172.23.199.17	4900_17
Switch G	18	172.23.199.18	4900_18



<b>Switch H</b>	19	172.23.199.19	4900_19
<b>Switch I</b>	20	172.23.200.20	4900_20
<b>Switch J</b>	21	172.23.200.21	4900_21
<b>Switch K</b>	22	172.23.200.22	4900_22
<b>Switch L</b>	23	172.23.200.23	4900_23

### Configure FC switches

To configure FC switches:

1. Set the switch name and fabric parameters for switch A.

**Note:** The following configurations need to be done with the switch *disabled*.

2. Configure the fabric parameters.
  - a. From the switch prompt, enter **switchdisable** to disable the switch.
  - b. From the switch prompt, enter **configure** to enter the configuration parameter menu.
  - c. Enter **Y** at the **Fabric Parameters** prompt.
  - d. Enter **4** for desired domain ID (as per the table above) at the Domain prompt and press **Enter**.
  - e. The R\_A\_TOV should be automatically set to 10000. If it is not, enter **10000** at the prompt and press **Enter**.
  - f. The E\_D\_TOV should be automatically set to 2000. If it is not, enter **2000** at the prompt and press **Enter**.
  - g. Accept the following defaults for the rest of the fields under the **Fabric Parameters** menu by pressing **Enter** after each prompt:
    - WAN\_TOV = 0
    - MAX\_HOPS = 7
    - Data field size = 2112
    - Sequence Level Switching = 0
    - Disable Device Probing = 0
    - Suppress Class F Traffic = 0
    - Switch PID Format = 1
    - Per-frame Route Priority = 0
    - Long Distance Fabric = 0

- BB\_Credit = 16

---

**Note:** For this case study, there is no long distance between any of the switches. The ISLs connecting the two are less than 10 km.

---

- h. At the **Insistent Domain ID Mode** prompt, enter **y** to accept the **Insistent domain ID** setting. When this mode is set, the switch attempts to acquire the domain number programmed in its **Switch Fabric Settings** from the fabric.
- i. Accept the default values from the remaining **Fabric Parameter Menu** items by pressing **Enter** after each prompt:
  - Virtual Channel parameters (yes, y, no, n): **[no]**
  - F\_Port login 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]**
  - ssl attributes (yes, y, no, n): **[no]**
  - http attributes (yes, y, no, n): **[no]**
  - snmp attributes (yes, y, no, n): **[no]**
  - rpcd attributes (yes, y, no, n): **[no]**
  - cfgload attributes (yes, y, no, n): **[no]**
  - web tools attributes (yes, y, no, n): **[no]**

---

**Note:** You may also press **CNTRL+D** after making the last change in the menu to exit and save the changes. This will eliminate the need to accept the default values for the rest of the menu items.

---

- j. Repeat from [Step a](#) for the ED-48000B switches B, C, and D, and for the DS-4900B switches E, F, G, H, I, J, K, and L using the values shown in the tables above, especially while entering the Domain IDs in [Step d](#).
3. In this case study, we are setting the Switch A as the principal switch. We have a choice of setting any of the core ED-48000B switches in this configuration as the principal switch. In order to do this:
    - a. Telnet into switch A.
    - b. At switch prompt, enter **fabricprincipal 1**. This will set the switch A as the principal switch. Verify the “switch mode” by running the **switchshow** command.

4. Ports on the switch may be configured if desired. By default the port type and port speed is set to auto.
  - a. Issue the **portcfgport** *<port number>* **1** to configure the port type to E\_Port for an ISL connection and to lock it as an E\_Port.
  - b. Issue the **portcfggport** *<port number>* **1** to configure as a generic G\_Port.
  - c. Issue the **portcfgspeed** command to configure the port speed.

### Connect cables

To connect the cables:

1. Connect ISLs.

Attach Fiber cable between switches as shown in [Figure 15 on page 115](#). Also refer to the port tables shown in [Step 2](#) for each individual switch. The index denotes the index number that is seen on running a **switchshow** on the Brocade switch.

2. Connect host and storage ports.

Attach fiber cable between switches and N\_Ports. (Refer to the port tables listed in this step for each individual switch. The index denotes the index number that is seen on running a **switchshow** on the Brocade switch).

The port connections need to be made as follows:

- For Switch A:

Index	Slot/Port#	Name
0	1/0	ISL to domain 21
1	1/1	ISL to domain 20
2	1/2	Green Host HBA 1
16	1/16	ISL to domain 10
32	2/0	ISL to domain 21
33	2/1	ISL to domain 20
48	4/0	ISL to domain 10
64	4/17	ISL to domain 17
65	4/18	ISL to domain 16
80	4/33	ISL to domain 13

Index	Slot/Port#	Name
81	4/34	ISL to domain 7
96	5/0	ISL to domain 17
97	5/1	ISL to domain 16
112	5/17	ISL to domain 13
113	5/18	ISL to domain 7

- For Switch B:

Index	Slot/Port#	Name
0	1/0	ISL to domain 22
1	1/1	ISL to domain 23
2	1/2	Blue Host HBA 1
3	1/3	Blue storage 1
16	1/16	ISL to domain 13
32	2/0	ISL to domain 22
33	2/1	ISL to domain 23
48	4/0	ISL to domain 13
64	4/17	ISL to domain 19
65	4/18	ISL to domain 18
80	4/33	ISL to domain 10
81	4/34	ISL to domain 4
96	5/0	ISL to domain 19
97	5/1	ISL to domain 18
112	5/17	ISL to domain 10
113	5/18	ISL to domain 4

- For Switch C:

Index	Slot/Port#	Name
0	1/0	ISL to domain 17
1	1/1	ISL to domain 16
2	1/2	Green Host HBA 2
16	1/16	ISL to domain 4

Index	Slot/Port#	Name
32	2/0	ISL to domain 17
33	2/1	ISL to domain 16
48	4/0	ISL to domain 4
64	4/17	ISL to domain 21
65	4/18	ISL to domain 20
80	4/33	ISL to domain 7
81	4/34	ISL to domain 13
96	5/0	ISL to domain 21
97	5/1	ISL to domain 20
112	5/17	ISL to domain 13
113	5/18	ISL to domain 7

- For Switch D:

Index	Slot/Port#	Name
0	1/0	ISL to domain 18
1	1/1	ISL to domain 19
2	1/2	Blue Host HBA 2
3	1/3	Blue storage 2
16	1/16	ISL to domain 7
32	2/0	ISL to domain 18
33	2/1	ISL to domain 19
48	4/0	ISL to domain 7
64	4/17	ISL to domain 23
65	4/18	ISL to domain 22
80	4/33	ISL to domain 4
81	4/34	ISL to domain 10
96	5/0	ISL to domain 23
97	5/1	ISL to domain 22
112	5/17	ISL to domain 4
113	5/18	ISL to domain 10

- For Switch E:

Port#	Name
0	ISL to domain 10
1	ISL to domain 10
2	ISL to domain 4
3	ISL to domain 4
4	Red Host HBA 1

- For Switch F:

Port#	Name
0	ISL to domain 10
1	ISL to domain 10
2	ISL to domain 4
3	ISL to domain 4
4	Red Storage 1

- For Switch G:

Port#	Name
0	ISL to domain 13
1	ISL to domain 13
2	ISL to domain 7
3	ISL to domain 7
4	Green Storage 1

- For Switch H:

Port#	Name
0	ISL to domain 13
1	ISL to domain 13
2	ISL to domain 7
3	ISL to domain 7
4	Green Storage 2

- For Switch I:

Port#	Name
0	ISL to domain 4
1	ISL to domain 4
2	ISL to domain 10
3	ISL to domain 10
4	Red Host HBA 2

- For Switch J:

Port#	Name
0	ISL to domain 4
1	ISL to domain 4
2	ISL to domain 10
3	ISL to domain 10
4	Red Storage 2

- For Switch K:

Port#	Name
0	ISL to domain 7
1	ISL to domain 7
2	ISL to domain 13
3	ISL to domain 13
4	Green Storage 4

- For Switch L:

Port#	Name
0	ISL to domain 7
1	ISL to domain 7
2	ISL to domain 13
3	ISL to domain 13
4	Green Storage 3

3. Verify port login status.
  - a. After all cables are connected, use the **switchshow** CLI command to verify all of the ports logged into the switch. The ports with ISLs must log in as E\_Ports while the ports connected to the HBA or storage must log in as F\_Ports.

### Zone hosts and storage

To zone hosts and storage:

1. Using a web browser, enter the IP address of Switch A (172.23.199.4) in the URL/address field.
2. Enter username and password at the prompt.
3. Click the **Zone Admin** icon in the lower left corner at bottom of web page.
4. The **Zone Admin** pop-up window should appear. You may be asked for login credentials again.
5. Under the **Alias** tab, click **New Alias**. A **Create New Alias** window appears. Enter **Red\_HBA1** as the alias name.



6. Expand the WWNs folder (if needed) by clicking on the + sign. Click on the WWN of Red\_HBA1 (**10:00:00:00:c9:38:e5:54**) and click **Add Member** to move the WWN to **Alias Members** column.

This will create an alias for Blue\_HBA1 using one of the WWNs of the host. All zoning created with this alias will now use its WWN for zoning.

7. Click **New Alias** again and in the **Create New Alias** window, enter **Red\_Storage1** (50:06:04:82:cc:19:bf:87) as the alias name.
8. Expand the WWNs folder (if needed) by clicking on the + sign. Select the WWN of Blue\_Storage\_A (**50:06:04:82:cc:19:bf:87**) and then select **Add Member** to move the WWN to the **Alias Members** column.

This will create an alias for Red\_Storage1 using the WWN of the storage device. All zoning created with this alias will now use its WWN for zoning.

9. Under the **Zone** tab, click **New Zone**. A **Create New Zone** window appears. Enter "**RedHBA1\_1470\_8aa**" for Zone Name. This case zones Red Host and Red Storage 1.
10. Expand the Alias folder (if needed) by clicking on the + sign. Select the "**Red\_HBA1**" alias and then click **Add Member** to move alias to the **Zone Members** column. Select the "**Red\_Storage1**" alias and then select **Add Member** to move alias to the **Zone Members** column.

This will create a zone for the Red\_HBA1 to allow access to the Red\_Storage1.

11. Under the **Zone Config** tab, click **New Zone Config**. A **Create New Config** window appears. This case uses the date as the name. Enter "**Oct\_31\_06\_1140**" as the config name.
12. Expand the Zones folder (if needed) by clicking on the + sign. Select the "**RedHBA1\_1470\_8aa**" zone and click **Add Member** to move zone to the **Zone Config Members** column.  
  
This will create a zone with the "**RedHBA1\_1470\_8aa**" zone as a zoneset member. This zoneset will need to be enabled before becoming effective.
13. Repeat [Step 5](#) through [Step 12](#) to create aliases and a zoneset for the other Red HBA and Red Storage port, Green HBAs, Green Storages, Blue HBAs and Blue Storage as stated below.

14. Click the **Zoning Actions** pull-down menu on top of window. Click **Enable Config** and select "**Oct\_31\_06\_1140**" zone and then click **OK**. This will push the new zone out to the fabric and make it effective.
15. When completed, the zone set, when running the **cfgactvshow** command from CLI, should be similar to what is shown below:

**Effective configuration:**

```
Cfg: Oct_31_06_1140
Zone: RedHBA1_1470_8aa
    10:00:00:00:c9:38:e5:54
    50:06:04:82:cc:19:bf:87

Zone: RedHBA2_1470_9aa
    10:00:00:00:c9:38:e5:55
    50:06:04:82:cc:19:bf:88

Zone: BlueHBA1_1489_8aa
    21:01:00:e0:8b:8a:c7:6d
    50:06:04:82:cc:19:c4:47

Zone: BlueHBA2_1489_9aa
    21:01:00:e0:8b:aa:c7:6d
    50:06:04:82:cc:19:c4:48

Zone: GreenHBA1_AllGreenStorage
    10:00:00:00:c9:39:a0:51
    50:06:04:82:cc:19:c4:07
    50:06:04:82:cc:19:c4:08
    50:06:04:82:cc:19:c4:c7
    50:06:04:82:cc:19:c4:c8

Zone: GreenHBA2_GreenStorage
    10:00:00:00:c9:39:a0:52
    50:06:04:82:cc:19:c4:07
    50:06:04:82:cc:19:c4:08
    50:06:04:82:cc:19:c4:c7
    50:06:04:82:cc:19:c4:c8
```

**Enabling the Switch Connection Policy (SCC)**

To enable the Switch Connection Policy:

1. At the switch prompt, enter **fdccfg -fabwideset "SCC:S;DCC"**.
2. Press **Enter**.

This command will set a strict SCC and tolerant DCC fabric-wide consistency policy.

---

**Note:** When a switch is joined to a fabric with a strict Switch Connection Control (SCC) or Device Connection Control (DCC) fabric-wide consistency policy, the joining switch must have a matching fabric-wide consistency policy. If the strict SCC or DCC fabric-wide consistency policies do not match, the switch cannot join the fabric and the neighboring E\_Ports will be disabled. If the strict SCC and DCC fabric-wide consistency policies match, the corresponding SCC and DCC access control list (ACL) policies are compared.

---

- To verify that the policy has been set, the **fddcfg --showall** command can be run on any switch in the fabric. Any switch on the fabric should show output similar to:

```
switch:admin> fddcfg --showall
Local Switch Configuration for all Databases:-
DATABASE - Accept/Reject
-----
SCC - accept
DCC - accept
PWD - accept
Fabric Wide Consistency Policy:- "SCC:S;DCC"
```

### Completing the SAN setup

At this point the SAN is ready to pass I/O from host to storage. Other steps, such as configuring LUN masking and modification of host configuration files, are required before the SAN setup is complete. Refer to the OS configuration guide for more details.

---

### Connectrix MDS example

This section contains information on this Connectrix MDS example.

General layout

Figure 16 shows four MDS 9513s in a full mesh configuration with edge switches attached.

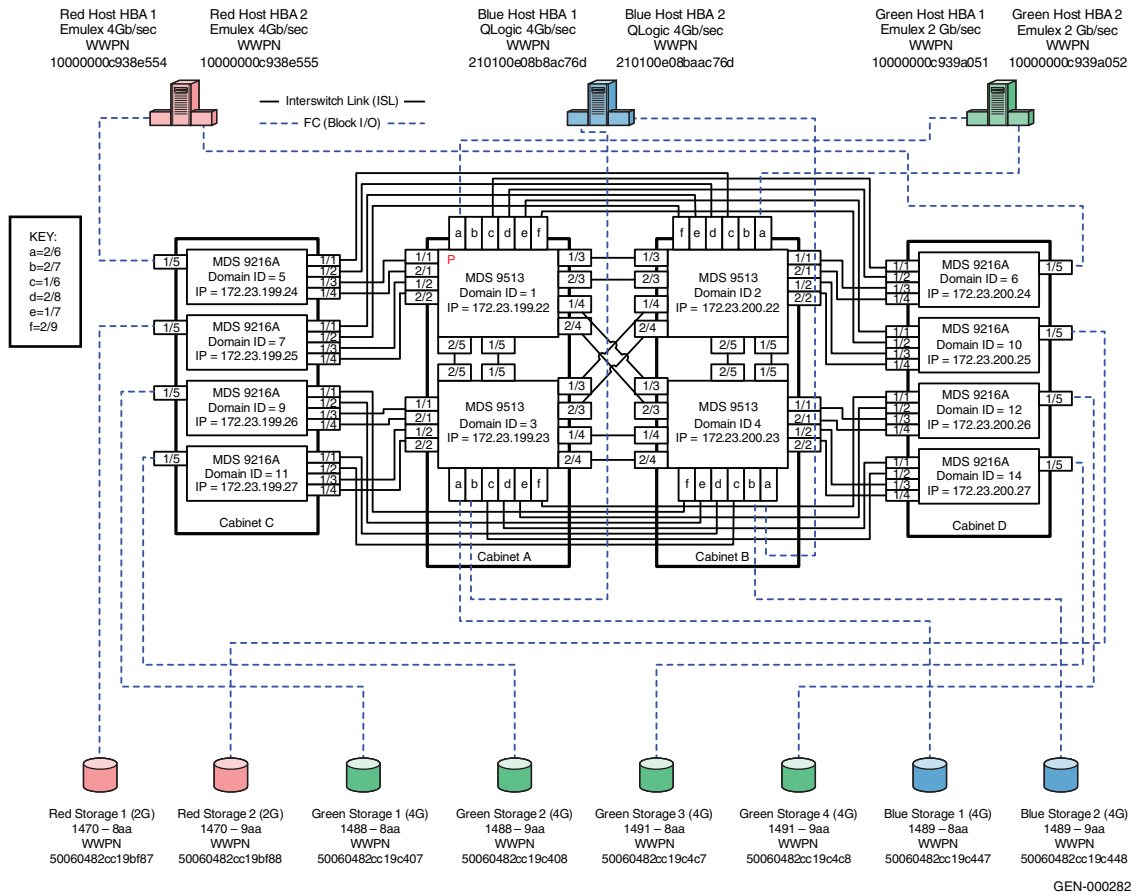


Figure 16 Four MDS 9513s in full mesh configuration with edge switches attached

Best practices

For general information on best practices for a compound core edge switch, refer to “Best practices” on page 114. Specific information for this example follows.

By default, utilization is set to 80%.

Host and storage layout

For general information on host and storage layout a compound core edge switch, refer to “Host and storage layout” on page 114. Specific information for this example follows.

Line Rate Cards have no special considerations. Over-subscribed cards should be used for hosts only.

### Switch and fabric management

For general information on switch and fabric management a compound core edge switch, refer to [“Switch and fabric management” on page 114](#). Specific information for this example follows.

Cisco Fabric Manager can be used for complex fabrics.

### Security

For general information on security a compound core edge switch, refer to [“Security” on page 114](#). Specific information for this example follows.

Enable switch and Port Binding for security.

### Setting up this topology

#### Assumptions specific to this case study:

- ◆ The switches are installed in an EMC-supplied cabinet.
  - For installation instructions, see *Connectrix EC-1500 Cabinet Installation and Setup Manual* which can be accessed from EMC Online Support at <https://support.emc.com>.
- ◆ The proper power receptacles have been provided by the customer.
  - For switch power requirements, refer to refer to *EMC Connectrix SAN Products Data Reference Manual*, available on the [E-Lab Navigator](#), **Documents > Topology Resource Center**.
  - For Cabinet power requirements, refer to *Connectrix EC-1500 Cabinet Installation and Setup Manual*, which can be accessed from EMC Online Support at <https://support.emc.com>.
- ◆ The switches have *not* been connected to the power source and are *not* powered on.
- ◆ Network drops, IP addresses, subnet mask, and gateway have been provided by the customer.

For switch or cabinet network requirements, refer to refer to *EMC Connectrix SAN Products Data Reference Manual*, available on the [E-Lab Navigator](#), **Documents > Topology Resource Center**.

---

**Note:** Connectrix MDS switches can be directly connected to the customer’s LAN. The switches can be placed on either a public or private network. There are advantages in both configurations. For more information, refer to refer to "Public versus private" in the *Networked*

*Storage Concepts and Protocols TechBook*, available on the [E-Lab Navigator, Documents > Topology Resource Center](#).

This example assumes that the customer has provided us with two Ethernet cables and that one of them is on the 172.23.199.x network and that the other is connected to the 172.23.200.x network.

- ◆ The proper number of line cards have been installed into the Connectrix MDS 9513s.

For help in determining how many ports are required, refer to "Determining customer requirements" in the *Networked Storage Concepts and Protocols TechBook*, available on the [E-Lab Navigator, Documents > Topology Resource Center](#).

- ◆ License keys have been obtained.
  - Go to the URL listed on the transaction code certificate that shipped with the product to get the license key.
- ◆ A laptop, connected to a Connectrix MDS serial port, is used to configure the IP addresses of the switches.
- ◆ The customer has provided a temporary password that is used as the default password when configuring the IP address.
- ◆ Cisco CLI, Fabric Manager, and Device Manager are used.

### Configure the IP address

To configure the IP address:

1. Power up the cabinet by connecting the power cords to the power receptacles provided by the customer.
2. Select one of the switches to configure and set the IP to 172.23.199.22.
3. Supply a network connection to the appropriate subnet.
4. Using an RS232 serial cable, connect to the serial port of the switch with a baud rate of 9600, 8 data bits, no parity, 1 stop bit and no flow control.

The **login** prompt should display.

5. Log in the first time with a username *admin* and password *admin*.  
You should be prompted to supply a new strong password for CLI user admin.
6. For this example, select **no** when prompted to run setup.

**Note:** This example starts with the switch that will have a Domain ID of 1 and an IP address of 172.23.199.22.

### CLI commands to configure the IP and gateway

- ◆ Switch# *config terminal*  
Enter configuration commands, one per line.  
Switch(config)# *interface mgmt 0*  
Switch(config-if)# *IP address 172.23.199.22 255.255.255.0*  
End with CNTL/Z.
- ◆ Switch# *config terminal*  
Enter configuration commands, one per line.  
Switch(config)# *ip default-gateway 172.23.199.2*  
End with CNTL/Z.

To authorize access on a switch for Device and Fabric Manager, run this command on every switch while supplying a username (nnn) and password (ppp):

- ◆ Switch#*conf t*  
Switch(config)# *snmp-server user nnn network-admin auth md5 ppp*  
Switch(config)#*end*  
Switch# *copy running-config startup-config*  
Switch# *exit*

### Configure the IP addresses of the switches in the cabinet

To configure the IP address of the switches in the other cabinet:

Follow [Step 4](#) through [Step 6](#) in “Configure the IP address” on [page 134](#).

Use the IP addresses found in [Figure 13](#) on [page 101](#) with subnet mask (255.255.255.0), and gateway (172.23.199.2 and 172.23.200.2).

### Configure the rest of the switch IP addresses in the cabinet

To configure the next switch, follow [Step 4](#) through [Step 6](#) in “Configure the IP address” on [page 134](#).

Use the IP address of (172.23.199.23), subnet mask (255.255.255.0), and gateway (172.23.199.2).

### **Install Fabric Manager and Device Manager**

To install Fabric Manager and Device Manager:

1. Open your web browser.
2. Enter the IP address of the switch into the address bar.
3. Follow the prompts and accept all defaults to install both Fabric Manager and Device Manager.

### **Configure a VSAN followed by a domain**

To configure a VSAN:

1. Open the Device Manager for the switch with an IP address of **172.23.199.22**.
2. Open the **VSAN** dialog box by selecting the **VSAN** menu item.
3. Click **Create**.
4. Enter the value **100** in the **VSAN ID** field.
5. Set the **VSAN Name** to "**Red\_Vsan\_100**".
6. Use the default interop mode.
7. Click **Create**.
8. Enter the value **200** in the **VSAN ID** field.
9. Set the next **VSAN Name** to be "**Green\_Vsan\_200**".
10. Click **Create**.
11. Enter the value **300** in the **VSAN ID** field.
12. Set the next **VSAN Name** to "**Blue\_VSAN\_300**".
13. Click **Create**, and then click **Close**.



14. From the **Device Manager** menu, select **FC/Domain Manager/Configuration** and set a static Domain ID for the switches as shown in [Table 2](#) and [Table 3](#).

**Table 2** 172.23.199.22 through 172.23.199.27

IP	Domain	VSAN_ID	VSAN_ID	VSAN_ID
172.23.199.22	1	100	200	300
172.23.199.23	3	100	200	300
172.23.199.24	5	100	200	300
172.23.199.25	7	100	200	300
172.23.199.26	9	100	200	300
172.23.199.27	11	100	200	300

**Table 3** 172.23.200.22 through 172.23.200.27

IP	Domain	VSAN_ID	VSAN_ID	VSAN_ID
172.23.200.22	2	100	200	300
172.23.200.23	4	100	200	300
172.23.200.24	6	100	200	300
172.23.200.25	10	100	200	300
172.23.200.26	12	100	200	300
172.23.200.27	14	100	200	300

### Configure FC switches

To configure FC switches:

1. Configure the switch ports.
  - a. Open the Device Manager of the switch with an IP address of 172.23.199.22, or the next switch, by double-clicking its icon in Fabric Manager.
  - b. From the **Configure** menu, select the **Device menu** item.

- c. Admin up and configure the ports as shown in the following tables:

Slot #	Port #	Name	VSAN ID
1	1	TE ISL to Domain 5	1
1	2	TE ISL to Domain 7	1
1	3	TE ISL to Domain 2	1
1	4	TE ISL to Domain 4	1
1	5	TE ISL to Domain 3	1
1	6	TE ISL to Domain 6	1
1	7	TE ISL to Domain 10	1
1	8		

Slot #	Port #	Name	VSAN ID
2	1	TE ISL to Domain 5	1
2	2	TE ISL to Domain 7	1
2	3	TE ISL to Domain 2	1
2	4	TE ISL to Domain 4	1
2	5	TE ISL to Domain 3	1
2	6	Green Host HBA 1	200
2	7		
2	8	TE ISL to Domain 6	1
2	9	TE ISL to Domain 10	1

- d. Following [Step a](#) through [Step c](#), configure the ports of the switch with an IP address of 172.23.199.23 as shown in the following tables:

Slot #	Port #	Name	VSAN ID
1	1	TE ISL to Domain 9	1
1	2	TE ISL to Domain 11	1
1	3	TE ISL to Domain 2	1
1	4	TE ISL to Domain 4	1
1	5	TE ISL to Domain 1	1
1	6	TE ISL to Domain 14	1
1	7	TE ISL to Domain 12	1
1	8		

Slot #	Port #	Name	VSAN ID
2	1	TE ISL to Domain 9	1
2	2	TE ISL to Domain 11	1
2	3	TE ISL to Domain 2	1
2	4	TE ISL to Domain 4	1
2	5	TE ISL to Domain 1	1
2	6	Blue Storage 1	300
2	7	Blue Host HBA 1	300
2	8	TE ISL to Domain 14	1
2	9	TE ISL to Domain 12	1

- e. Following [Step a](#) through [Step c](#), configure the ports of the switch with an IP address of 172.23.200.22 as shown in the following tables:

Slot #	Port #	Name	VSAN ID
1	1	TE ISL to Domain 6	1
1	2	TE ISL to Domain 10	1
1	3	TE ISL to Domain 1	1
1	4	TE ISL to Domain 3	1
1	5	TE ISL to Domain 4	1
1	6	TE ISL to Domain 5	1
1	7	TE ISL to Domain 7	1
1	8		

Slot #	Port #	Name	VSAN ID
2	1	TE ISL to Domain 6	1
2	2	TE ISL to Domain 10	1
2	3	TE ISL to Domain 1	1
2	4	TE ISL to Domain 3	1
2	5	TE ISL to Domain 4	1
2	6	Green Host HBA 2	200
2	7		
2	8	TE ISL to Domain 5	1
2	9	TE ISL to Domain 7	1

- f. Following [Step a](#) through [Step c](#), configure the ports of the switch with an IP address of 172.23.200.23 as shown in the following tables:

Slot #	Port #	Name	VSAN ID
1	1	TE ISL to Domain 1	1
1	2	TE ISL to Domain 2	1
1	3	TE ISL to Domain	1
1	4	TE ISL to Domain	1
1	5	TE ISL to Domain 2	1
1	6	TE ISL to Domain 11	1
1	7	TE ISL to Domain 9	1
1	8		

Slot #	Port #	Name	VSAN ID
2	1	TE ISL to Domain 12	1
2	2	TE ISL to Domain 14	1
2	3	TE ISL to Domain 1	1
2	4	TE ISL to Domain 3	1
2	5	TE ISL to Domain 2	1
2	6	Blue Host HBA 2	300
2	7	Blue Storage 2	300
2	8	TE ISL to Domain 9	1
		TE ISL to Domain 11	1

- g. Following [Step a](#) through [Step c](#), configure the ports of the switch with an IP address of 172.23.199.24 as shown in the following table:

Slot #	Port #	Name	VSAN ID
1	1	TE ISL to Domain 2	1
1	2	TE ISL to Domain 2	1
1	3	TE ISL to Domain 1	1
1	4	TE ISL to Domain 1	1
1	5	Red Host HBA 1	100
1	6		
1	7		
1	8		

- h. Following [Step a](#) through [Step c](#), configure the ports of the switch with an IP address of 172.23.200.24 as shown in the following table:

Slot #	Port #	Name	VSAN ID
1	1	TE ISL to Domain 1	1
1	2	TE ISL to Domain 1	1
1	3	TE ISL to Domain 2	1
1	4	TE ISL to Domain 2	1
1	5	Red Host HBA 2	100
1	6		
1	7		
1	8		

- i. Following [Step a](#) through [Step c](#), configure the ports of the switch with an IP address of 172.23.199.25 as shown in the following table:

Slot #	Port #	Name	VSAN ID
1	1	TE ISL to Domain 2	1
1	2	TE ISL to Domain 2	1
1	3	TE ISL to Domain 1	1
1	4	TE ISL to Domain 1	1
1	5	Red Storage 1	100
1	6		
1	7		
1	8		

- j. Following [Step a](#) through [Step c](#), configure the ports of the switch with an IP address of 172.23.200.25 as shown in the following table.

Slot #	Port #	Name	VSAN ID
1	1	TE ISL to Domain 1	1
1	2	TE ISL to Domain 1	1
1	3	TE ISL to Domain 2	1
1	4	TE ISL to Domain 2	1
1	5	Red Storage 2	100
1	6		
1	7		
1	8		

- k. Following [Step a](#) through [Step c](#), configure the ports of the switch with an IP address of 172.23.199.26 as shown in the following table:

Slot #	Port #	Name	VSAN ID
1	1	TE ISL to Domain 4	1
1	2	TE ISL to Domain 4	1
1	3	TE ISL to Domain 3	1
1	4	TE ISL to Domain 4	1
1	5	Green Storage 1	200
1	6		
1	7		
1	8		

- l. Following [Step a](#) through [Step c](#), configure the ports of the switch with an IP address of 172.23.200.26 as shown in the following table:

Slot #	Port #	Name	VSAN ID
1	1	TE ISL to Domain 3	1
1	2	TE ISL to Domain 3	1
1	3	TE ISL to Domain 4	1
1	4	TE ISL to Domain 4	1
1	5	Green Storage 4	1
1	6		
1	7		
1	8		



- m. Following [Step a](#) through [Step c](#), configure the ports of the switch with an IP address of 172.23.199.27 as shown in the following table:

Slot #	Port #	Name	VSAN ID
1	1	TE ISL to Domain 4	1
1	2	TE ISL to Domain 4	1
1	3	TE ISL to Domain 3	1
1	4	TE ISL to Domain 3	1
1	5	Green Storage 2	200
1	6		
1	7		
1	8		

- n. Following [Step a](#) through [Step c](#), configure the ports of the switch with an IP address of 172.23.200.27 as shown in the following table:

Slot #	Port #	Name	VSAN ID
1	1	TE ISL to Domain 3	1
1	2	TE ISL to Domain 3	1
1	3	TE ISL to Domain 4	1
1	4	TE ISL to Domain 4	1
1	5	Green Storage 3	200
1	6		
1	7		
1	8		

### Connect cables

To connect the cables:

1. Connect ISLs.
  - a. Attach Fiber cable between switches as shown in [Figure 13 on page 101](#).

- b. After all cables are connected, use Fabric Manager to verify that all ISL connections are up.
  - c. Re-arrange icons to accurately reflect the switch configuration.
2. Connect host and storage ports.
  - a. Attach fiber cable between the switches and N\_Ports.

### Configure domains

Using Fabric Manager set the static Domain IDs for each VSAN and switch.

### Zone hosts and storage

To zone hosts and storage:

1. Open the **Zoning** dialog box in Connectrix Manager by right-clicking the appropriate fabric topology and selecting the **Zoning** menu item.
2. Create a zone by clicking **New Zone** under the **Zones Tree**.
3. Provide a descriptive name for the zone. This example zones "Red host HBA 1" and "Red Storage 1". Type "**RedHBA1\_1470\_8aa**" and press **Enter**.
4. Select "**Red Host HBA 1**" (WWPN 10000000c938e554) in the **potential zone members list**.
5. Click the right-pointing arrow on the divider between the **potential members list**, and the **zones list** to add the HBA to the zone.
6. Select "**Red Storage 1**" (WWPN 50060482cc19bf87) in the **potential zone members list**.
7. Click the right- pointing arrow on the divider between the **potential members list** and the **zones list** to add the storage port to the zone.
8. Repeat [Step 2](#) through [Step 6](#) for all host and storage pairs in the environment.
9. Create a zone set by clicking **New Set** under the **Zone Sets Tree**.
10. Provide a descriptive name for the zone set as shown in the example following [Step 12](#).

11. Add all of the new zones to the zone set of the proper VSAN.  
When completed, the zone sets should be similar to what is shown in the example following these steps.
12. Activate the **VSAN Zone Set**.

```

Zone set name = "Red_Oct_31_06_1140"

Zone name = "RedHBA1_1470_8aa"
Zone Member = "10000000c938e554"
Zone Member = "50060482cc19bf87"

Zone name = "RedHBA2_1470_9aa"
Zone Member = "10000000c938e555"
Zone Member = "50060482cc19bf88"

Zone set name = "Green_Oct_31_06_1140"

Zone name = "GreenHBA1_AllGreenStorage"
Zone Member = "10000000c939a051"
Zone Member = "50060482cc19c407"
Zone Member = "50060482cc19c408"
Zone Member = "50060482cc19c4c7"
Zone Member = "50060482cc19c4c8"

Zone name = "GreenHBA2_AllGreenStorage"
Zone Member = "10000000c939a052"
Zone Member = "50060482cc19c407"
Zone Member = "50060482cc19c408"
Zone Member = "50060482cc19c4c7"
Zone Member = "50060482cc19c4c8"

Zone set name = "Blue_Oct_31_06_1140"

Zone name = "BlueHBA1_1489_8aa"
Zone Member = "210100e08b8ac76d"
Zone Member = "50060482cc19c447"

Zone name = "BlueHBA2_1489_9aa"
Zone Member = "210100e08baac76d"
Zone Member = "50060482cc19c448"

```

### Complete the SAN setup

At this point the SAN is ready to pass I/O from host to storage. Other steps, such as configuring LUN masking and modification of host configuration files, are required before the SAN setup is complete. Refer to the OS configuration guide for more details.

## Heterogeneous switch interoperability

As SANs become larger and topologically more complex, it is increasingly advantageous to use switches from different vendors that work together. *Interoperability* is the term used to describe a Fibre Channel fabric that contains switches from more than one vendor. Interoperability has countless definitions in the IT connectivity space. For the context of this chapter, it refers to FC switch interoperability, which increases the ability and flexibility to design complex SANs. It takes into account the various features that different vendor switches provide as well as the specific features that end users look for while designing SANs.

This section provides an in-depth description for setting up an EMC-supported Fibre Channel SAN comprising of FC switches (director class, distribution class, departmental switches and blade server FC switch modules) from different switch vendors, such as Brocade, Cisco, and Brocade M series. Switch migration procedures to move customers' existing SAN topology from one EMC-supported switch vendor type to another EMC-supported vendor type are also provided.

Support for Native Connectivity between Brocade B-series and M-series platforms remains largely unchanged up to and including FOS v6.4.x. All previously supported configurations and capabilities carry forward to FOS v6.4. Also, there are no restrictions or limitations to interop support due to the addition of the FC8-64 blade.

FOS v7.x does not support direct E\_port connectivity with an M-EOS platform either in interop mode 2 or interop mode 3. A switch that is running FOS v7.x can only operate in Brocade Native mode (Interop mode 0).

Connectivity between a fabric that has FOS v7.x switch and an M-EOS edge fabric can still be achieved via Brocade FCR.

---

**Note:** EX\_ports on FOS v7.1 platforms support only interop mode 0. Users must use pre-FOS v7.1 if they need to configure EX\_ports in interop mode 2 or interop mode 3.

---

Refer to [Table 4 on page 190](#) for a list of all the supported switch operating modes that must be set for interop connectivity between multi-vendor switches.

The information provided in this section has been obtained from vendor product documents and testing experiences in the EMC E-Lab qualification labs.

---

## Interoperability overview

This section explains the concept and significance of heterogeneous interoperability in the FC SAN world.

### E\_Port interoperability

The approval of the FC-SW-2 standard has created an open standard for switch-to-switch communication, allowing end-users to select best-in-class products with assurance that these products will work together. Currently, most SAN switch vendors offer FC-SW-2 compliant switches including, but not limited to, Brocade, Cisco, and Brocade M series.

Multi-vendor FC switch interoperability is also referred to as E\_Port interoperability, since two switches are connected to each other using their respective E\_Ports per FC standards.

### Heterogeneous interoperability in EMC context

EMC supports the following vendor switches in a SAN deploying switch interoperability: Brocade, Cisco, and Brocade M series. Most of the E\_Port switch connectivity-based testing between these multi-vendor switches, for different kinds of switch hardware and firmware, is conducted in the EMC E-Lab qualification labs.

The testing mainly involves the validation of:

- ◆ Link initialization between switches from different vendors
- ◆ Name server distribution
- ◆ Zoning changes
- ◆ Routing protocols
- ◆ Fabric management application.

Interoperability support for switches that are compatible with the FC-SW standard is listed in the "Switched Fabric Topology Parameters" section of the *EMC Support Matrix*.

This section comprises of the following attributes (or columns):

- ◆ Name of the switch
- ◆ Name of the switches it is interoperable with
- ◆ Firmware versions tested on the switch and the interoperable switches

- ◆ Switch (fabric) management application revision
- ◆ Maximum number of domains per fabric
- ◆ Maximum number of hops
- ◆ Maximum domain-to-domain ISLs supported by EMC

This information is based on the qualifications conducted at E-Lab.

Although the support matrix can only directly represent a two-way interop, it may also be used to represent a three-way interop. Consider there are three switches, each from a different vendor, running firmware versions A, B, and C, respectively.

If the following entries are present in the support matrix:

- ◆ A is interoperable with B
- ◆ A is interoperable with C
- ◆ B is interoperable with C

then one can infer that switches running firmware versions A, B, and C are interoperable and can co-exist in the same fabric.

### Significance of a multi-vendor switch configuration

There are at least three user scenarios that indicate the significance of a multi-vendor switch environment:

#### User scenario 1: Complex SANs

As SANs become larger and topologically more complex, it is getting increasingly useful to get switches from different vendors to work together. As discussed above, there are various vendors in the FC industry, designing and manufacturing switch hardware with different hardware features, protocol related-features, and performance (2 G/4 G/10 G). In such a case, end users must be able to select the best in-class products to design a SAN that meets their needs in terms of features and performance.

#### User scenario 2: Blade servers

Blade server technology is becoming widely accepted, but it is still evolving. At the time of this publication, only some switch vendors, such as Brocade, and Brocade M series are manufacturing FC switch modules that can be plugged into the backend of a blade server. A blade server accesses a switched fabric or storage it is being hooked up to using FC switch modules (as discussed in detail in "Blade servers" in the *Non-EMC SAN Products Data Reference Manual*, available on the [E-Lab Navigator](#), **Documents > Topology Resource Center**). Hence, if end users have a switch-vendor fabric which is

neither of the above switch module manufacturing vendors, interoperability between multi-vendor switches is an important aspect to consider before building such a fabric.

### User scenario 3: Switch migration

Depending on their requirements, some end users move their SAN from one switch vendor type to another switch vendor type. This switch migration procedure needs to be executed in phases. One of the transitional phases requires the vendor type switches (one type before and the other after the migration) to be in the same fabric, making interoperability between multi-vendor switches an aspect to be considered before starting a switch migration procedure.

---

## Heterogeneous SAN design

This section provides details on setting up an interoperable SAN environment, and studies different combinations of some of the EMC-supported heterogeneous SANs.

In addition to providing a detailed step-by-step approach to design some specific type of SANs, this section is intended to give you an insight into the kind of configurations previously tested and supported by E-Lab, and to provide some of the best practices and caveats that must be considered while designing these SANs.

### Components needed to create a heterogeneous design

The concept of simple SANs and Complex SANs has been discussed in detail in [“Best practices” on page 88](#). Most of the discussed topologies comprising of at least two switches are applicable to designing heterogeneous SANs. It is important to note that:

- ◆ In the heterogeneous SANs, the switched fabric must be comprised of switches from two or more different vendors.
- ◆ It is essential to set the appropriate operational mode on the switches so that they can communicate with any other switch from a different vendor in the same fabric.

---

## How to set up an interoperable switched fabric topology

In order to address the individual vendor switch settings for supported switch interoperability configurations, a cookbook approach is used to explain a seven phase switch vendor migration process. The assumptions for this seven phase migration process follows.

**Assumptions for the 7 phase migration process**

- ◆ In all cases, the case studies start with a homogeneous single vendor fabric.
- ◆ Compound core edge topology is used for these case studies.
- ◆ The switch migration process involves a step to move the edge switches from the same vendor type core switches to a different vendor type core switches.

---

**Note:** This process is helpful to end users who want to move their fabrics from one vendor switch core type to another vendor switch core type.

---

- ◆ One of the steps in the migration involves a stage when different vendor switches co-exist in the same fabric.

---

**Note:** This can help end users set up a supported interoperability fabric based on the switch configuration settings from different vendors that can co-exist in a stable fabric.

---

**Phases**

The seven phase process that has been appointed across all the case studies is explained in detail in the following sections. This example shows a transition from a switched fabric topology using Vendor A switches, to a topology with Vendor B switches.

**Phase 1:**

This is the base configuration phase with the pre-existing Vendor A switched fabric topology. Some default settings on these switches need to be changed before introducing a Vendor B switch in this fabric.

**Phase 2:**

In this phase, a Vendor B core switch is added to the homogeneous Vendor A fabric. This implies that all the Vendor A edge switches now have ISLs to this core Vendor B switch.

**Phase 3:**

Half of the zoned initiator (hosts) and target (storage) pairs from a Vendor A core switch are moved to the Vendor B core switch.

**Phase 4:**

All the other initiator-target pairs from the Vendor A core switches are completely moved to the Vendor B core switch.



---

**Note:** Phase 3 and 4 are executed in steps to avoid any disruption of traffic between the host and storage during the transition and to avoid any downtime.

---

**Phase 5:**

One or more Vendor B switches are added to the edge of the fabric.

**Phase 6:**

Hosts and storage are moved from any Vendor A edge switches to the new Vendor B edge switches.

**Phase 7:**

More Vendor B switches are added, if required, to the core of the fabric.

**Case studies**

Based on the switch vendor types and the interoperability modes currently tested and supported by EMC, five case studies are described in this section.

*Case study #1* **Migrating from a Connectrix B homogeneous fabric in Connectrix B native mode to a Connectrix MDS fabric**

**Assumptions specific to this case study:**

- ◆ Interoperability mode settings on the switches are:
  - *Brocade native mode (interopmode 0)* mode on the Connectrix B switches in the fabric.
  - *Cisco Interop-3* mode on the Connectrix MDS switches in the fabric.
- ◆ Fabric management applications used for managing the fabric:
  - Cisco Fabric Manager, Brocade Web Tools.

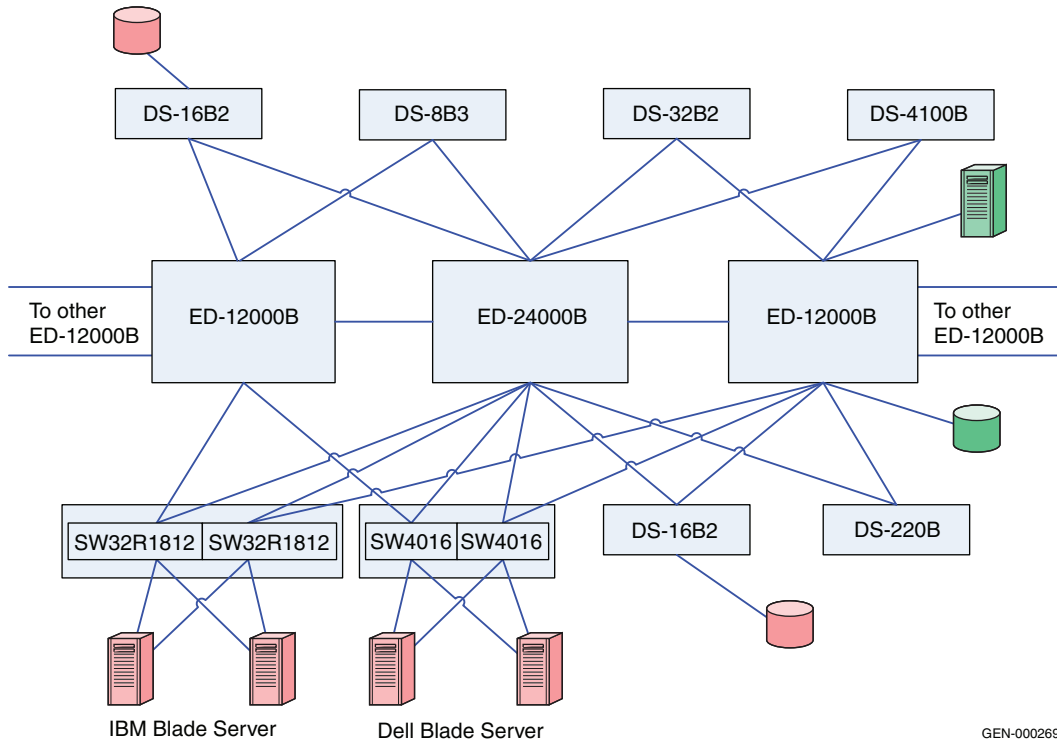
---

**Note:** The Brocade Fabric Manager is not supported for management of a Connectrix B-Connectrix MDS heterogeneous fabric.

---

## Phase 1: Base configuration - Pre-existing Connectrix B core-edge fabric

### Topology



GEN-000269

**Figure 17 Phase 1: Basic configuration**

As shown in [Figure 17](#), the two ED-12000Bs and the ED-24000B director switches are at the core of the fabric while the departmental switches, the DS-4100B, DS-220B, DS-16B3, DS-8B2, and Connectrix B-based blade server switch modules from IBM (IBM 32R1813) and Dell (SW4016), are at the edge. Please refer to the “Switched Fabric Topology Parameters” section of the *EMC Support Matrix* to obtain a list of all the other Brocade/EMC Connectrix B switches that can be supported in a heterogeneous set up as well as the operating modes specified above for this case study.

The specific configuration settings, best practices, host and storage layouts, and topology design to withstand failures for a Connectrix B core-edge homogeneous fabric discussed in the “[Connectrix B](#)

example” on page 90 applies to this base topology.

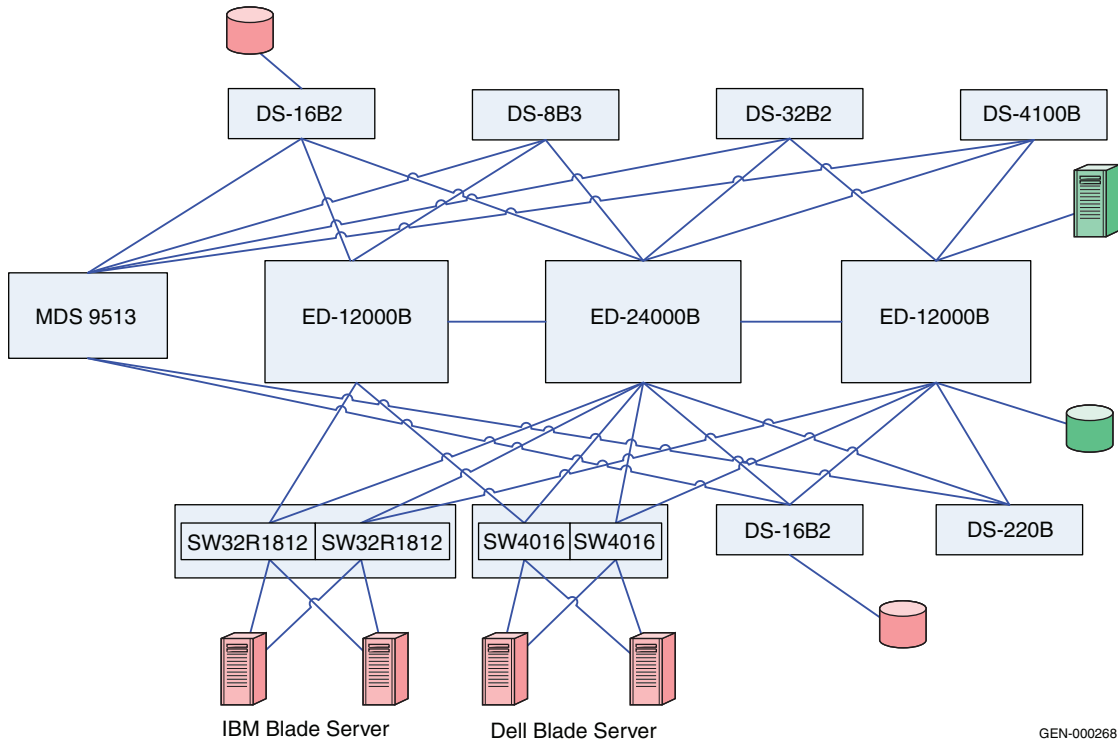
### Checkpoints

Before adding the Cisco core Connectrix MDS 9513 director to the homogeneous Connectrix B fabric, verify the following fabric characteristics using either the Connectrix B switch management application, Web Tools, or the Brocade CLI.

- ◆ **Proper distribution of the Name Server information.**  
Verify that all the host HBA ports and storage ports logged into the fabric are listed in the name server.
- ◆ **Proper distribution of Zoning information.**  
Verify that the active zoneset comprises of the zones that contain the desired mapping of host and storage ports.
- ◆ **Proper display of fabric and N\_Port elements on the management applications.**  
Verify that the physical topology and domain count of the fabric is as desired.
- ◆ **No disruption in data transfer.**  
Verify that the data traffic is running appropriately through the SAN.

## Phase 2: Adding the Connectrix MDS 9513 to the core of the fabric

### Topology



GEN-000268

**Figure 18** Phase 2: Adding Connectrix MDS 9513

As illustrated in [Figure 18](#), the Connectrix MDS 9513 director is added to the core of the fabric. Also it is important to note that there is no ISL between the Connectrix B cores and the Connectrix MDS 9513. Please refer to the “Switched Fabric Topology Parameters” section of the *EMC Support Matrix* to obtain a list of all the Connectrix MDS series switches that can be supported in a heterogeneous fabric with the Connectrix B switches, as well as the operating modes specified above for this case study.

Before adding the MDS 9513 to the fabric, the following steps need to be executed on the Connectrix MDS switch:

1. Create a VSAN on the Connectrix MDS switch.

A Virtual SAN can be used to create multiple logical SANs over the same physical infrastructure.

- a. Click **Create VSAN** on the Cisco Fabric Manager.
- b. When the **Create VSAN** window appears, check the switches that must be included in this VSAN. In this case, only the Connectrix MDS 9513 would be included.
- c. Fill in the **VSAN ID** field with an unused ID number and the VSAN name filed with an appropriate name. In this case study, the **VSAN ID number = 150** is assigned.
- d. For the VSN attributes leave everything other than the **InterOperValue** and the **AdminState** at default.
- e. Set **InterOperValue** to **Interop-3** and the **AdminState** to **Suspended**.
- f. Assign interfaces to the legacy switch interop mode 2 VSAN, i.e., statically assign VSAN membership for an interface using Fabric Manager, by selecting **Interfaces > FC Physical** from the **Physical Attributes** pane. The interface configuration in the **Information** pane appears. Click the **General** tab on this window and double-click and complete the **PortVSAN** field by entering the **VSAN ID** number (150) for every port desired to be used by this fabric.

For more details on the VSAN settings for Connectrix MDS 9000 Family switches refer to the Cisco document located at <http://www.cisco.com>.

For configuring Interop-3 specific settings on the Cisco CLI refer to the Cisco documentation located at <http://www.cisco.com>.

The legacy switch interoperability mode 3 for Connectrix B switches with more than 16 ports (and a core PID =1) was introduced with Connectrix MDS SAN-OS Release 1.3. With this VSAN-based interop mode, Connectrix B switches do not have to be altered from their native mode and can be seamlessly added to a new or existing Connectrix MDS SAN-OS VSAN.

2. Lower the Connectrix MDS 9000 switch ISL buffer-to-buffer credits (BB\_Credits) to match the Brocade BB\_Credits, because the Connectrix B switch cannot handle more than 16 BB\_Credits on an ISL.

3. Configure the Connectrix B switches.

Now that the Connectrix MDS 9513 director is configured, there will be no configuration or disruption to the Connectrix B fabric. All that is required is to enable the new ISL ports. A

**portcfgislmode** command must be run against all the ports on the Connectrix B edge switches that will be linked with an ISL to the Connectrix MDS 9513 core switch.

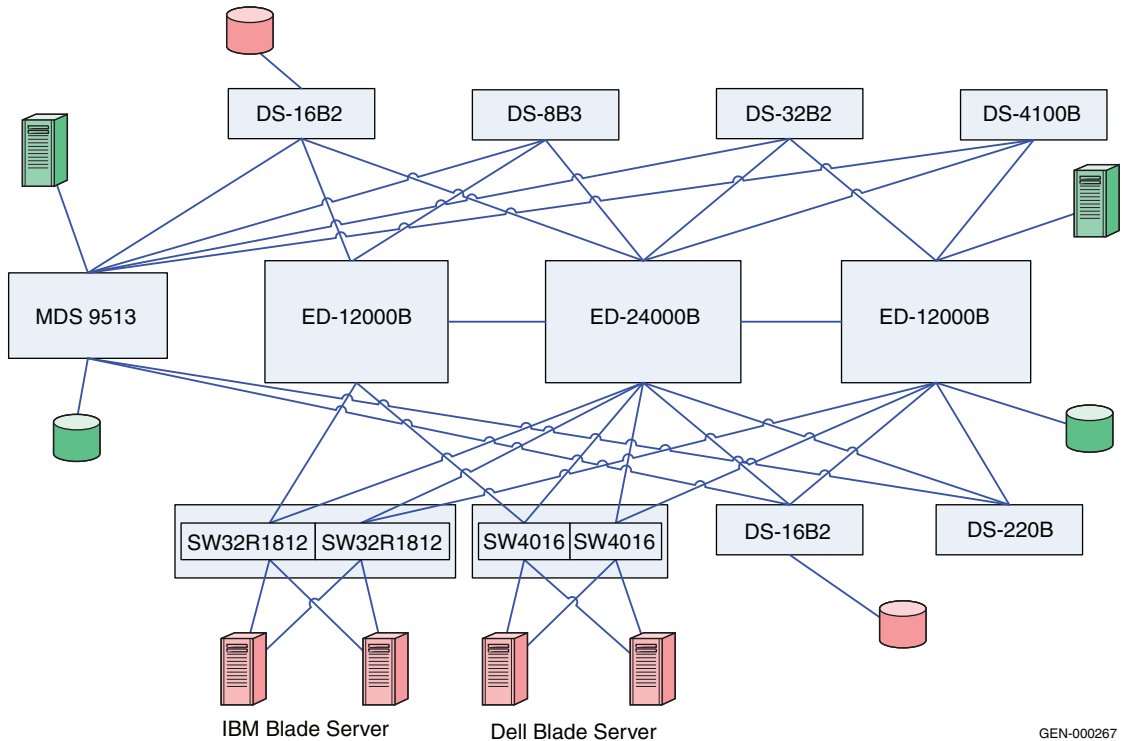
4. Before connecting the Connectrix MDS switch to the Connectrix B switches through an ISL, it is important to verify that all the Connectrix B switches and the Connectrix MDS switch are running the supported firmware versions for the respective interop modes. This can be checked by referring to the *EMC Support Matrix* entries for these switches
5. Create an ISL between the Connectrix MDS 9513 and all the Connectrix B edge switches as per [Figure 18 on page 157](#).
6. Verify connectivity of the fabric by validating the same set of “[Checkpoints](#)” on [page 156](#). In this case, the Cisco Fabric Manager and Cisco CLI can be used to verify the fabric topology and name server information, in addition to the Brocade CLI and Brocade Web Tools.

### **IMPORTANT**

**At the end of this phase, it is highly recommended to check that the zoning information has been distributed appropriately to the Connectrix MDS core. The active zoneset on the Connectrix B and Connectrix MDS fabric management tools must be compared to verify that there are no differences.**

### Phase 3: Moving half of the host and storage ports from the Connectrix B core to the Connectrix MDS 9513

#### Topology



GEN-000267

**Figure 19** Phase 3: Moving half the host and storage ports

As illustrated in [Figure 19](#), the same host-target pair is now connected to both core switches: the ED-12000B and the Connectrix MDS 9513.

This phase is an intermediate phase to completely moving the host and storage ports to the Connectrix MDS9513. This phase ensures that the zones are appropriately pushed from the Connectrix B to the Connectrix MDS switches and that the traffic between the host and storage zoned together is not disrupted. The host and storage ports can be moved successfully to another core switch without any downtime.

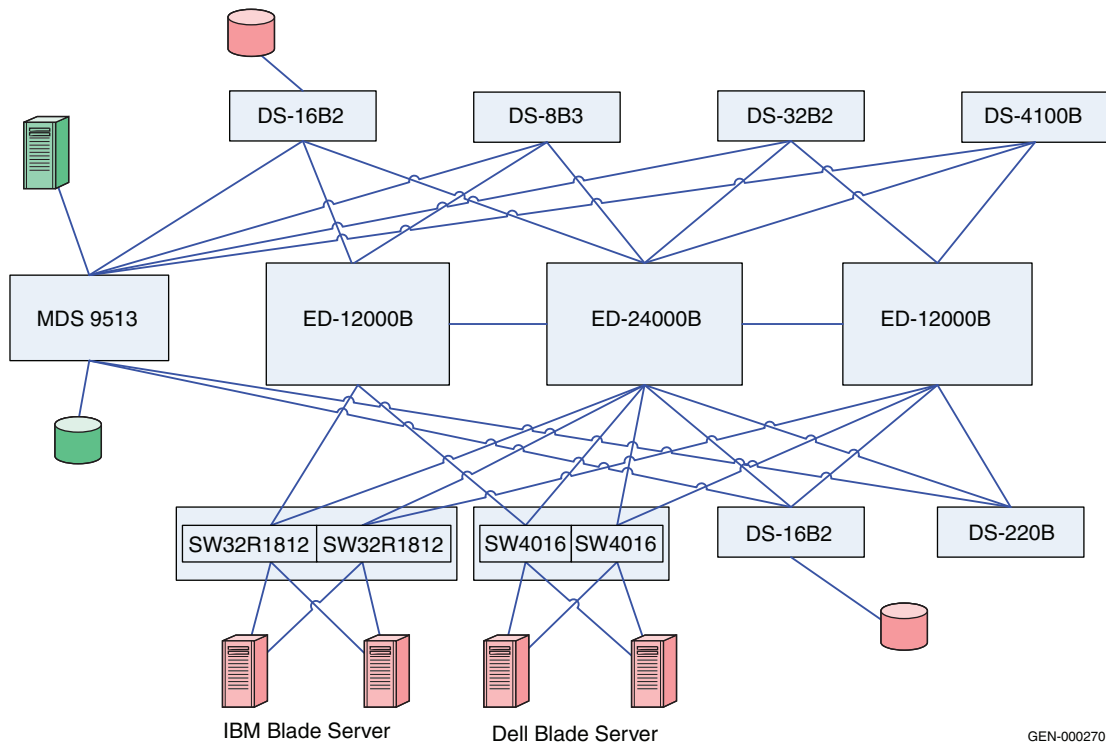
There are no specific steps to be executed at this phase other than physically pulling cables from one switch and plugging them into the



other core switch. However, it is recommended to follow the “Checkpoints” on page 156 and validate that this transition did not affect the connectivity and functioning of the fabric.

#### Phase 4: Completely moving the host and storage ports from the Connectrix B core to the Connectrix MDS 9513

##### Topology



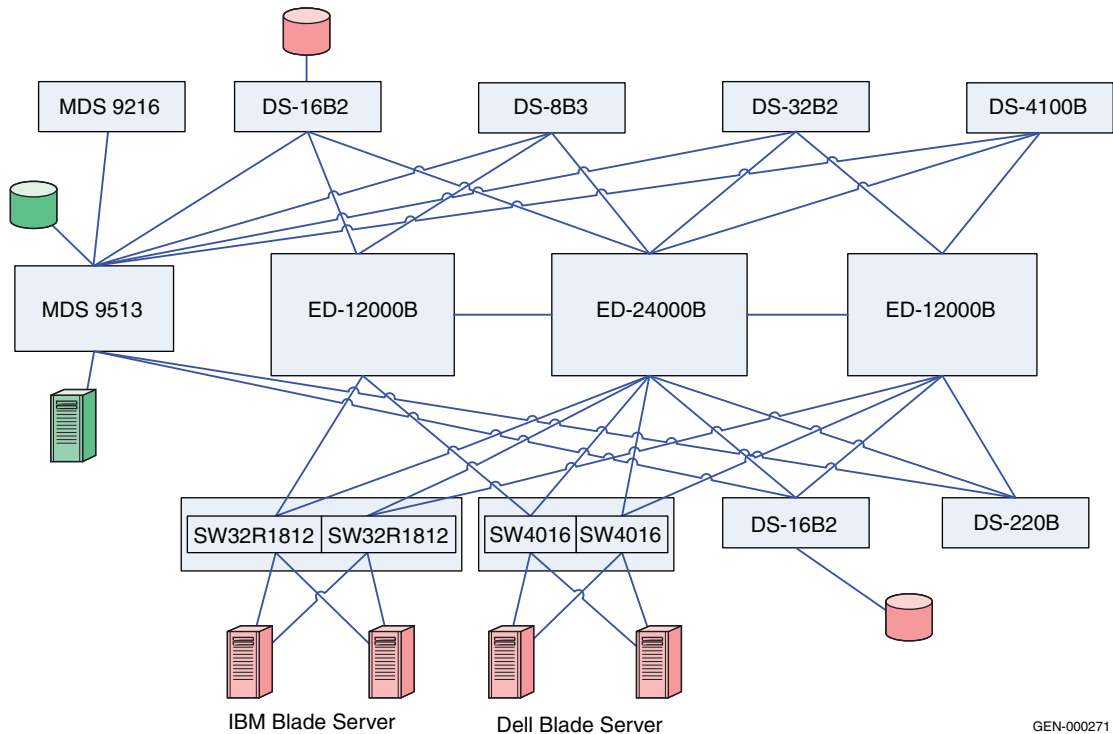
GEN-000270

**Figure 20 Phase 4: Completely moving host and storage ports**

This is an extension of “Phase 3: Moving half of the host and storage ports from the Connectrix B core to the Connectrix MDS 9513”. On validating that a stable fabric exists at the completion of Phase 3, Phase 4 can be executed by pulling the remaining host-storage pairs from the ED-12000B switch and transferring them to the Connectrix MDS 9513 core.

Again, it is recommended to go follow the “Checkpoints” on page 156 using both Connectrix B and Connectrix MDS supported fabric management tools.

### Phase 5: Adding an Connectrix MDS 9216 to the edge Topology



GEN-000271

**Figure 21 Phase 5: Adding Connectrix MDS 9216**

To configure the Connectrix MDS 9216:

1. Create a VSAN on the Connectrix MDS switch.

A Virtual SAN can be used to create multiple logical SANs over the same physical infrastructure.

- a. Select **Create VSAN** on the Cisco **Fabric Manager**.
- b. Once the **Create VSAN** window appears, check the switches that must be included in this VSAN. In this case, only the Connectrix MDS 9513 will be included.
- c. Fill in the **VSAN ID** field with an unused ID number and the **VSAN name** field with an appropriate name. In this case study, **VSAN ID number = 801** is assigned.

- d. For **VSAN** attributes, leave everything other than the **InterOperValue** and the **AdminState** at default. VSAN attributes can be seen in the information pane.
- e. Set **InterOperValue** to **Interop-1** and the **AdminState** to **Suspended**.
- f. Statically assign VSAN membership for an interface using Fabric Manager, by choosing **Interfaces > FC Physical** from the **Physical Attributes** pane. The interface configuration appears in the **Information** pane. Select the **General** tab on this window, and double-click and complete the **PortVSAN** field by entering the **VSAN ID** number (801) for every port desired to be used by this fabric.

For more details on the VSAN settings and for configuring the Interop-1 specific settings for Connectrix MDS 9000 family switches refer to the following Cisco documentation located at <http://www.cisco.com>.

2. To set port settings on the Device Manager for the Connectrix MDS 9513:
  - a. Select the desired switch, in this case, the Connectrix MDS 9513. The **Device Manager** for this switch appears.
  - b. Select the ports participating in this fabric, and then set the port speed to **automax 2 G**. Leave the other settings as default.
  - c. Set the **Admin** option to **up**.
3. Unsuspend the VSAN.


After configuring all the settings as stated above, go back to the **VSAN attributes**, and change the **AdminState** to **Active**.

4. Using an ISL, link the Connectrix MDS 9513 to all the Connectrix M edge switches.

When connecting these switches with matching VSAN IDs, the zoning information merges. After executing the steps above, you must link this edge switch with an ISL to both the Connectrix MDS and Connectrix M core switches.

## Phase 6: Moving hosts and storage to a new edge

## Topology



GEN-000272

**Figure 22** Phase 6: Moving hosts and storage to a new edge

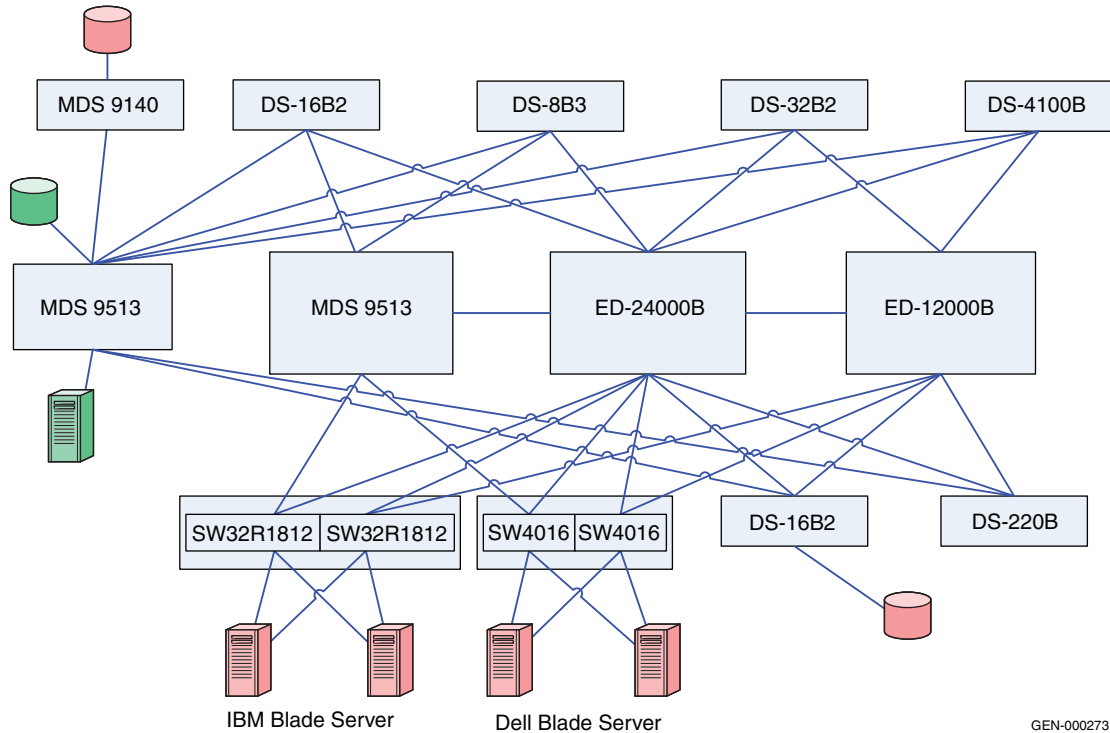
Hosts and storage ports shared by the other Connectrix B edge switches in the fabric (except for the Server Blade N\_Ports logged into the Brocade blade server switch modules) can be completely or partially moved to the Connectrix MDS 9216 edge switch. It is evident in this transitional phase that both Connectrix B and Connectrix MDS switches, can co-exist in a stable fabric with Connectrix B operating in Connectrix B native mode and Connectrix MDS operating in Interop-3 mode.

The settings on the switches in this phase can be used to implement configuration settings for a Connectrix B/Connectrix MDS interop fabric. All the non-default settings discussed in [“Phase 2: Adding the Connectrix MDS 9513 to the core of the fabric” on page 157](#) for Connectrix MDS 9513, Connectrix MDS 9216 and Connectrix B

switches apply to setting up a heterogeneous Connectrix MDS-Connectrix MDS fabric from the ground up.

### Phase 7: Adding a Connectrix MDS switch to the core

#### Topology



GEN-000273

**Figure 23** Phase 7: Adding Connectrix MDS switch to core example

As illustrated in [Figure 23](#), you can add a Connectrix MDS 9513 to the core with similar settings as the previous Connectrix MDS 9513, and you can also create an ISL to the existing core Connectrix MDS 9513 director and other edge switches in the fabric. It is important that you configure the new VSAN with the same VSAN ID and attributes as the existing VSAN to ensure a clean fabric merge.

### Complete migration to Connectrix MDS

At the end of case study #3 you have completed a migration from a Connectrix B-only fabric to a Connectrix B-Connectrix MDS fabric with Connectrix MDS switches at the core, connected through ISLs to every edge switch in the fabric. The host and storage ports (except for

the blade server host ports) can be completely moved to the Connectrix MDS switches as specified in Phase 3, 4 and 6. The Connectrix B edge switches (except for the blade server Brocade switch modules) and the Connectrix B core switches can now be pulled out of the fabric. A fully operational Connectrix MDS-only fabric now exists. This is a complete migration from one switch type (Connectrix B) to another type (Connectrix MDS).

### **IMPORTANT**

**In a Connectrix MDS-only fabric, it is recommended that all Connectrix MDS switches operate in their native (default) mode. Connectrix MDS switches are in interopmode at the end of the migration.**

You need not reboot the switch to change the interopmode on a Connectrix MDS. The **Interop mode** attribute on currently active VSANs must be changed from **Interop-3** to **Default**.

You cannot modify active zones that were pushed to Connectrix MDS switches from Connectrix B switches; zoning must be reconfigured. It is highly recommended that you backup the configuration to avoid losing all zoning information in the active zoneset on the Connectrix MDS switches.

### **Warnings or caveats**

Consider the following:

- ◆ Refer to EMC Knowledgebase solution emc149735 for all interop issues.
- ◆ In MDS interop mode 3 and Brocade Native mode, a host attached to a Brocade switch does not receive RSCN after zoneset activation/de-activation and therefore will not immediately discover newly added LUNs. Zoneset activation scenarios include a new zoneset activation and modification of existing zoneset and reactivation. A zoneset may be a regular zoneset or an IVR zoneset. This issue is seen on EMC-supported Brocade releases: v5.3.x; v6.1.x. The workaround is to bounce (disable and enable) the host port attached to Brocade switches after zoneset activation in order to update the devices. (Only the Host port that added the new target port needs to be bounced.)
- ◆ Zoning changes cannot be activated from a Brocade switch. The workaround is to use MDS switches to activate zoning changes. This caveat is specific to EMC-supported Brocade FOS v6.1.x.

*Case Study #2***Setting up a heterogeneous switched fabric with Connectrix B switches in the interop mode (interopmode 1), and the Connectrix MDS switches in Interop-1 mode****Assumptions specific to this case study:**

- ◆ Interoperability mode settings on the switches:
  - Interop mode: *interopmode 1* on the Connectrix B switches in the fabric.
  - Cisco *Interop-1* mode on the Connectrix MDS switches in the fabric.
- ◆ Fabric management applications used for managing the fabric:
  - Cisco Fabric Manager
  - Web Tools only for Connectrix B switches

**IMPORTANT**

Brocade's interopmode 1 has been replaced with interopmode 3 on Brocade FOS v6.0.x and higher. Therefore, if the Brocade switch is running FOS v6.0.x or higher, the following assumptions will apply to this case study:

**For Interoperability mode settings on the switches:**

- Interop mode: *interopmode 3* on the Connectrix B switches in the fabric
- Cisco *Interop-1* mode on the Connectrix MDS switches in the fabric

**For Fabric management applications used for managing the fabric:**

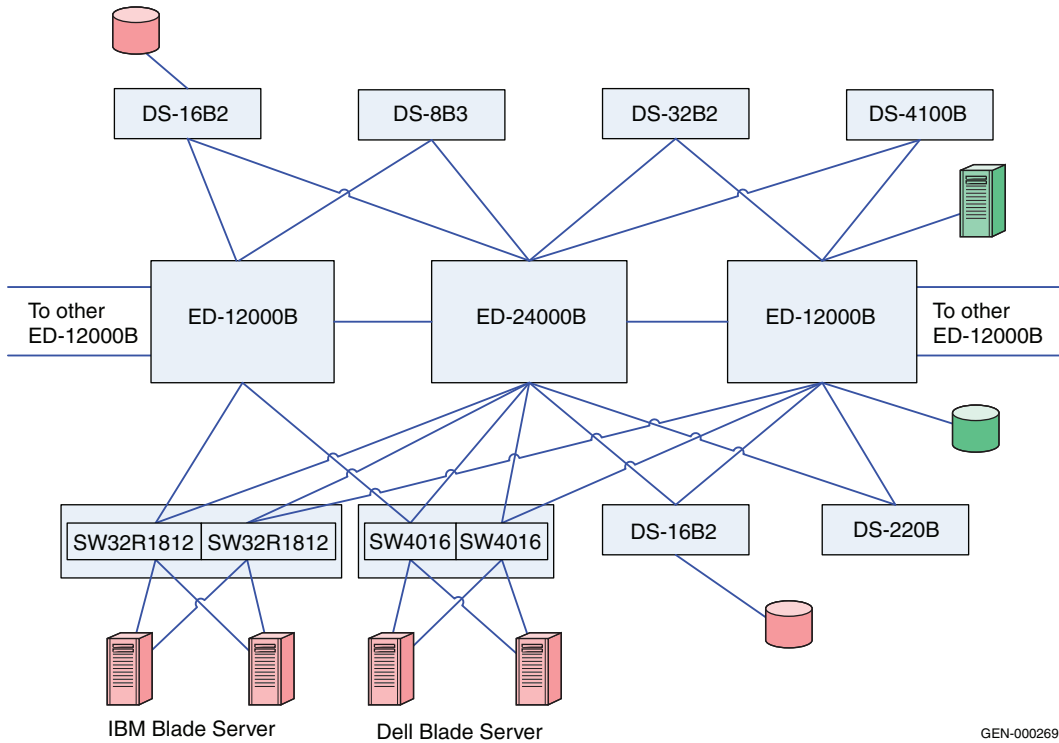
- Cisco Fabric Manager
- Web Tools only for Connectrix B switches

Refer to the *EMC Support Matrix* for the latest supported firmware versions on the Cisco MDS switches for interop with Brocade/Connectrix B switches running FOS v6.0.x and higher.

When migrating from a Connectrix B to a Connectrix MDS fabric, keep the fabric in Connectrix B native mode and follow the migration process explained in “[Case study #1](#)” on page 154. This case study shows the settings that you must configure on Connectrix B and Connectrix MDS switches before they are linked with an ISL. This is a required step when migrating from a Connectrix B to a Connectrix MDS fabric.

## Phase 1: Base configuration – Pre-existing Connectrix B core-edge fabric

### Topology



GEN-000269

**Figure 24** Phase 1: Basic configuration

As shown in [Figure 24 on page 168](#), the two ED-12000B, and ED-24000B director switches are at the core of the fabric. Edge hardware includes DS-4100B, DS-220B, DS-16B3, DS-8B2, Brocade-based blade server switch modules from IBM (IBM 32R1813), and Dell (SW4016). Refer to the “Switched Fabric Topology Parameters” section of the *EMC Support Matrix* for a list of other Brocade/EMC Connectrix switches supported in a heterogeneous setup, and for relevant operating modes.

The following information discussed in the “[Connectrix B example](#)” on [page 90](#) applies to this base topology:

- ◆ Specific best practices
- ◆ Host and storage layouts



- ◆ Design to withstand failures for a Connectrix B core-edge homogeneous fabric

This case study, unlike “[Case study #1](#)” on page 154, sets Connectrix B switches to interop mode, links them with an ISL to Connectrix MDS switches (operating in their supported Interop-1 mode), and then pulls out the Connectrix B cores from the fabric.

Since this is a Connectrix B-only fabric with all switches operating in Connectrix B native mode by default, the following non-default settings must be configured on the Connectrix B switches using the Brocade CLI:

1. Telnet into all the Connectrix B switches in the fabric; one at a time.
2. Run the **interopmode** command at the switch command prompt.
3. If interopmode is **On**, the switch can become part of a fabric that can accept the addition of a Connectrix MDS. If it is **Off**, perform [Step 4](#) through [Step 8](#).
4. Disable the switch by running the **switchdisable** command at the switch prompt.
5. Issue the **interopmode 1** command.
6. Enter a confirmation yes: **y** to switch the interopmode to **on**. Enter **y** in the warning box.
7. Reboot the switch to activate configuration settings.
8. After reboot, repeat [Step 1](#) through [Step 3](#) to verify that interopmode has been set to **On**.
9. Run the **version** command to verify that the switch is running the supported firmware version for interop with Connectrix MDS.
10. The **msplmgmtdeactivate** command must explicitly be run prior to connecting from a Connectrix B switch to a Connectrix MDS 9000 Family switch.

### Checkpoints

Before adding the Connectrix MDS core Connectrix MDS 9513 director to the Connectrix B fabric, verify the following fabric characteristics using either the Brocade switch management application, Web Tools, or the Brocade CLI.

- ◆ **Proper distribution of the Name Server information.**

Verify that the all host HBA ports and storage ports that are logged into the fabric are listed in the name server.

- ◆ **Proper distribution of Zoning information.**

Verify that the active zoneset comprises zones that contain the desired mapping of host and storage ports.

- ◆ **Proper display of fabric and N\_Port elements on the management applications.**

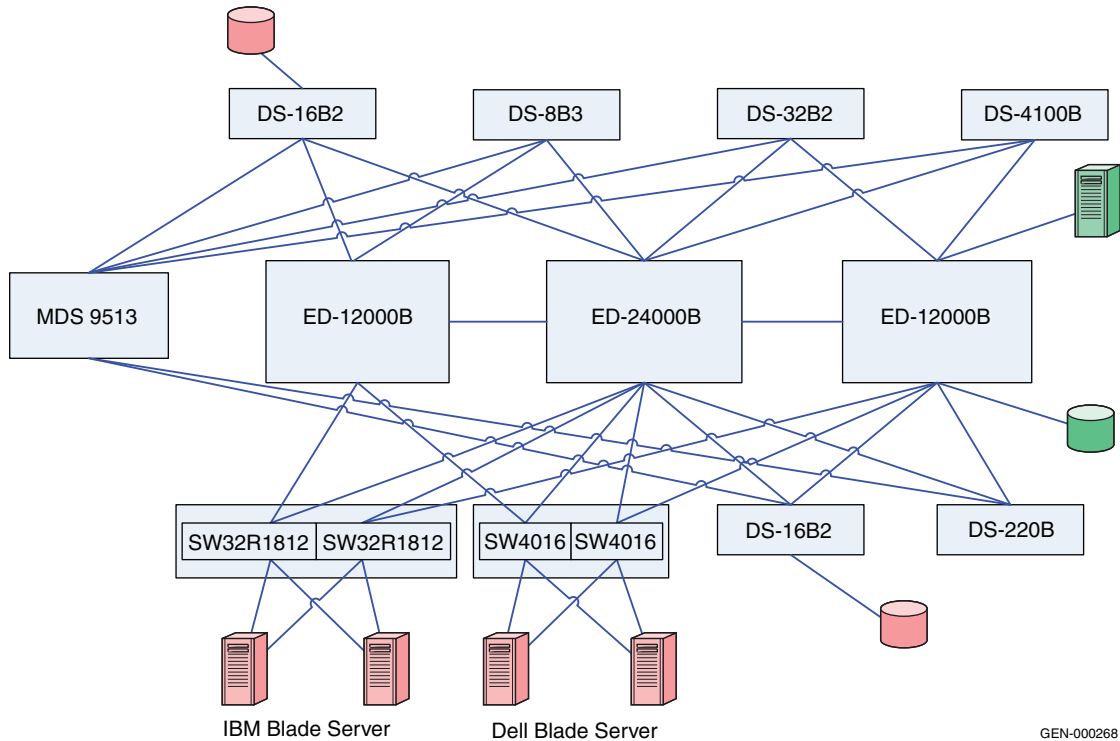
Verify that the physical topology and domain count of the fabric are correct.

- ◆ **No disruption in data transfer.**

Verify that the data traffic is running appropriately through the SAN.

## Phase 2: Adding the Connectrix MDS 9513 to the core of the fabric

### Topology



**Figure 25** Phase 2: Adding Connectrix MDS 9513

As shown in [Figure 25](#), the Connectrix MDS 9513 director is added to the core of the fabric. It is important to note that there is no ISL between the Connectrix B cores and the Connectrix MDS 9513. Please refer to the “Switched Fabric Topology Parameters” section of the [EMC Support Matrix](#) for a list of Cisco/Connectrix MDS switches that are supported in a heterogeneous configuration with Connectrix B switches. In addition, operating modes are listed.

Before adding a Connectrix MDS 9513 to the fabric, follow these steps for a Connectrix MDS switch:

1. Create a VSAN on the Connectrix MDS.

A Virtual SAN can be used to create multiple logical SANs over the same physical infrastructure.

  - a. Click **Create VSAN** on the **Cisco Fabric Manager**.
  - b. When the **Create VSAN** window appears, check the switches you want to include in the VSAN. In this case, only Connectrix MDS 9513 switches are included.
  - c. Fill in the **VSAN ID** field with an unused ID number, and the **VSAN name** field with an appropriate name. In this case study, the **VSAN ID number = 801** is assigned.
  - d. For VSAN attributes, leave everything other than the **InterOperValue** and the **AdminState** at default. VSAN attributes can be seen in the **information** pane.
  - e. Set **InterOperValue** to **Interop-1** and the **AdminState** to **Suspended**.
  - f. Statically assign VSAN membership for an interface using Fabric Manager, by choosing **Interfaces > FC Physical** from the **Physical Attributes** pane. The interface configuration appears in the **Information** pane. Select the **General** tab.
  - g. Double-click and complete the **PortVSAN** field by entering the **VSAN ID number (801)** for every port that will be used by this fabric.

For more details on VSAN settings and for configuring Interop-1 specific settings for Connectrix MDS 9000 Family switches, refer to the documentation located at <http://www.cisco.com>.
2. Set port settings on the Device Manager for the Connectrix MDS 9513:
  - a. Select the desired switch, in this case the Connectrix MDS 9513. The **Device Manager** for this switch displays.
  - b. Select the ports participating in this fabric and set the port speed to **automax 2 G**. Leave the other settings as **default**.

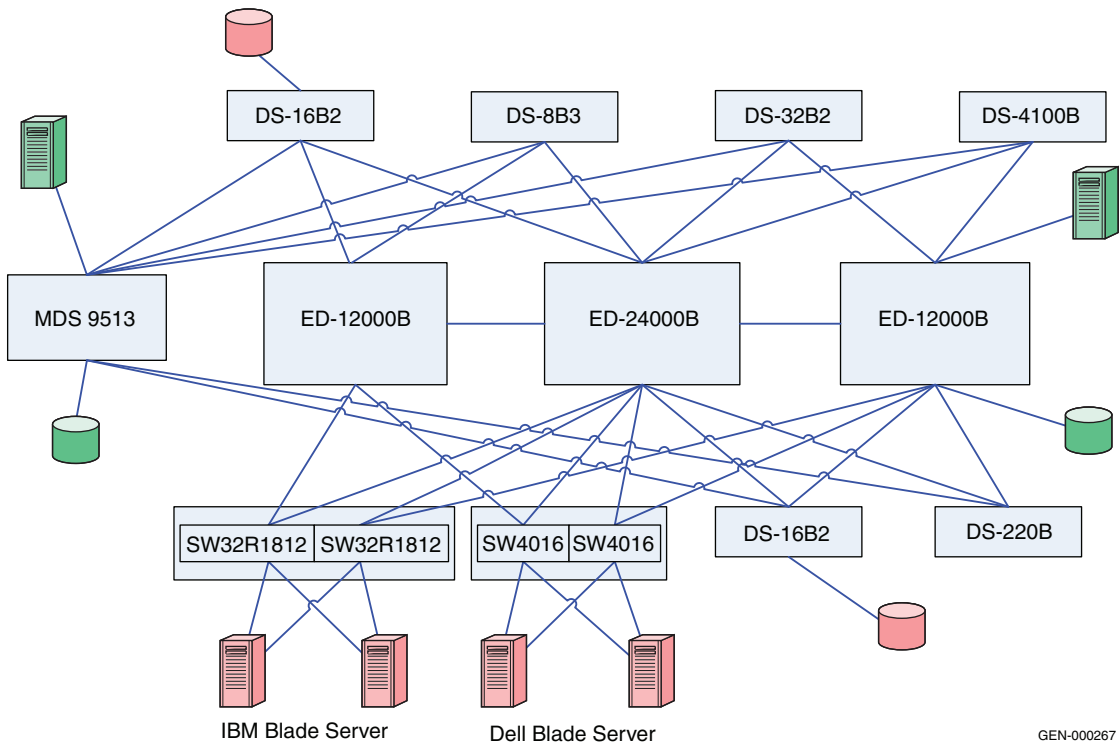
- c. Set the **Admin** option to **up**.
3. Unsuspend the VSAN.  
After configuring all the settings as stated above, go back to the **VSAN attributes** and change the **AdminState** to **Active**.
4. Using an ISL, link the Connectrix MDS 9513 to all of the Connectrix B edge switches that appear in the topology diagram.
5. Verify connectivity of the fabric by validating the same set of **“Checkpoints”** on page 169. In this case, the Cisco Fabric Manager and Cisco CLI can be used to verify the fabric topology and name server information.

### **IMPORTANT**

**At the end of this procedure, make sure the zoning information was distributed appropriately to the Connectrix MDS core. Compare the active zoneset on the Connectrix B and Connectrix MDS fabric management tools to verify that no differences exist.**

### Phase 3: Moving half of the host and storage ports from the Connectrix B core to the Connectrix MDS 9513

#### Topology



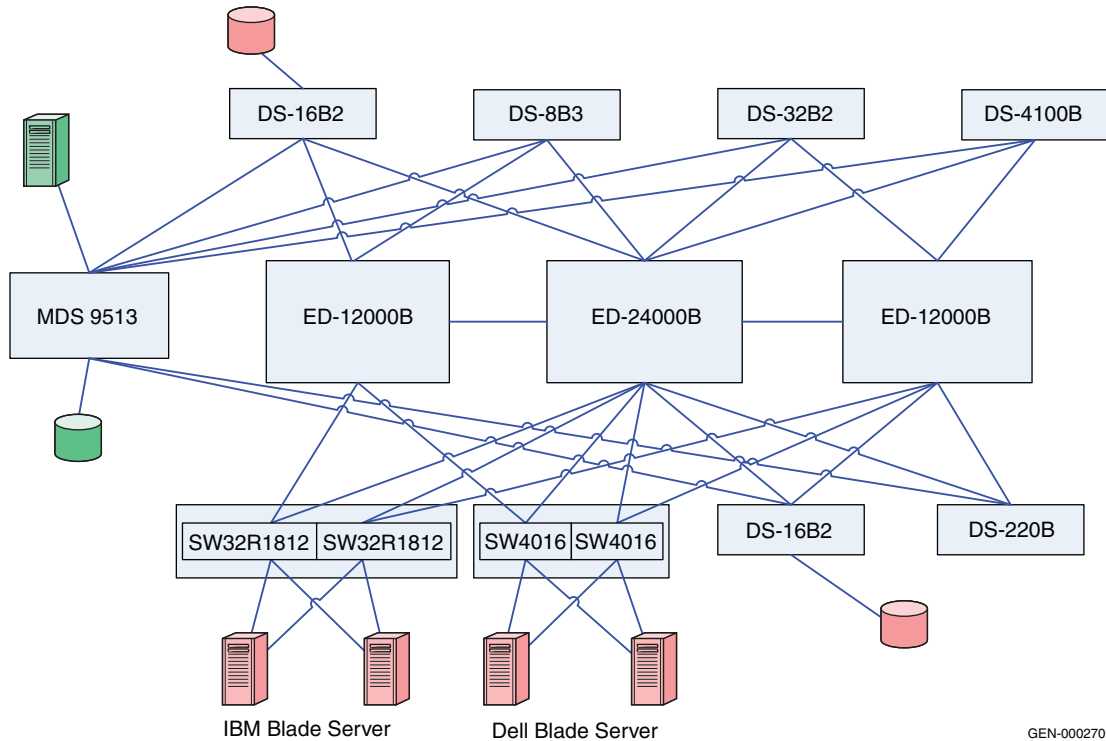
**Figure 26** Phase 3: Moving half the host and storage ports

As shown in [Figure 26](#), the same host-target pair is now connected to both core switches: the ED-24000B and the Connectrix MDS 9513. This is an intermediate phase when moving the host and storage ports to the Connectrix MDS 9513. It ensures that the zones are appropriately pushed from the Connectrix B to the Connectrix MDS switches and that the traffic between the host and storage zone is not disrupted. The host and storage ports can be successfully moved to another core switch without any downtime.

Next, you should physically pull cables from one switch and plug them into the other core switch. Additionally, you should review the [“Checkpoints” on page 169](#), to validate that this transition did not affect the connectivity and functioning of the fabric.

### Phase 4: Completely moving the host and storage ports from the Connectrix B core to the Connectrix MDS 9513

#### Topology



**Figure 27** Phase 4: Completely moving host and storage ports

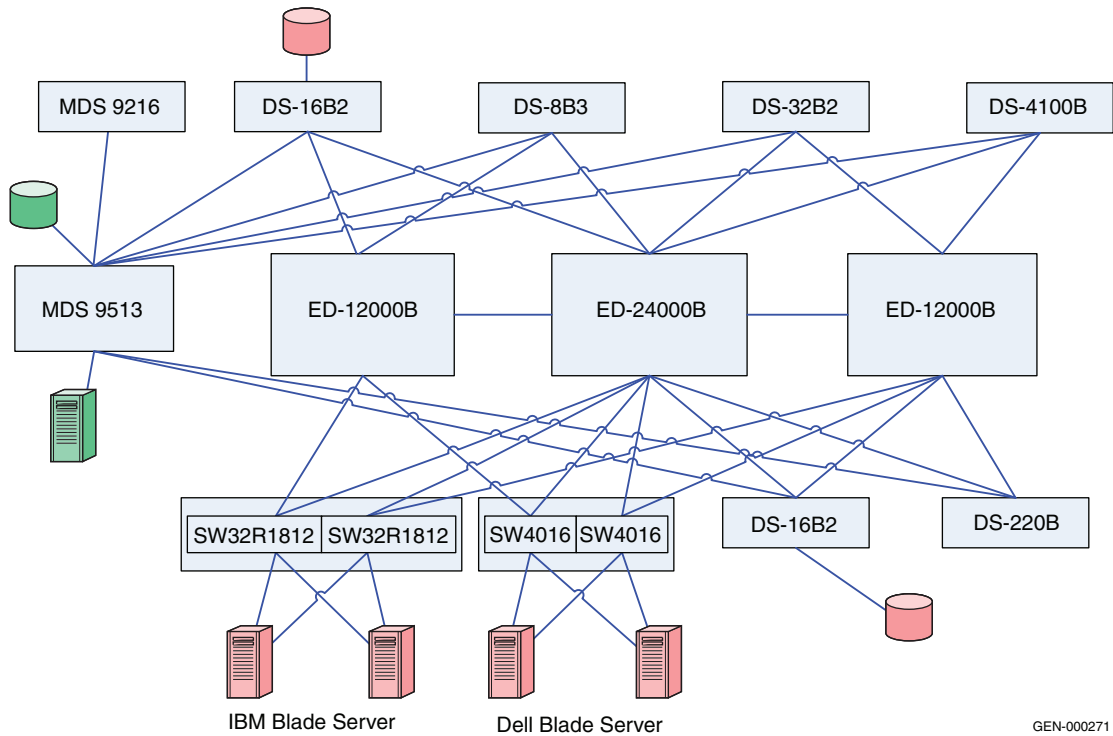
This phase is an extension of “Phase 3: Moving half of the host and storage ports from the Connectrix B core to the Connectrix MDS 9513” on page 174. After validating that a stable fabric exists, pull the remaining host-storage pairs from the ED-24000B switch and transferring them to the Connectrix MDS 9513 core.

Review the “Checkpoints” on page 169 using Connectrix MDS fabric management tools.

GEN-000270

## Phase 5: Adding a Connectrix MDS 9216 to the edge

## Topology



GEN-000271

**Figure 28** Phase 5: Adding Connectrix MDS 9216

Perform the following procedure. After connecting these switches with matching VSAN IDs, the zoning information merges. After executing these steps, this edge switch must be linked using an ISL to both the Connectrix MDS and Connectrix B core switches.

1. Create a VSAN on the Connectrix MDS switch.

A Virtual SAN can be used to create multiple logical SANs over the same physical infrastructure.

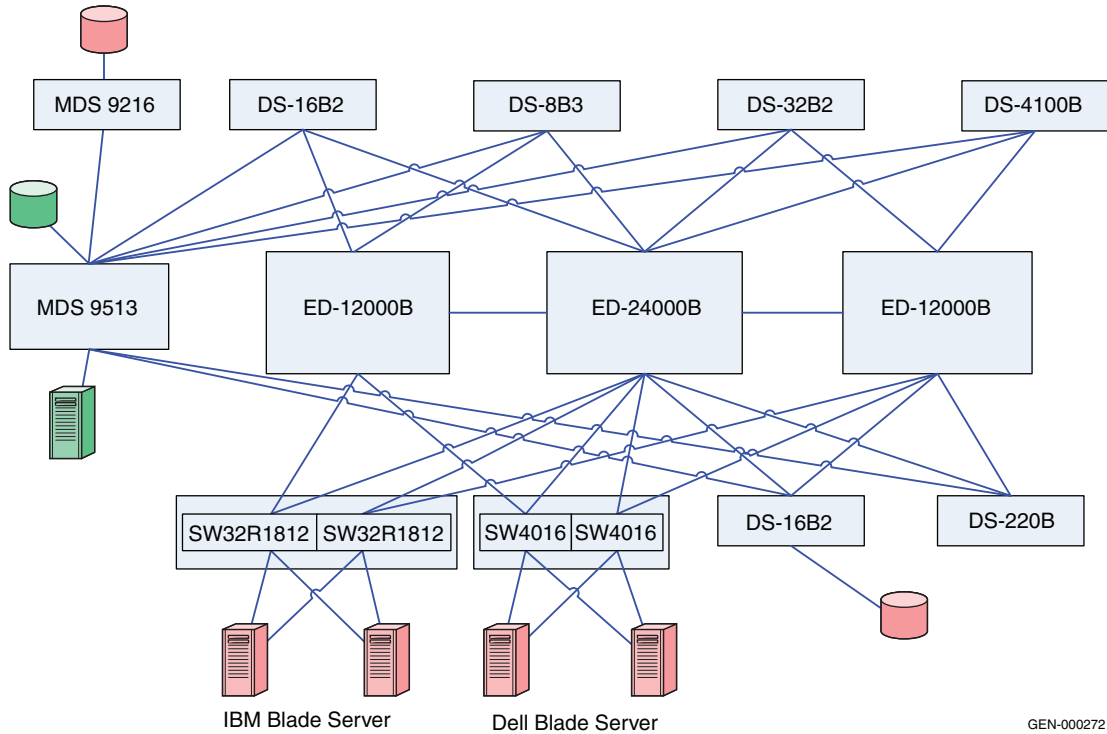
- a. Click **Create VSAN** on the **Cisco Fabric Manager**.
- b. When the **Create VSAN** window appears, select the switches that you want to include in the VSAN. In this case study, only the Connectrix MDS 9513 is included.



- c. Fill in the **VSAN ID** field with an unused ID number and the **VSAN name** field with an appropriate name. In this case study, the **VSAN ID number = 801** is assigned.
  - d. For VSAN attributes, leave everything other than the **InterOperValue** and the **AdminState** at default. VSAN attributes can be seen in the **Information** pane.
  - e. Set **InterOperValue** to **Interop-1** and the **AdminState** to **Suspended**.
  - f. Statically assign VSAN membership for an interface using Fabric Manager, by selecting **Interfaces > FC Physical** from the **Physical Attributes** pane.  
The interface configuration appears in the **Information** pane.
  - g. Select the **General** tab on this window, double-click and complete the **PortVSAN** field by entering the **VSAN ID number (801)** for every port desired for this fabric.  
  
For more details on VSAN settings and for information on configuring the Interop-1 specific settings for Connectrix MDS 9000 Family switches refer to the documentation located at <http://www.cisco.com>.
2. Set port settings on the Device Manager for the Connectrix MDS 9513:
    - a. Select the desired switch, in this case the Connectrix MDS 9513. The **Device Manager** for this switch appears.
    - b. Select the ports participating in this fabric, and then set the port speed to **automax 2 G**. Leave the other settings as **default**.
    - c. Set the **Admin** option to **up**.
  3. Unsuspend the VSAN.  
After configuring all the settings as stated above, go back to the **VSAN attributes** and change the **AdminState** to **Active**.
  4. Link, using an ISL, the Connectrix MDS 9513 to all the Connectrix B edge switches referring to the topology diagram.

After connecting these switches with matching VSAN IDs, the zoning information merges. After executing the steps above, this edge switch must be linked with an ISL to both the Connectrix MDS switches.

## Phase 6: Moving hosts and storage to a new edge Topology



GEN-000272

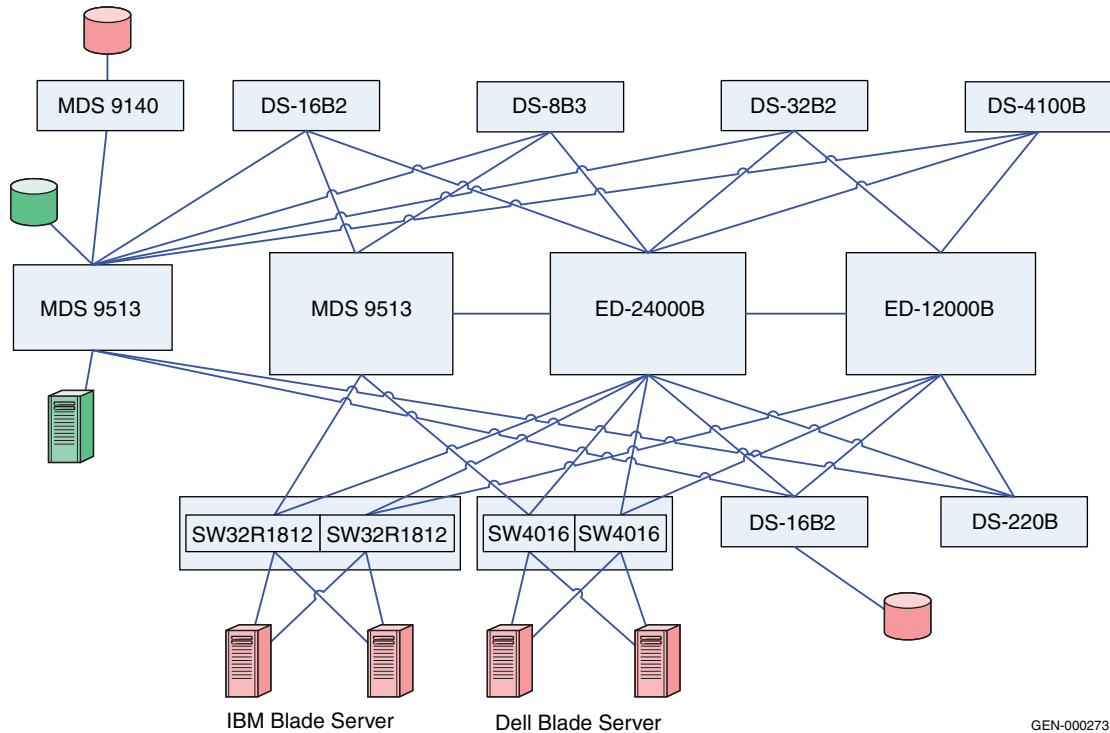
**Figure 29** Phase 6: Moving hosts and storage to a new edge

Hosts and storage ports shared by the other Connectrix B edge switches in the fabric can be completely or partially moved to the Connectrix MDS 9216 edge switch. It is evident in this transitional phase that both the switches, Connectrix B and Connectrix MDS, can co-exist in a stable fabric with the Connectrix B operating in **interopmode 1** and the Connectrix MDS operating in **Interop-1 mode**.

The settings on the switches in this phase can be used to dictate the configuration settings when setting up a Connectrix B–Connectrix MDS interop fabric. All the non-default settings previously discussed for the Connectrix MDS 9513, Connectrix MDS 9216, and the Connectrix B switches apply to setting up a heterogeneous Connectrix MDS–Connectrix B fabric from the ground up.

## Phase 7: Adding a Connectrix MDS switch to the core

## Topology



GEN-000273

**Figure 30** Phase 7: Adding Connectrix M switch to the core

As shown in [Figure 30](#), another Connectrix MDS switch, the Connectrix MDS 9513, can be added to the core with similar settings as the previous Connectrix MDS 9513, and can be linked with an ISL to the existing core Connectrix MDS 9513 director and the other edge switches in the fabric. Again, it is important to note that the VSAN must be configured with the same VSAN ID and attributes as the existing VSAN to ensure a clean fabric merge.

### Complete migration to Connectrix MDS

At the end of case study 4, a migration from a Connectrix B-only fabric to a Connectrix B-Connectrix MDS fabric with Connectrix MDS switches at the core, connected using ISLs to every edge switch in the fabric, is complete. The host and storage ports (except for the blade server host ports) can be completely moved to the Connectrix MDS

switches as specified in Phase 3 (refer to “Phase 3: Moving half of the host and storage ports from the Connectrix B core to the Connectrix MDS 9513” on page 174), Phase 4 (refer to “Phase 4: Completely moving the host and storage ports from the Connectrix B core to the Connectrix MDS 9513” on page 175) and Phase 6 (refer to “Phase 6: Moving hosts and storage to a new edge” on page 178). The Connectrix B edge switches (except for the blade server Brocade switch modules), and the Connectrix B core switches can then be pulled out from the fabric. This results in a fully operational Cisco-only fabric. This is a complete migration from one switch type (Connectrix B) to another type (Connectrix MDS).

### **IMPORTANT**

**In a Connectrix MDS-only fabric, all the Connectrix MDS switches should operate in native (default) mode. The Connectrix MDS switches are in interopmode at the end of the migration.**

You need not reboot Connectrix MDS switches to change the interopmode. The **Interop mode** attribute on the currently active VSANs must be changed from **Interop-1** to **Default**.

It is highly recommended that you backup the configuration to avoid losing all zoning information that was present in the active zoneset for Connectrix MDS switches.

### **Warnings or caveats**

Consider the following:

- ◆ Please refer to EMC Knowledgebase solution emc149735 for all interoperability issues.
- ◆ Zoning changes cannot be activated from a Brocade switch. The workaround is to use MDS switches to activate zoning changes. This caveat is specific to EMC-supported Brocade FOS v6.1.x.

---

## **Cisco Inter VSAN Routing (IVR) in a heterogeneous environment**

Cisco Inter VSAN Routing allows the Connectrix MDS series switches to consolidate multiple physical fabrics of different modes into a single physical fabric separated by VSANs while leveraging IVR to provide the necessary connectivity between the VSANs on a per-device basis. IVR provides additional benefits, including the ability to connect a core PID 0 Connectrix B switch to a core PID 1 Connectrix B switch. Normally, an outage would be required to

change the core PID of one of the switches. IVR allows the Connectrix MDS switch to act as a conduit between two different legacy switch interop modes.

### CAUTION

**RDI mode must be explicitly enabled in IVR configurations which will use interop mode 1 to connect to non-Cisco FC switches. RDI mode is automatically enabled in configurations that use interop mode 2, 3 or 4. According to documentation from Cisco:**

**When third-party switches are involved in an IVR configuration, use the `ivr virtual-fcdomain-add` command to enable RDI mode, as third-party switches do not often query the remote name server unless it is present in the domain list. Issuing this command causes all virtual domains to appear in the domain list. This command should only be issued in the VSANs (border VSANs) in which the specified switch is present. Issuing the `ivr virtual-fcdomain-add` command is required only in edge or border VSANs. It is not required in MDS native VSANs.**

### Set up IVR with interop mode

To set up IVR with interop mode VSANs:

1. Establish ISL connectivity between the Connectrix MDS switch and the third-party switch. VSANs, ISLs, and third-party switches should be configured as specified in the appropriate sections of this document.
2. Enable IVR and configure IVR topologies, zones, and zone sets on the IVR-enabled Connectrix MDS switch. IVR zoning must be managed from the Connectrix MDS only.

**Note:** In all cases, only manage IVR Zoning from the Connectrix MDS.

To configure IVR and IVR zones using the IVR Zone Wizard in Cisco Fabric Manager:

- a. Click the **IVR Zone Wizard** icon in the **Zone** toolbar.
- b. Select the VSANs that will participate in IVR in the fabric.
- c. Select the end devices that will communicate over IVR.

---

**Note:** Fabric Manager displays an error message if all the switches participating in IVR do not have unique Domain IDs. These switches must be reconfigured before configuring IVR.

---

- d. Enter the **VSAN ID** of the VSAN you want to use as the transit VSAN between the VSANs selected for the IVR zone, and then click **Next**.
- e. Optionally, configure a unique **AFID** for switches in the fabric that have non-unique VSAN IDs in the **Select AFID** dialog box.
- f. Verify the transit VSAN, or configure the transit VSAN, if Fabric Manager cannot find an appropriate transit VSAN.
- g. Set the **IVR zone** and **IVR zone set** (now or later) using the **Edit Local Full Zone** database from the **Zone** menu.

Please note:

- A zone corresponding to each active IVR zone is automatically created in each edge VSAN specified in the active IVR zone. All WWPNs in the IVR zone are members of these zones in each VSAN.
  - Sometimes zone names beginning with prefix **IVRZ** and a zone set with name **nozonest** appear in logical view. The zones with prefix **IVRZ** are IVR zones that get appended to regular active zones. The prefix **IVRZ** is appended to active IVR zones by the system. Similarly the zone set with name **nozonest** is an IVR active zone set created by the system if no active zone set is available for that VSAN, and if the **ivrZonestActivateForce** flag is enabled on the switch.
- h. Verify all steps that Fabric Manager uses to configure IVR in the fabric.
  - i. Click **Finish** if you want to enable the configured IVR topology and the associated IVR zones and IVR zone set.
  - j. The **Save Configuration** dialog box appears. The configuration of the master switch to be copied to other IVR-enabled switches may be saved. Depending on what is required, click **Continue Activation** or click **Cancel**.
  - k. Click **Finish**.

For more detailed information on setting up different types of IVR or IVR-NAT configurations please refer to the documentation located at <http://www.cisco.com>.

### Warnings or caveats

Please refer to EMC Knowledgebase solution emc149735 for all interop issues.

- ◆ Brocade switches cannot see IVR-enabled devices if NAT is enabled. IVR-1 or IVR without NAT works fine.
- ◆ If the IBM Blade Center is involved in an IVR configuration, use the **ivr virtual-fcdomain-add** command since specified switches do not query the remote name server unless it is present in the domain list. Issuing this command causes all virtual domains to appear in the domain list. This command should only be issued in the VSANs (border VSANs) in which the specified switch is present. With Connectrix MDS SAN-OS Release 1.3(4a), issue the **ivr virtual-fcdomain-add** command only in edge or border VSANs. Do not issue the command in transit VSANs. If multiple IVR-enabled switches exist in an edge or border VSAN, the following error message may be displayed in the system messages: (DEVICE\_ON\_WRONG\_NATIVE\_VSAN). This message can be ignored.
- ◆ When routing between an interop mode 1 or 4 VSAN and a non-interop mode 1 or 4 VSAN, make sure that there are no domains in the native, or legacy switch interop modes 2 or 3 VSANs that are outside of the 97 to 127 range. Domains outside of this range cannot be created in the interop mode 1 VSAN. This is not a limitation of IVR-1, but a restriction imposed upon the configuration by interop mode 1. Since IVR-2 rewrites the Domain ID and FCID of the device in the non-interop mode 1 or 4 VSAN, this is not an issue.
- ◆ If an IVR zone set is active in a VSAN running in an interop mode, regular zone changes (such as zone set activation, reactivation or deactivation for zones that do not require IVR services) should always be initiated from the IVR border switch (the IVR-enabled Connectrix MDS 9000 switch) for that VSAN. This prevents disruption to ongoing traffic between IVR initiators and targets that are allowed by both the previous active zone set and the new active zone set. Otherwise, traffic disruption occurs, and a zone change can also fail because only the border switch has the IVR configuration for that VSAN.

- ◆ Domain conflicts can occur between switches in different VSANs and thus the Domain ID of a switch needs to be reset to avoid a Domain ID clash with a switch in a different VSAN (which is accessible with IVR).
- ◆ No zone names on a third-party switch participating in an interop environment with IVR can start with **IVRZ**. All IVR zoning must be managed through the Cisco switches.

---

## Vendor-specific switch settings for interop

This section provides a more generic set of instructions or non-default settings that need to be configured on all switches belonging to a specific vendor type before they are introduced into a heterogeneous switched fabric. In addition, this section also provides the techniques, best practices, caveats, and unavailable features for a switch interoperability implementation.

### Checklist

This checklist can be used to serve as a tool for completing all steps *before* actually merging fabrics.

- ◆ Verify that each switch has a unique Domain ID. When merging fabrics, ensure that there are no duplicate Domain IDs among all switches that will be part of the merged fabric.
- ◆ Verify that all switches have been set up to work in a supported interop mode.
- ◆ Verify that the E\_D\_TOV and R\_A\_TOV are set the same on all switches that will be part of the new fabric. (By default, they should all be the same; if necessary, refer to the appropriate user manual for information on how to set up operating parameters.)

---

**Note:** Switches use different units to represent the same values; for example a value of 2000 on a Brocade switch or Cisco switch is the equivalent of 20 on a Brocade M series switch.

---

- ◆ Verify that the active zone set has been checked (with the respective switch fabric management tools) and does not contain illegal characters.



---

**Note:** There are limitations on what characters can be used for zone names. The name needs to start with a letter and then can be any combination of upper and lowercase letters as well as numbers and underscores. Everything else, including dashes, periods, or spaces should be avoided.

---

- ◆ If a switch is not operational and the zoning definition on that switch is not required, be sure to clear the zoning configuration on that switch.
- ◆ If a switch is operational and the zoning configuration on that switch is required, be sure to check that there are no duplicate active zone names. If there are duplicate zone names, rename one of the zones.
- ◆ Ensure that all switches are configured with WWN zoning.
- ◆ Ensure that all switches comply with proper zone naming.
- ◆ Back up the switch configuration by issuing the appropriate commands.

Some of the items on this checklist can be considered as recommended best practices to configure a stable interoperable environment with minimal disruption to existing data flow, if any.

### Recommended best practices

- ◆ Set the Domain IDs rather than allowing the fabric to set them.
- ◆ Set the *core* switch as a *principal* switch. This reduces Class F traffic by ensuring that it goes directly from core to edge. For example, if an edge switch is the principal switch, build fabric traffic must go through the core to get from edge to edge.
- ◆ In a vendor switch migration from either Brocade or Brocade M series, leave Brocade or Brocade M series switches in their native modes and use Cisco's legacy modes, Interop-3 and Interop-4 (only if supported) for a non-disruptive and seamless migration.

---

**Note:** At the time of this release, EMC supports Interop-4 using RPQ for migration purposes only. Refer to the *EMC Support Matrix* for the most up-to-date support.

---

- ◆ Manage all IVR-based zoning in a heterogeneous environment with Cisco switches using the Cisco Fabric Manager. No zones or zone sets created on the Brocade or Brocade M series switch, or any other vendor switch in an IVR setup can start with an IVRZ prefix.

### Configuration of non-default settings on supported vendor switches

- ◆ Refer to the specific vendor interoperability caveats before setting up an interoperable environment. Use only the correct and supported version of switch firmware and Fabric Management software.

The vendor-specific non-default settings that must be enabled on each of the supported vendor switches before introducing them into an interop environment follow:

#### Configure a Brocade switch

Telnet can be used to change the Domain ID or set the interop mode on a Brocade switch. These actions require disabling the switch; therefore, plan accordingly.

#### Principal switch setting

To configure a Brocade director switch in the core of a fabric as the *principal* switch, issue the **fabricprincipal** command on the switch CLI running FOS 4.x and higher. This setting becomes active on the next reboot, or after the build fabric (BF) event.

#### Configure the Domain ID

To set the preferred Domain ID:

1. Start a Telnet session.
2. Enter **switchdisable** to disable the switch.
3. Enter **configure**.
4. Answer **yes** to the **Fabric parameters** questions.
5. The first field, **Domain**, is the Domain ID, in the range 97 through 127, with these exceptions:
  - The lowest numbers are generally allocated to the Brocade M series switches if any, in the fabric.
  - Domain ID 104 (decimal 8) is reserved by HP.

Starting with 97, allow one number for each Brocade M series switch. Then assign the next available number to the Brocade switch, and number the remaining Brocade switches in order (skipping 104).

Accept all default values for the other parameters in this section, and answer **No** to all other sections until the switch enters the phase where it commits the new Domain ID to flash.

The Domain ID becomes the value that the switch requests when attempting to join the fabric from an offline mode.

6. Enter **switchenable**.

### Setting the interopmode

To set the interop mode:

1. Start a Telnet session
2. Enter **switchdisable**.
3. Enter **interopmode 1** to enable interopmode: **interopmode 1**.
4. Enter **reboot** or **fastboot**.

### Configure a Cisco switch

To configure an MDS series switch the Domain ID and persistent FCIDs must be set, in addition to creating an interop VSAN. If you configure the switch as a core switch, you must set the switch priority.

To create an interop VSAN:

1. Log into Cisco Fabric Manager.
2. In the left pane, expand **All VSANs**.
3. Select **VSAN Attributes**.
4. In the right pane, select **Create row**.
5. Uncheck any MDS series switch that will not participate in the new VSAN.
6. In the **VSAN** field, assign a VSAN ID, 2 through 4093.
7. If desired, give the VSAN a name.
8. Leave **Loadbalancing** at default.
9. Set **InterOperValue** to **interop-1**.

---

**Note:** Interop-2 and Interop-3 are the legacy Brocade modes and Interop-4 is the legacy Brocade M series mode, and can be set depending on the requirement.

At the time of this release, EMC supports Interop-4 using RPQ for migration purposes only. Please refer to the *EMC Support Matrix* for the most up-to-date support.

---

10. Click **Create**. The new VSAN appears in the left pane.

11. Click **Close** in the VSAN creation window.

### Configure the Domain ID and setting the switch priority

1. Log into the Cisco Fabric Manager.
2. In the left pane, expand the VSAN that has been designated for interop VSAN; for example, VSAN0002.
3. Select **Domain Manager**.
4. In the right pane, select the **Configuration** tab.
5. Set a Domain ID under the DomainID field, 97 through 127, (except 104, which is reserved for HP).
6. If configuring the switch as a core switch, the **Priority** field must be set to **1**. (If the switch will be an edge switch, the field must be left at the default setting of 128).
7. Under **Restart**, select **nondisruptive**.
8. Click **Apply Changes**.

### Unavailable features

Some of the features that are available prior to enabling the interoperability mode on FC switches are subsequently disabled. For example, when **interop** mode is enabled on Brocade switches, it disables domain/port-based zoning, Virtual Channel flow control, trunking, etc. This creates operational challenges for storage administrators who have to give up functionality in order to build heterogeneous fabrics.

Some features that are not available on the specific vendor switches when operating in Interop Fabric Mode include:

#### Brocade switches

- ◆ QuickLoop
- ◆ QuickLoop Fabric Assist
- ◆ Remote Switch
- ◆ Extended Fabrics
- ◆ Trunking
- ◆ Secure Fabric OS
- ◆ Alias Server
- ◆ Platform Service
- ◆ Virtual Channels
- ◆ FCIP

### Cisco switches

- ◆ TE\_ports (trunking expansion ports) and Port-Channels cannot be used to connect MDS to non-MDS switches. However, TE\_ports and Port Channels can still be used to connect an MDS to other MDS switches even when in Cisco Fabric mode VSANs.
- ◆ The Quality of Service feature is intended to provide nodes with high bandwidth needs and greater access to the fabric resources. Quality of Service is applied end to end (host to storage), and can be implemented only if host and storage are attached to MDS models.

### Supported interoperability modes

Table 4 on page 190 shows EMC-tested modes. Read the table as follows:

The vendor switch operating in the respective mode listed in the two left columns can interoperate with the vendor switch in the top row operating in the mode specified by the intersection of that row and column.

---

**Note:** For the most current support information, refer to the *EMC Support Matrix*.

---

Table 4 EMC-tested modes

		Connectrix B/ Brocade	Connectrix MDS/ Cisco
<b>Connectrix B/ Brocade</b>	Interopmode 0	Interopmode 0	Interop-2 (if Brocade core PID=0)  Interop-3 (if Brocade core PID=1)
	Interopmode 1 (Supported on switches running Brocade FOS versions up to v5.3.x)	Interopmode 1	Interop-1
	Interopmode 2 (supported on switches running Brocade FOS v6.0.x and higher)	Interopmode 2	N/A
	Interopmode 3 (supported on switches running Brocade FOS v6.0.x and higher)	Interopmode 3	Not supported
<b>Connectrix MDS/ Cisco</b>	Native	N/A	Native
	Interop-1	Interopmode1	Interop-1
	Interop-2	Interopmode0 (if Brocade core PID=0)	Interop-2
	Interop-3	Interopmode0 (if Brocade core PID=1)	Interop-3
	Interop-4	N/A	Interop-4

## Heterogeneous interoperability test information

### Interoperability test plan

A list of tests that EMC E-Lab performs to validate vendor switch interoperability follows. These tests were performed for most of the applicable phases of the seven phase switch migration procedure discussed earlier in “[Case studies](#)” beginning on [page 153](#).

### Generic switch features/performance tests

- ◆ I/O runs in an interop environment
- ◆ Switch reboots and power cycle
- ◆ Repeated Switch reboot

- ◆ Port type and Speed sensing (1/2/4 G performance based tests)
- ◆ NDCLA: Online code loads
- ◆ Firmware downgrade/upgrade backward compatible to two (2) firmware revisions
- ◆ Threshold alert-based tests
- ◆ Routing Table Management for switches
- ◆ Block and unblock ports

#### **Tests with a cable puller**

- ◆ Constant N\_Port failures (Cable pull test)
- ◆ Constant E\_port failures: Constant ISL breaks

#### **Jammer and analyzer tests**

- ◆ RSCN format/Domain format check
- ◆ N\_Port login process into fabric
- ◆ Sending corrupted frames to switch
- ◆ Sending out of order frames to switch
- ◆ Loss of synchronization

#### **Zoning**

- ◆ Member addition to/removal from zone
- ◆ Mixed port and WWN zoning tests
- ◆ Creating/deleting zones
- ◆ Fabric merge/conflict tests for same zone set name, different zone names, same zone members
- ◆ Fabric merge/conflict tests for different zone tests, different zone names, different zone members
- ◆ Fabric merge/conflict tests for same zone set names, same zone names, different zone members

#### **Fabric: Principal switch test**

- ◆ Principal switch selection with priority settings
- ◆ Principal switch selection with no priority settings

#### **Command line**

- ◆ Performance monitoring

- ◆ Switch configuration

#### **External Host based**

- ◆ Host power cycle
- ◆ Host Reboot

#### **Automated patch panel**

Throughout this FC-SW based interoperability section, several topologies have been discussed that can be stepped through sequentially. Normally, the initial setup and incremental changes require a number of hours to configure. In the past, this setup time proved to be a significant barrier to performing frequent interoperability testing, and customer-requested migration processes at E-Lab.

To reduce setup time and help facilitate the migration testing, E-Lab recently purchased a Brocade M series UCS 2910 Automated Patch Panel. This enables E-Lab to approach an interop deployment in stages, similar to the way deployment is performed by an end user. E-Lab can also now take point-in-time snapshots of any configuration for future reference when confirming a fix or reproducing a problem.

#### **More about the UCS 2910**

The UCS 2910 is an Optical to Electrical to Optical (O-E-O) physical layer switch that handles all serial communication protocols up to 2 GB/s. (4 GB/s interfaces are in development.) It allows multiple users to control subsets of ports. The user interface allows for simple point-and-click connections for single connections as well as configuration backup and restore functions for multiple connections.





















## Monitoring your SAN

This chapter provides the following information to monitor your SAN to limit errors that can impact performance:

- ◆ Introduction ..... 202
- ◆ Switch-based error types..... 204
- ◆ Fabric resiliency features and recommendations ..... 210
- ◆ Brocade fabric resiliency concepts ..... 218
- ◆ Configuring FS features case study ..... 225
- ◆ Summary ..... 240

---

## Introduction

As storage needs continue to grow, SANs are not only expanding in size, but are also becoming more complex and more application-intensive. Over the years, high density, high accessibility, and high scalability have been the quality aspects associated with SANs. So, do the same best practices used for designing your SAN still apply for these large and complex SANs?

The answer is *yes*. However, while reliability of data and prevention of outages have always been a priority when it comes to designing and administering a SAN, *performance* is also a critical factor to take into consideration. With the advent of newer technologies (such as installation of SSDs and deployment of FAST VP), not only are data response times expected to get shorter, but there is bound to be an increase in IOPS. This can have an impact on the available bandwidth and could become a potential source of congestion if the SAN is not well-designed and monitored. Just as proper SAN administering and design are considered best practices to manage your SAN, proper SAN monitoring should not be neglected. This chapter will detail recommended best practices to monitor your SAN to ensure higher performance.

Effective SAN monitoring not only assists in detecting any existing error conditions in a SAN, but also makes performance adjustments and aids in decisions for future capacity planning.

This chapter provides recommended best practices for monitoring your SAN. These best practices are essential for large and complex SANs that manage critical data and have optimal performance requirements.

Some of the fabric resiliency-based SAN monitoring features that will be reviewed have been added in the newer switch firmware versions and are meant to assist users in detecting high latency and congestion scenarios to prevent fabric-wide impact in a SAN. It is important to note that these features do not eliminate these error conditions.

---

**Note:** To better understand the concept of congestion, refer to the *Congestion and backpressure* section in the "FC SAN Concepts" chapter of the *Networked Storage Concepts and Protocols TechBook*, available on the on the [E-Lab Navigator, Documents > Topology Resource Center](#).

---

This chapter first discusses different switch error types, how they originate in a SAN, and how they can be detected in a Brocade or Connectrix B switch fabric. Next, current firmware features that can be used to detect the different types of switch errors, along with recommendations to prevent error impact on other switch components and device functionality in the SAN, are offered. A case study shows how a user or SAN administrator can configure these fabric resiliency features on a switch.

## Switch-based error types

Before reviewing the fabric resiliency features available on FC switches today, it is important to examine the different kind of switch-based errors, how they originate in a SAN, and how they can be monitored using the switch command line interface (CLI). This section provides different types of errors, plausible causes, and error displays.

### Error type

#### **Invalid Cyclic Redundancy Checks (CRC), or frame errors**

CRC errors indicate framing or bit errors that can occur on any link with media or transmission problems. The following are some specific areas within the FC frame that can cause a CRC error:

- ◆ Bad or missing SOF/EOF values

These errors cover frames with SOF/EOF that have invalid delimiter values. (For more information on standard delimiter values, refer to the "Ordered sets" section in the "FC SAN Concepts" chapter of the *Networked Storage Concepts and Protocols TechBook*, available on the [E-Lab Navigator](#), **Documents > Topology Resource Center**.)

- ◆ Improperly truncated frames

A frame with incomplete data or not enough bytes to fill the frame header information, such as the source/destination address, is considered an improperly truncated frame. These frames generally occur as a result of an interruption in the transmission of data, which is common while recovering from a link bounce event.

- ◆ Bit errors in the payload

These types of errors cover the more generic data corruption errors.

#### **Possible causes**

The possible causes of CRC errors in a SAN are poor physical connections or defective components. The external defective components comprise of bad transceivers on switch ports or bad cables or cable ends.

---

**Note:** For more information, refer to the "Optics" section in the "FC SAN Concepts" chapter of the *Networked Storage Concepts and Protocols TechBook*, available on the on the [E-Lab Navigator](#), **Documents > Topology Resource Center**.

---

Signal integrity errors on an internal link between the switch ASIC and SERDES can also trigger a CRC. It is important to look at all aspects of a physical connection to resolve CRC-based errors.

### Error display

Brocade/Connectrix B series switches have three fields or counters that indicate the occurrence, origin, and type of CRC errors when a **porterrshow** command is run on the Brocade CLI.

The counters that should be examined are:

- ◆ *crc err* counter– Frames with CRC errors
- ◆ *crc g\_eof* counter– Frames with CRC errors and a good end-of-frame (EOF) delimiter

Brocade tags the EOF of a frame containing a bad CRC with an 'ni' or an 'a' and forwards it without dropping the frame.

In a large fabric, comprising of ISLed switches, there is a possibility that the frame with the error will pass through multiple switches and the same erroneous frame will increment the *crc\_err* counter on all of these switches. Therefore, it is difficult to find where the CRC error originated.

The *crc g\_eof* counter was added in Brocade FOS v6.3 to indicate the origin of where the CRC error was first detected. The *crc g\_eof* counter will be incremented by '1' on the first switch that forwards the frame with the CRC error and where the EOF frame gets tagged with an 'ni' or 'a'.

- ◆ *bad eof* counter – Frames with bad end-of-frame delimiters

The *crc\_err* counter will also increment as the *bad eof* counter gets incremented.

### Error type **Link failure**

There are three primary reasons for a link failure condition: link reset, loss of sync, and loss of signal, each discussed further in this section.

If a port remains in a "link reset (LR) Receive State" for a period of time greater than the timeout period (R\_A\_TOV), an LR protocol timeout is detected, which results in a link failure.

Similarly, a link failure may also indicate that a *loss of signal* or *loss of sync* lasting longer than the R\_A\_TOV value was detected while the switch was in an online state.

### Possible causes

Possible causes include:

#### ◆ Link reset (LR)

A link reset indicates a buffer-to-buffer credit problem between two connected FC ports. The objective of a link reset is to reset the outstanding credit balance between the connected ports. When an Nx\_Port has no buffer-to-buffer credit available and has exceeded the link timeout period (E\_D\_TOV), a link timeout is detected. When a link timeout is detected, the Nx\_Port or Fx\_Port begins the Link Reset protocol. A switch or device port that cannot transmit frames due to the lack of credits received from the destination port uses the LR to reset the credit offset. The state of the port waiting to transmit during the E\_D\_TOV period is called the "LR Receive State". If the LR is not received after the E\_D\_TOV timeout has elapsed, a link failure occurs.

The generic causes of a link initialization process (often incorrectly referred to as a *link reset process*) could be cable pulls, server reboots, or port resets, but a link credit reset can be an indication of a buffer credit starvation issue that occurs as an after-effect of a frame congestion or backpressure condition.

The terms congestion and backpressure are sometimes used interchangeably, but although they are closely related, they are very different:

- Congestion occurs at the point of restriction.
- Backpressure is the effect on the environment leading up to the point of restriction.

Refer to the *Networked Storage Concepts and Protocols TechBook*, available on the [E-Lab Navigator](#), **Documents > Topology Resource Center**, for a detailed explanation and example of the congestion and backpressure concept.

#### ◆ Loss of sync

A loss of sync condition indicates synchronization failures on either bit or transmission-word boundaries. These can represent three back-to-back words with bad KChars, incorrect disparity, code violations etc. The transmitting port to which the receiver is unable to synchronize records the loss of sync.

- ◆ Loss of signal

A loss of signal condition implies that the light energy being received on a switch port, which has been transmitted by the attached device (initiator or target), has fallen below the loss of signal threshold. This is also referred to as a *link down* condition.

Similar to a link reset, a loss of sync and loss of signal are generally caused as a result of cable pulls, server reboots, or port resets. If a loss of sync occurs in the absence of these conditions, it would be worthwhile to check the link for loose physical connections and defective components, such as bad transceivers or cable ends.

### Error display

Brocade and Connectrix B series switches have three counters that indicate the occurrence of these link based errors when a **porterrshow** command is run on the Brocade CLI. The counters that should be examined are:

- ◆ *link fail* counter – Link failures (LF1 or LF2 states)
- ◆ *loss sync* counter – Loss of synchronization
- ◆ *loss signal* counter – Loss of signal

### Error type

#### Class 3 discard frames (C3TXO)

When a frame is received into the switch buffer, but the switch is unable to process the frame or route it within a pre-calculated hold time, a C3 discard condition occurs wherein these frames are dropped. In such instances, when the switch cannot keep up with traffic rates since it is holding onto frames that it is unable to process, discarding an undeliverable frame prevents this condition from completely congesting the switch for extended periods of time.

---

**Note:** The hold timer is the amount of time that a switch will allow a frame to sit in a queue without being transmitted. For more information, refer to the hold timer information in the "FC SAN Concepts" chapter of the *Networked Storage Concepts and Protocols TechBook*, available on the on the [E-Lab Navigator, Documents > Topology Resource Center](#).

---

During the recovery process, the initiator device connected to the switch and discarding the frames sends an ABTS (abort sequence) or

Status Check condition. This causes the affected frame and associated data sequence to be resent.

### Possible causes

A discard can occur for a frame with an incorrect or corrupt DID which cannot be routed, or due to the presence of devices sending frames without first using a fabric login (FLOGI).

Another potential cause that is easier to plan for, and therefore eliminate, is:

- ◆ Presence of high latency or slow drain devices

Occasionally, due to an architectural limitation, overall system load, poor volume layout, or a malfunctioning internal component, it is not possible for a destination Nx\_Port to process frames at the same rate as they are being received from the fabric. An example is a target port without enough cache. Such devices are called *high latency* or *slow drain* devices. For more information, refer to “[Latency and congestion bottleneck conditions](#)” on [page 218](#).

It is not always possible for frames to be transmitted as quickly as they are received so the fabric is bound to experience congestion. If the congestion persists long enough, backpressure will result, further causing a fabric-wide impact, leading switches to discard frames since they cannot hold onto them for more than the pre-calculated hold time.

### Error display

Brocade and Connectrix B series switch ports have one counter that indicates the occurrence of these C3 discards when a **porterrshow** command is run on the Brocade CLI. The counter that should be examined is:

*disc c3* counter – Discarded class 3 errors (switch is holding onto the frame longer than the hold time allows)

### Error type

#### Invalid Transmission Words (ITW)

ITWs indicate 8b/10b encoding errors not associated with frames but with ordered sets such as IDLEs, R\_RDYs, and other primitives.

Ordered sets are used to distinguish between data and the various types of control information. Ordered sets are four character/byte transmission words that all begin with the special character/byte K28.5. The next three transmission bytes indicate what control information is being transmitted.



An ITW is a word which does not match the definition of the different ordered sets. These errors may be in the form of code violations in one of the characters, the special character K28.5 in the wrong position in the ordered set, or an incorrect disparity.

#### **Possible causes**

The detection of invalid transmission words is an indication that the receiver is out of synchronization. The possible causes for this condition would be the same as the causes for the “loss of sync” errors which would be due to the presence of bad physical media, server reboots or link bounces.

#### **Error display**

Brocade and Connectrix B series switch ports have a counter that indicates the occurrence of these invalid transmission words or ordered sets when a `portstatshow <port number>` command is run on the Brocade CLI. The counter that should be examined is:

`er_bad_os` counter – Invalid ordered sets

- |                   |   |
|-------------------|---|
| <b>Error type</b> | <p><b>State Changes (ST)</b><br/>State changes include the link failure, mainly <i>loss of sync</i> type of errors. The possible causes for these errors are the same as discussed in <a href="#">“Loss of sync” on page 206</a>, including bad physical media, server reboots, and link bounces.</p>   |
| <b>Error type</b> | <p><b>Protocol Errors (PE)</b><br/>Protocol errors indicate FC 2 layer issues with establishing connectivity. These are issues that are likely to cause flow control errors which prevent transactions from being sent. They mainly occur at the time of FC initialization when two SAN components are trying to establish connectivity. Misreporting of the FC login parameters, including buffer-to-buffer credits, time-out values (R_A_TOV, E_D_TOV), and switch domain IDs, can prevent ISLs from forming and get recorded by the switch as protocol errors.</p> |

---

## Fabric resiliency features and recommendations

As SANs continue to scale and expand, it is difficult to keep a track of misbehaving components and to predict their effects on the rest of the SAN. Switch vendors have recently added certain features and functionality in the newer versions of the firmware that can be used to detect and notify the user about the occurrence of the errors and, in some cases, prevent their adverse impact across the SAN.

EMC recommendations have been designed to prevent data corruption, unavailability, and data loss conditions due to any unpredictable occurrence of the errors described in [“Switch-based error types” on page 204](#). The recommended topologies and SAN design tips discussed in this section have been drafted to enhance load balancing and data failover so as to prevent a SAN wide impact of bad physical media, bouncing links, or rebooting servers.

This section includes the following information:

- ◆ [“Brocade SAN resiliency features” on page 210](#)
- ◆ [“Fabric resiliency thresholds” on page 212](#)
- ◆ [“Quick reference for steps to address switch-based errors issues” on page 213](#)

---

### Brocade SAN resiliency features

Brocade introduced new features to address switch errors. These features are highly recommended for deployment in Brocade-based SANs. They include:

- ◆ [“Fabric Watch and port fencing” on page 210](#)
- ◆ [“Bottleneck Detection” on page 210](#)
- ◆ [“Edge Hold Time” on page 211](#)

#### Fabric Watch and port fencing

Fabric Watch is an optional (licensed) feature that was enhanced in Brocade FOS v6.1.0 with the addition of port fencing. The port fencing capability allows a switch to monitor specific behaviors and protect a switch by blocking a port when specified error thresholds have been reached on that port. The user can either accept the default threshold numbers or specify customized threshold values.

#### Bottleneck Detection

Bottleneck Detection was introduced in Brocade FOSv 6.3.0 with monitoring for device latency conditions and then enhanced in

Brocade FOS 6.4.0 with added support for congestion detection on both E\_Ports and F\_Ports (This enhancement was back-ported into Brocade Brocade FOS v6.3.1b and later for F\_Ports only). The Brocade FOS 6.3.1b release (and later) included enhancements in the algorithm for detecting device latency, making it more accurate. Bottleneck Detection does not require a license and is supported on both 4 and 8 Gb/s platforms.

## Edge Hold Time

Edge Hold Time (EHT) configuration is a new capability added in the Brocade FOS v6.3.1b release. Frames are dropped in switches if they have been held in the switch buffers for longer than an established Hold Time, a value calculated from several configurable fabric parameters (R\_A\_TOV, E\_D\_TOV, WAN\_TOV, or MAX\_HOPs). Unless any of these fabric parameters have been changed from their defaults, the Hold Time is calculated to be 500 ms.

By default, the hold time on all switches in a fabric tend to match since the other parameters used to calculate it have to be consistent throughout the fabric, but the EHT values can be different on switches in the same fabric.

In a core-edge topology, when congestion conditions cause frames to drop in the core of the fabric, which tend to process more traffic flows, there is bound to be more disruption.

To reduce frame drops on E\_Ports on core switches, the edge switches that host the end devices can be configured to have a shorter Hold Time compared to the core switches by using the Edge Hold Time feature (available in Brocade FOS v6.3.1b and later). This setting overwrites the edge switches' calculated Hold Time in effect. Since the Hold Time on the edge of the network is lowered, the blocked frames get discarded by the ASIC sooner than before. This reduces the likelihood of frame loss on the core of the network, effectively mitigating the impact of the misbehaving device. However, it is important to note that an I/O retry will be required for each of the dropped frames, so this solution will not completely address high latency device issues.

EMC recommends that you enable the Edge Hold Time feature on the edge switches in a core-edge topology by following these design guidelines provided by Brocade:

- ◆ The Edge Hold Time feature is recommended primarily for initiators (hosts). Extreme care must be taken if you choose to apply EHT to target ports because a target port can service a large number of initiators. A large number of frame drops on a target

port could potentially affect a very large number of running applications. Those applications may be more tolerant to poor performance or to a large number of I/O retries.

- ◆ There is no calculation for determining the best value for Edge Hold Time. Edge Hold Time can be set from 100 to 500 milliseconds. The lower the value the more frame drops you can expect. We recommend taking a value around 250 milliseconds, observing the results, and then applying the EHT value.
- ◆ Edge Hold Time is less effective when initiators and targets share the same switch because the timeout value will apply equally to both storage and host ports.
- ◆ Although an Edge Hold Time value is set for an entire switch, it gets activated on an ASIC that has one or more F\_Ports. It is thus recommended that, if possible, ISLs should be placed on a different ASIC than the servers or F\_Ports. That will prevent the E\_Ports from using the newly set EHT value and go with the default value of 500ms.

---

## Fabric resiliency thresholds

It is important to note that all the fabric resiliency features are non-disruptive except for port fencing. With port fencing, the user is notified about the error type and the affected port gets disabled, blocking all the traffic going through it. User intervention is required to re-enable the fenced port. Disabling an active port may be undesirable to some users. They may want to be alerted before the port gets fenced in order to rectify the issue with no disruption.

As an example, consider the [“Invalid Cyclic Redundancy Checks \(CRC\), or frame errors” on page 204](#). CRC errors and Invalid Words can occur on any normal links. These have also been known to occur during certain transitions such as server reboots. It is only when these errors occur more frequently that can they cause a severe impact. While most systems can tolerate infrequent CRC errors or Invalid Words, other environments can be sensitive to even infrequent instances.

Therefore, the overall quality of the fabric interconnects and fabric design are key factors.

- ◆ Cleaner interconnects and fabrics following all the recommended design best practices can have *low*, or *aggressive*, thresholds since they are less likely to introduce errors on the links.

- ◆ Less clean interconnects or fabrics that do not follow design best practices can have *high*, or *conservative*, thresholds.

In most cases, if the interconnects cannot be classified into either of these categories, low or high, the moderate or default thresholds should be used.

The low and high thresholds were created to help users plan for a disruptive event due to port fencing. The low threshold plainly notifies or alerts the user that the error low threshold value has been reached without taking any action on the port. When the high threshold is reached, the port gets disabled or fenced.

Refer to [Table 5 on page 215](#) for error types and threshold values.

## Quick reference for steps to address switch-based errors issues

This section provides the high-level steps to be executed to address issues based on a given switch-based error type. For more in-depth explanation on the specific Brocade features recommended, refer to [“Brocade fabric resiliency concepts” on page 218](#). For the detailed steps to be executed on Brocade switches, refer to [“Case study: Brocade CLI and CMDCE” on page 225](#). In this section, the error-types for which the steps have been provided include:

- ◆ [“CRC errors, ITWs, State Changes, or Protocol errors” on page 213](#)
- ◆ [“Link reset of C3 discards” on page 214](#)

The following information is also provided:

- ◆ [“Quick reference tables” on page 214](#)

### CRC errors, ITWs, State Changes, or Protocol errors

When switch ports record an unexpected number of errors such as CRC errors, Invalid Transmission words, State Changes, or Protocol Errors, it is important to detect and isolate those ports by taking the following steps:

- ◆ Set threshold values for error counts using Brocade Fabric Watch and enable *port fencing* to isolate the ports experiencing an increased number of errors.
- ◆ Take the necessary action based on the error type to eliminate the possible cause of the issue. For example, replace a bad cable, transceivers, etc., or resolve any server or storage related issues.

EMC recommends using a fiber inspection and cleaning kit provided by JDSU to address optical contamination based problems in optical networks.

(<http://www.jdsu.com/en-us/Test-and-Measurement/Products/markets/fiber-inspection/Pages/default.aspx>)

- ◆ Validate that the SAN design allows for data failover while the issue is being fixed and that no downtime has to be scheduled unless it is absolutely necessary. Some examples or ways to achieving this include:
  - Add more ISLs to create more routes for the data to flow between end devices
  - Enable ISL trunking
  - Configure multipath I/O applications, such as EMC PowerPath, on the servers
  - Enable data recovery

### Link reset of C3 discards

When switch ports experience link resets or C3 discards due to the presence of a high-latency, slow drain device, it is vital to detect this device and the effect it has on the different switch ports across the fabric to reduce its fabric-wide impact. Take the following measures:

- ◆ Enable *Edge Hold Time* on the edge switches in a core-edge topology to prevent the core, and therefore the entire, fabric from being affected by the congestion issues on the edge switches that have the slow drain devices attached.
- ◆ Enable Brocade *Bottleneck Detection* to detect the number of ports experiencing latency and congestion issues across the fabric.
- ◆ Set threshold values for error count due to C3 discards and Link Resets using Brocade Fabric Watch, with *port fencing* enabled, to isolate the ports experiencing an increased number of errors.
- ◆ Fix the misbehaving or faulty device that is the root cause of the congestion and backpressure scenario in the fabric.

### Quick reference tables

This section provides tables that can be used as a quick reference for configuring and understanding the notification, issue isolation and rectification steps for the different error types discussed in “Switch-based error types” on page 204. Tables include:

- ◆ Table 5, “Configuration reference table,” page 215
- ◆ Table 6, “Notification reference table,” page 216
- ◆ Table 7, “Issue isolation and rectification reference table,” page 217

### Configuration reference table

Table 5 provides only basic information on which feature should be enabled to detect and notify a particular error type, along with some of the standard threshold values for those error types. The actual commands or syntax used to configure these features is described in “Case study: Brocade CLI and CMDCE” on page 225. The threshold numbers provided apply to SAN topologies that follow design best practices.

**Table 5 Configuration reference table**

Error types	Port fencing	Bottleneck Detection	Edge Hold Time
CRC errors	Enable Moderate Thresholds: Low 5 High 20 Aggressive Threshold: High 2 Conservative Thresholds: Low 5 High 40	N/A	N/A
ITW	Enable Moderate Thresholds: Low 25 High 40 Aggressive Threshold: High 25 Conservative Thresholds: Low 25 High 80	N/A	N/A
State Changes	Enable Threshold: 7	N/A	N/A
Link failures	Enable (for Link Resets) Threshold: 5	N/A	N/A
C3 Discards	Enable (for C3_TX_TO) Threshold: 5	Enable Default values Threshold: 0.1 Time: 300s Quiet time: 300s	Enable Recommended value Edge hold time: 250ms

### Notification reference table

Table 6 lists the notification that the user should expect, but does not provide the exact syntax of the error messages.

**Table 6** Notification reference table

Error types	Port fencing	Bottleneck detection	Edge hold time
CRC errors	User will be notified based on the kind of event notification selected (email, etc.), but the affected port will also get disabled due to <b>CRC</b> error count exceeded.	N/A	N/A
ITW	User will be notified based on the kind of event notification selected (email, etc.), but the affected port will also get disabled due to <b>Invalid words</b> error count exceeded.	N/A	N/A
State Changes	User will be notified based on the kind of event notification selected (email, etc.), but the affected port will also get disabled due to <b>SC</b> error count exceeded.	N/A	N/A
Link failures	User will be notified based on the kind of event notification selected (email, etc.), but the affected port will also get disabled due to <b>Link reset</b> error count exceeded.	N/A	N/A
C3 Discards	User will be notified based on the kind of event notification selected (email, etc.), but the affected port will also get disabled due to <b>C3 discards</b> count exceeded.	Notification through warnings on CLI and the master log on CMDCE which specifies whether the bottleneck condition is due to latency, severe latency (stuck VC), or a credit loss.	No notification about frames being dropped within the configured Edge Hold Time.



### Issue Isolation and rectification reference table

Table 7 provides the necessary steps that the user should take if notified about a fenced port or a bottleneck condition (with or without Edge Hold Time).

**Table 7 Issue isolation and rectification reference table**

Error types	Port fencing	Bottleneck Detection	Edge Hold Time
CRC errors	Recommended rectification steps: Look into the physical connections and components, fix any apparent issues, clear the errors on the switch and re-enable the fenced or disabled port .	N/A	N/A
ITW	Recommended first steps: Look out for bad physical connections or link bounces, resolve the issue detected, clear the errors on the switch and re-enable the fenced or disabled port .	N/A	N/A
State Changes	Recommended first steps: Look out for bad physical connections or link bounces, resolve the issue detected, clear the errors on the switch and re-enable the fenced or disabled port.	N/A	N/A
Link failures	Recommended steps: Look out for port resets or server reboots, resolve the issue detected, clear the errors on the switch ,and re-enable the fenced or disabled port.	N/A	N/A
C3 Discards	Recommended steps: Identify the congested points in the fabric and the source of congestion, fix the issue, clear the errors on the switch, and re-enable the fenced or disabled port.	Identifying the source of congestion, clearing the bottleneck condition and the errors on the switch.	Clearing the bottleneck condition and the errors on the switch.

---

## Brocade fabric resiliency concepts

This section provides an example of a simplified two-hop, three-switch topology with a slow drain device attached to one of the switches, to better explain the following concepts:

- ◆ The difference between latency and congestion
- ◆ How latency on a switch can propagate through the fabric
- ◆ Different types of latencies based on severity
- ◆ How Bottleneck Detection, along with other best practices, can be used to mitigate latency effects

The following sections will help clarify these concepts:

- ◆ [“Latency and congestion bottleneck conditions” on page 218](#)
- ◆ [“Latency severities” on page 221](#)
- ◆ [“Latency detection, notification, isolation, and mitigation ” on page 222](#)

---

### Latency and congestion bottleneck conditions

A device that experiences latencies and responds more slowly than expected is known as a *slow drain*, or *latency*, device. This device does not return buffer credits (through R\_RDY primitives) to the transmitting switch to which it is directly attached fast enough to support the offered load, even though the offered load is less than the maximum physical capacity of the link connected to the device. An example is shown in [Figure 31 on page 219](#).

---

**Note:** For a detailed explanation and example of a slow drain (latency) device, refer to the *Networked Storage Concepts and Protocols TechBook*, available on the [E-Lab Navigator](#), **Documents > Topology Resource Center**.

---

Once all available credits are exhausted, the switch port connected to the device needs to hold additional outbound frames until a buffer credit is returned by the slow drain device. When the device continues to not respond in a timely fashion, the transmitting switch is forced to hold frames for longer periods of time, resulting in high buffer occupancy. This, in turn, results in the switch lowering the rate it returns buffer credits to other transmitting switches attached to it. This effect propagates through switches, and potentially multiple switches with devices attempting to send frames to devices attached

to the switch with the high-latency device, ultimately impacting the fabric. This is an example of a *latency bottleneck*, wherein a port is unable to transmit frames at the offered rate because credits are not returned fast enough from the other end (receiver). The steps shown in Figure 31 explain the sequence of effects of a latency bottleneck port.

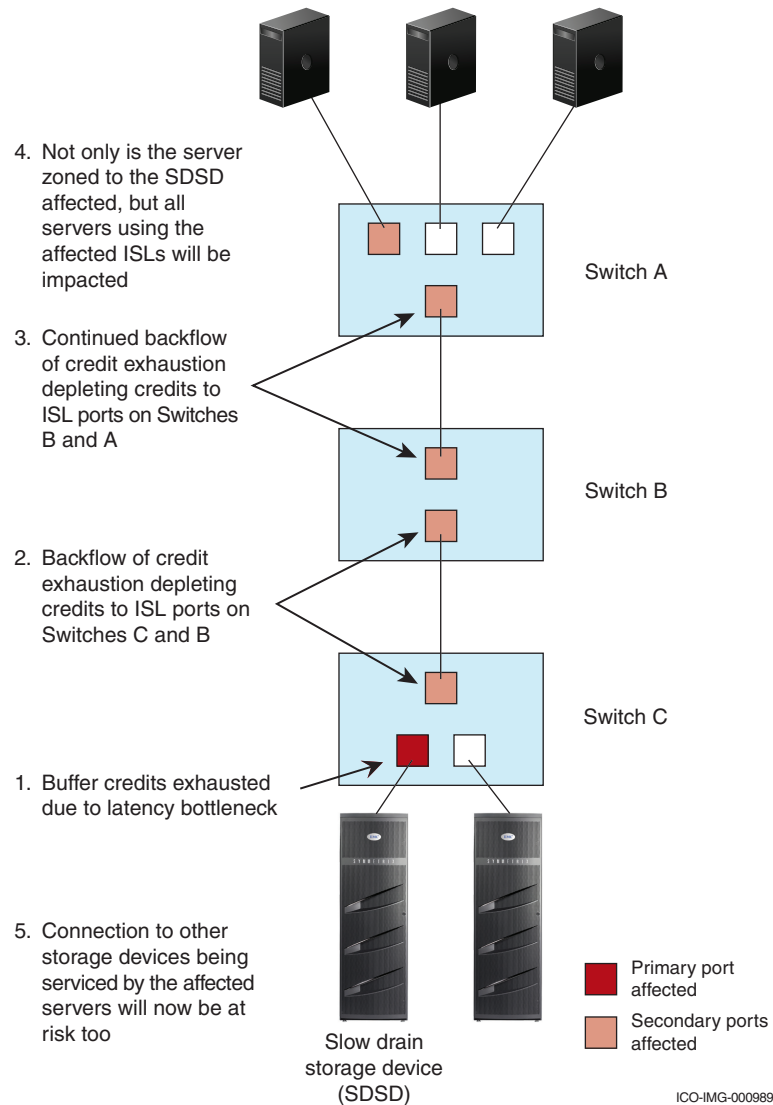


Figure 31 Fabric wide effects of a latency bottleneck condition

Another type of bottleneck is caused due to congestion wherein a port is unable to transmit frames at the offered rate because the offered rate is greater than the physical data rate of the line. This effect has been illustrated in Figure 32, where the ISL between Switch B and Switch C is a 4 Gb/s ISL expected to manage a higher data rate.

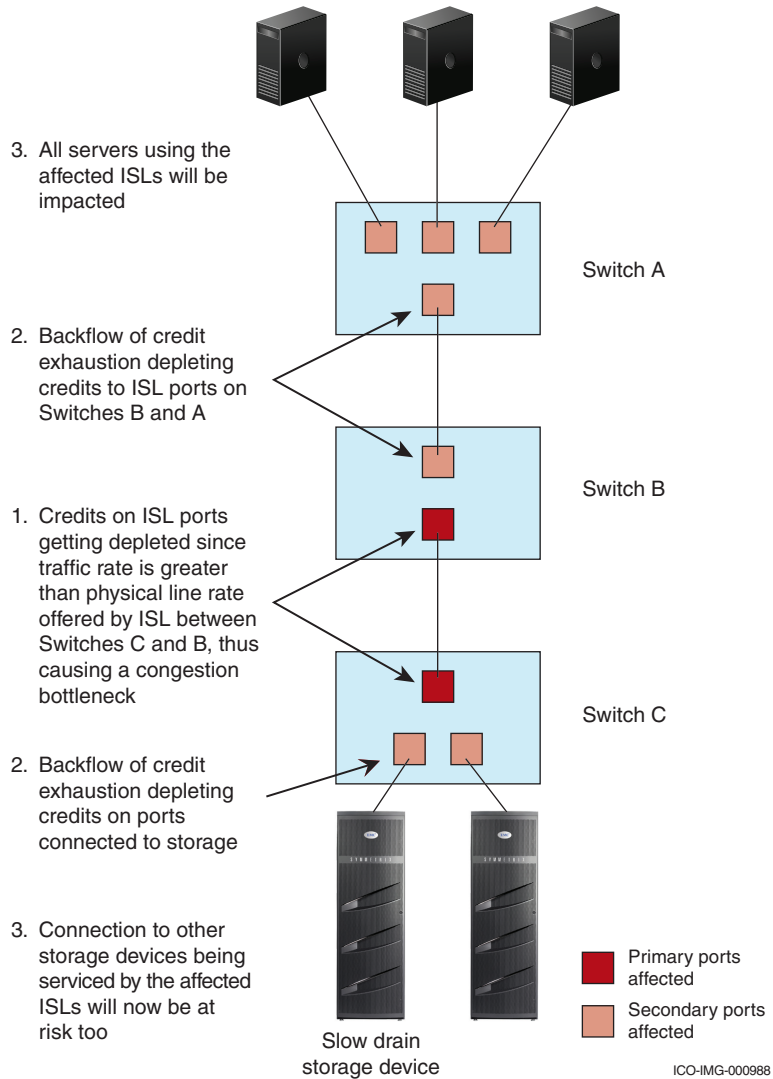


Figure 32 Fabric wide effects of a congestion bottleneck condition

---

## Latency severities

The presence of a slow drain device or an ISL with a low physical data rate, as compared to the actual offered traffic rate, can have a disruptive fabric-wide impact. However, the impact to the fabric, and other traffic flows, varies based on the severity of the latency exhibited by the device. The longer the delay caused by the device in returning credits to the switch, the more severe the problem.

The following latency severities are discussed in this section:

- ◆ [“Moderate device latencies”](#) on page 221
- ◆ [“Severe device latencies”](#) on page 221

### Moderate device latencies

Moderate device latencies are defined as those not severe enough to cause frame loss. Frame loss typically occurs above 100 ms.

If the time between successive credit returns by the device is between a few hundred microseconds to tens of milliseconds, then the device exhibits moderate latencies since this delay is not typically enough to cause frame loss. This causes a drop in performance of traffic flows using the fabric, but typically does not cause frame drops or I/O failures.

When a device exhibits moderate latency behavior, applications may see a drop in performance but not usually I/O failure. The higher the latency, the greater the chance that an end user will experience degraded performance.

### Severe device latencies

Severe device latencies result in frame loss, which triggers the host SCSI stack to detect failures and to retry I/Os. This process can take tens of seconds, or possibly as long as 30 – 60 seconds, which can cause a very noticeable application delay and potentially result in application errors.

If the time between successive credit returns by the device is in excess of 100 milliseconds, then the device is exhibiting severe latency. When a device exhibits severe latency, the switch is forced to hold frames for excessively long periods of time, possibly hundreds of milliseconds. When this time becomes greater than the established timeout threshold, the switch drops the frame (per Fibre Channel standards). This leads to the frame loss in switches or the C3 (Class 3) discards discussed in [“Switch-based error types”](#) on page 204.

Since the effect of device latencies often spreads through the fabric, frames can be dropped due to timeouts, not just on the F\_Port to

which the misbehaving device is connected but also on E\_Ports carrying traffic to the F\_Port.

Dropped frames typically cause I/O errors that result in a host retry and can result in significant decreases in application performance. The implications of this behavior are compounded and exacerbated by the fact that frame drops on the affected F\_Port (device) result not only in I/O failures to the misbehaving device, which would be expected, but also frame drops on E\_Ports may cause I/O failures for unrelated traffic flows involving other hosts, which would not typically be expected.

---

## Latency detection, notification, isolation, and mitigation

The following information is included in this section :

- ◆ [“Bottleneck Detection” on page 222](#)
- ◆ [“Fabric Watch for Timeout Notification on F\\_Ports ” on page 223](#)
- ◆ [“Port fencing to isolate a misbehaving bottlenecked port” on page 223](#)
- ◆ [“Edge Hold Time” on page 224](#)
- ◆ [“Mitigation action based on Bottleneck Detection ” on page 224](#)
- ◆ [“Proactive mitigation” on page 224](#)

### Bottleneck Detection

As discussed in [“Bottleneck Detection” on page 210](#), enabling Brocade’s Bottleneck Detection feature is a recommended best practice to detect devices that exhibit latency and congestion-based scenarios. Bottleneck Detection is a comprehensive feature that can be used to detect a wide range of device latencies from mild to severe.

Once Bottleneck Detection is enabled, the switch monitors F\_Ports for latency symptoms. Specifically, it looks for conditions in which the time delay between successive buffer credit returns from a device is higher than expected. When the condition is detected, Bottleneck Detection reports latency bottlenecks at F\_Ports based on user configurable thresholds. These reports can then be leveraged to:

- ◆ Determine the severity and duration of the latency behavior
- ◆ Determine the specific device port on which device latencies are occurring

- ◆ Determine the actual device latency in the range of 100 microseconds to hundreds of milliseconds

Detecting bottleneck scenarios, notification and alerting mechanisms, and isolation of the issue aid in preventing fabric-wide congestion. Fabric Watch, port fencing, and configuring the Edge Hold Time are other features and best practices to meet these objectives.

### **Fabric Watch for Timeout Notification on F\_Ports**

It is a recommended best practice to use Fabric Watch to detect frame timeouts, that is, frames that have been dropped because of severe latency conditions. The Fabric Watch C3TX\_TO, introduced in Brocade FOS v6.3.x for 8 Gb/s ports and available in Brocade FOS v6.3.1b/6.4.0 and later for 4 Gbp/ ports, should be used to track the number of frame timeouts. If the number of timed-out frames on an F\_Port exceeds the currently effective threshold settings, Fabric Watch notifies the user through one of the following mechanisms:

- ◆ Send an SNMP trap
- ◆ Log a RASlog message
- ◆ Send an email alert
- ◆ Log a SYSlog message

### **Port fencing to isolate a misbehaving bottlenecked port**

When a misbehaving device exhibits extremely high latencies causing frame timeouts, it is likely also causing a severe fabric impact and should be removed from the fabric. Port fencing, based on timeouts, is an optional feature that can be used to quarantine a high latency device and mitigate the impact on the fabric (8 Gb/s platform support available in Brocade FOS 6.3 and later; 4 Gb/s platform support available in Brocade FOS 6.3.1b and later). Brocade recommends enabling port fencing for transmit timeouts on F\_Ports.

Once port fencing is configured, when the number of frames dropped due to timeouts on an F\_Port reaches a user-configured threshold, the port is fenced (blocked). This disables the port, requiring user intervention to bring it back online. Once the F\_Port of the offending device is fenced, no further actions are required. The default or recommended threshold settings can safely disable the misbehaving device, preventing an impact to the fabric without causing a false trigger (fencing a port when there is not a high-latency device).

## Edge Hold Time

This applies primarily to a core edge switch topology. To reduce frame drops on E\_Ports on core switches, the edge switches can be configured to have a shorter Hold Time compared to the core switches by using the Edge Hold Time feature (available in Brocade FOS 6.3.1b and later). This setting lowers the Hold Time on the edge of the network, which reduces the probability of frame loss on the core of the network, effectively mitigating the impact of the misbehaving device.

It is a recommended best practice to enable the Edge Hold Time feature and to reduce timeouts on unrelated flows.

## Mitigation action based on Bottleneck Detection

Brocade FOS v6.4.0 and later includes an enhancement to Bottleneck Detection that allows the switch to provide some fabric-level mitigation when device latency is detected but port fencing thresholds have not yet been reached.

When latency is detected on a port, frames held in the transmit port connected to the misbehaving device are dropped for a short period of time. This allows the switch to return credits to other transmitting switches, allowing other traffic flows to move at a faster rate. This protects other flows from a severe performance drop resulting from a single misbehaving device. If a misbehaving device continues to exhibit latencies for several seconds, the port is disabled via port fencing, if port fencing has been enabled and configured.

These best practices need to be implemented once the fabric has been configured to adhere to the design or configuration best practices.

## Proactive mitigation

Fabrics can be architected to mitigate some impacts of device latency. Isolating the device flows (host/storage pair) that exhibit high latencies by either putting them in their own fabric or on their own blade/switch will contain the impact of the latencies to the fabric or blade/switch containing the high-latency device flows. Features, such as Brocade Integrated Routing (Fibre Channel Routing) and local switching, provide architectural-level solutions that limit the need for more complex monitoring and mitigation capabilities. However, using fabric design as a protection mechanism does require some knowledge of which devices are likely to exhibit latency.



## Configuring FS features case study

This section includes information for the following case study:

- ◆ [“Case study: Brocade CLI and CMDCE” on page 225](#)

### Case study: Brocade CLI and CMDCE

This case study example demonstrates how the Brocade CLI and the Connectrix Manager Data Center Edition (CMDCE) can be used to enable the resiliency features described in [“Fabric resiliency features and recommendations” on page 210](#). Snapshots of expected outcomes in the presence of a high latency device causing C3 discards or link errors are also provided.

It is important to note that:

- ◆ Bottleneck detection can be configured and monitored via CLI only (pre Brocade FOS v7.x).
- ◆ Edge Hold Time is easily configured via CLI on a switch-by-switch basis.
- ◆ Port fencing is more easily configurable using CMDCE.

This section contains the following examples:

- ◆ [“Two initiators and targets no congestion and backpressure” on page 225](#)
- ◆ [“Two Initiators and one slow drain causes congestion or backpressure” on page 226](#)

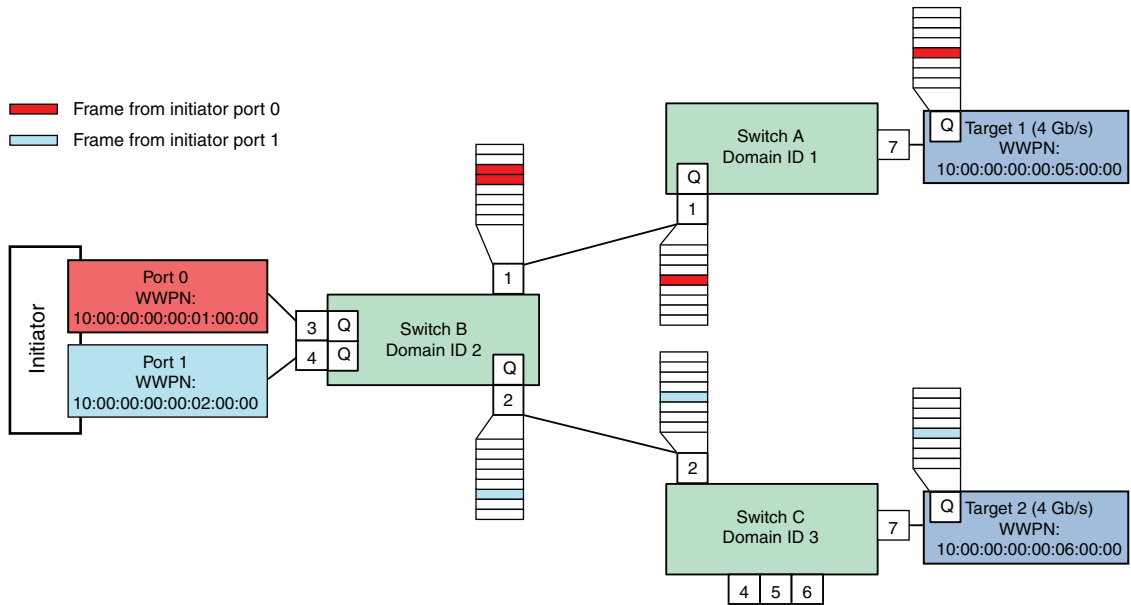
#### Two initiators and targets no congestion and backpressure

For this example, assume that a initiator with two HBA ports is zoned so that each HBA port has access to one target; for example:

- ◆ HBA port 1 is zoned to Target 1
- ◆ HBA port 2 is zoned to Target 2

[Figure 33 on page 226](#) shows an uncongested environment containing multiple initiators, each one transmitting to their own target.

**Note:** All Queues should be considered to have the same number of buffers even though they are not displayed that way in the illustration below. Although the Queues located near each port are intended to indicate shared memory type of buffers, the same type of issues can also be experienced in environments utilizing virtual output queues.



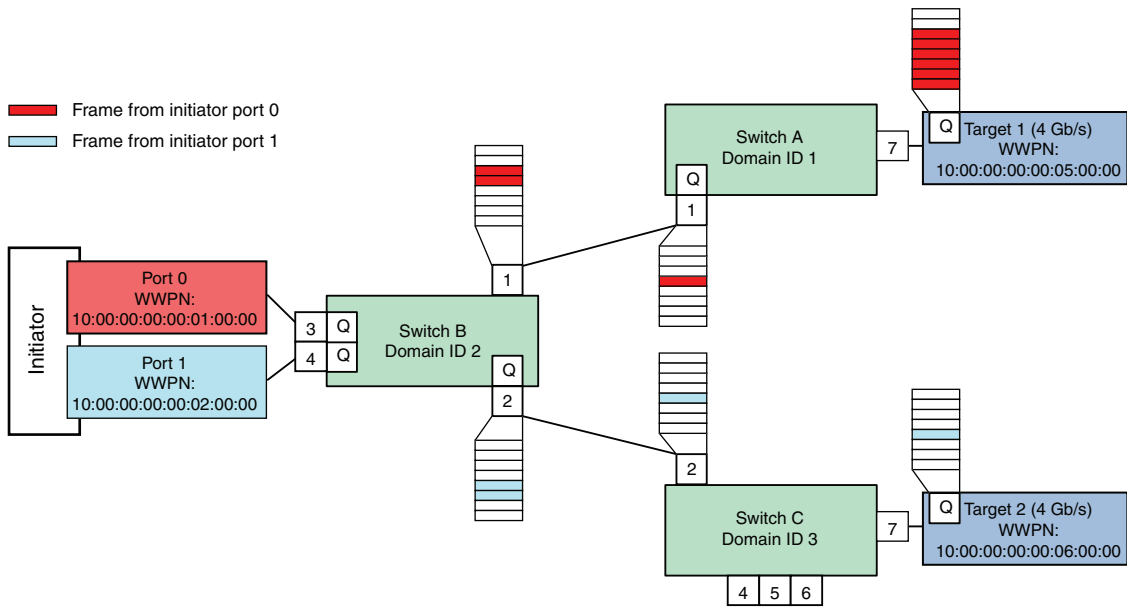
ICO-IMG-000369-TOP\_a

**Figure 33 Uncongested environment**

**Two Initiators and one slow drain causes congestion or backpressure**

Figure 34 on page 227 through Figure 36 on page 229 illustrate what can happen when a single slow drain device is present in a fabric. As shown in Figure 34, the queue on target 1 is full.

**Note:** For the sake of this example, assume that both initiators are transmitting at the same rate but that Target 1 is handling frames at a rate that is less than they are being transmitted by the initiator.




ICO-IMG-000370-TOP\_a

**Figure 34** Impact of a slow drain port

The impact a slow drain port will have on the rest of the fabric will first be felt on Switch A as shown in [Figure 34](#). Since the Queue on Target 1 is full, port 7 will have no transmit credits and will be unable to transmit any of the frames in Switch A, Port 1's queue for port 7. This, in turn, will affect the transmit queue on port 1 of switch B. However, since the Queue on Target 2 is not full, little of the transmit queue for core switch B, port 2 will be consumed.

If this condition persists, there is a possibility that the hosts application performance can deteriorate to a point where it can no longer respond to incoming frames in a sufficiently timely manner. More and more of the total number of buffers on Switch B, port 1 will be consumed with fewer and fewer buffers available in the shared memory for core switch B, which can also affect the transmit queue for port 2 on switch B in the long run, creating an undesirable fabric wide impact (see [Figure 35](#) on page 228).



ICO-IMG-000371-TOP\_a

Figure 35 Buffer Queue for port 7 continues to grow

The above situation has been illustrated in Figure 36, where edge switch A, port 1's Queue has been completely consumed by the frame destined for Target 1.

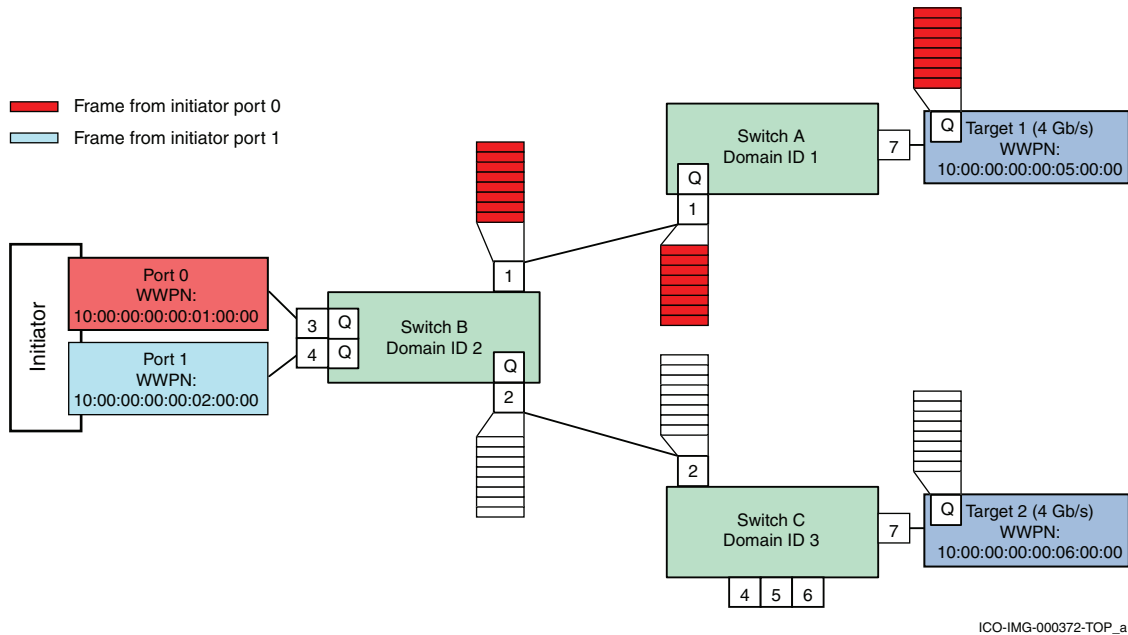


Figure 36 Buffer Queue for port 7 on Switch A, port 1

## Configuring Bottleneck Detection

Finding a slow drain, especially without the use of a protocol analyzer, is one of the most challenging performance problems to troubleshoot. A slow drain device causes a latency bottleneck and backpressure across the fabric and, in this case, leads to congestion bottlenecks. This section provides the steps to configure Bottleneck Detection.

**Note:** Bottleneck detection can be configured and monitored only using CLI (prior to Brocade FOS v7.0.x).

When Bottleneck Detection is enabled, RASlog alerts can also be enabled to be sent when the bottleneck conditions at a port exceed a specified threshold.

On the edge switches A and C with target port connections, complete the following steps:

1. Log in with *admin* level privileges.
2. Enter **bottleneckmon --enable** to enable Bottleneck Detection on the F\_Ports or ports where the target devices are attached by using the following command:

```
bottleneckmon --enable [ -alert ] [ -thresh threshold ] [ -time window ] [ -qtime
quiet_time] [slot/]portlist [[slot/]portlist]...
```

For this example, on switch A run:

```
Switch_A:admin> bottleneckmon -enable -alert 7
Switch_A:admin> bottleneckmon --enable -alert1
```

On switch C, run:

```
Switch_C:admin> bottleneckmon --enable -alert 7
```

If the alert parameter is not specified, alerts are not sent, but a history of bottleneck conditions for the port can be viewed. The thresh, time, and qtime parameters are also ignored if the alert parameter is not specified.

RASlog alerts can be enabled or disabled along with configuration of the following parameters:

- Threshold – The percentage of 1-second intervals required to generate an alert
- Time – The time window in seconds in which bottleneck conditions are monitored and compared against the threshold

---

**Note:** If the time parameter is changed, it should be set to 300 or higher.

---

- Quiet Time options – Time in seconds between bottleneck alerts

The default settings for Bottleneck Detection are the recommended settings. The default values for the threshold, time and qtime are 0.1, 300secs, and 300secs. With these default settings, alerts are logged when a port is experiencing a bottleneck condition for 10% of the time (default value) over any period of 300 seconds (default value) with a minimum of 300 seconds (default value) between alerts.

The settings are configurable in the event that a user has specific reasons for modifying them or if the design best practices have not been followed but, in most cases, the default settings should not be changed. There are several reasons they should not be changed. For example, the defaults include transient events that cause moderate congestion that are considered normal. Increasing the time or threshold may accommodate such events.

The following example shows how the threshold and other parameters can be changed:

On switch A, run the following settings for port 6 where no device is attached:

```
Switch_A:admin> bottleneckmon --enable -thresh 0.6 -time 420 6
```

- To validate that the bottleneck monitor has been enabled, the following command can be run on switch A. Only the ports on which the bottleneck monitor has been enabled will be displayed, as shown below. Even the trial setting on port 6 in the previous step has been displayed.

```
Switch_A:admin> bottleneckmon --status
```

Port	Alerts?	Threshold	Time(s)	Quiet Time(s)
1	Y	0.100	300	300
6	Y	0.600	420	300
7	Y	0.100	300	300

- The following command can then be run to monitor or to display a history of the bottleneck severity for a specific port.

The following is an example of displaying the bottleneck history for Switch A, port 7 due to the attached slow drain device in 5-second windows over a period of 30 seconds:

```
Switch_A:admin> bottleneckmon --show -interval 5 -span 30 7
=====
Mon Jun 15 18:54:35 UTC 2010
=====
From                To                Percentage of affected secs
=====
Jun 15 18:54:30    Jun 15 18:54:35    80.00%
Jun 15 18:54:25    Jun 15 18:54:30    40.00%
Jun 15 18:54:20    Jun 15 18:54:25    0.00%
Jun 15 18:54:15    Jun 15 18:54:20    0.00%
Jun 15 18:54:10    Jun 15 18:54:15    20.00%
Jun 15 18:54:05    Jun 15 18:54:10    80.00%
```

5. If the **bottleneckmon --enable --alert** option is selected, RASlog alerts will be sent when the bottleneck conditions at a port exceed a specified threshold.

If the alert parameter is not specified, alerts are not sent, but a history of bottleneck conditions for the port can be viewed. The following are kinds of alerts that the user can expect to see:

- The following is an example of a bottleneck detection alert on an F\_Port (Switch A, port 7):

```
2011/06/15-18:53:47, [AN-1003], 1, FID 128, WARNING, Switch_A, Latency bottleneck
at slot 0, port 7. 40.00 percent of last 300 seconds were affected. Avg. time
b/w transmits 677.6751 us.
```

- The following is an example of a congestion alert on an E\_Port (Switch A, port 1):

```
2011/06/15-18:55:32, [AN-1004], 2, FID 128, WARNING, Switch_A, Slot 0, port 1 is
a congestion bottleneck. 80.00 percent of last 300 seconds were affected by
this condition.
```

6. While enabling the bottleneck monitor and alerting, the bottleneck detection-based mitigation action can be enabled as follows on all F\_Ports in a switch:

```
Switch_A:admin> bottleneckmon --enable -act
```

- To enable/disable mitigation action after Bottleneck Detection has been enabled on all F\_Ports, use the following commands on Switch A:

```
Switch_A:admin> bottleneckmon --config -act OR
Switch_A:admin> bottleneckmon --config -noact
```

- To enable/disable mitigation action after enabling Bottleneck Detection for a specific port, for e.g. port 7 on switch A, use:

```
Switch_A:admin> bottleneckmon --config -act 7
Switch_A:admin> bottleneckmon --config -noact 7
```

All F\_Ports with Bottleneck Detection enabled and the -act flag set are subject to mitigation action. Ports excluded from Bottleneck Detection (using the --exclude operation) are also excluded from mitigation action.

Using Bottleneck Detection, the misbehaving ports can be identified. But, what if one does not want the affected port to impact the other ports? For example, in this case study, not only is the port attached to the slow drain device affected due to a latency condition, but it has also impacted the other port on the switch due to backpressure. This



backpressure effect can be prevented by fencing or disabling the impacted port. Implementing port fencing automatically activates this behavior.

## Enabling port fencing

Port fencing monitors ports for erratic behavior and disables a port if specified error conditions are met. It can be configured using the CLI and the Connectrix Manager Data Center Edition (CMDCE). This section discusses how both these interfaces can be used to configure port fencing.

### Using the CLI

The port fencing CLI is part of Fabric Watch and is used to enable error reporting on all ports of a specified type and configure the ports to report errors for a specific area. Supported port types include E\_Ports and F\_Ports. The specified port type can be configured to report errors for one or more areas.

This case study uses the **portfencing** command to configure port fencing for C3\_TX\_TO. As explained in [“Switch-based error types” on page 204](#), C3 discards can occur due to slow drain.

To configure port fencing for C3\_TX-TO using the CLI, complete the following steps:

1. To enable port fencing on all Fx\_Ports on the switch, for C3\_TX\_TO, run the following command on the CLI:

```
Switch_A:admin> portfencing --enable fop-port -area
C3TX_TO
```

2. Use portThconfig to customize port fencing thresholds:

For C3\_TX\_TO, threshold = 5 (and notification action selected is viewing port logs)

```
Switch_A:admin> portThConfig --set port -area C3TX_TO
-highthreshold -value 5 -trigger above -action
portlog
```

If the custom settings have to be applied, the following command should be run:

```
Switch_A:admin> portThConfig --apply port -area
C3TX_TO -action_level cust -thresh_level custom
```

3. To verify the settings in steps 1 and 2, execute the following steps:

- To display that port fencing has been enabled (with a sample output):

```
Switch_A:admin> portfencing -show
Port Type      Area          PF Status
-----
E-port        CRC           enabled
              ITW          enabled
              LR           disabled
              PE           disabled
              ST           disabled

FOP-port      CRC           disabled
              ITW          disabled
              LR           disabled
              C3TX-TO     enabled
              PE           disabled
              ST           disabled

Port          CRC           disabled
              ITW          disabled
              LR           disabled
              C3TX-TO     disabled
              PE           disabled
              ST           disabled
```

- To display or verify the port threshold configuration for all port types and areas:

```
Switch_A:admin> portthconfig --show
```

4. A fenced port can also be viewed by running the **switchshow** command. The fenced port will show up as disabled with the appropriate reason or error type.

**Using CMDCE** Step 1 through Step 4 in the “Using the CLI,” section can be executed using CMDCE by completing the following steps:

1. On the main interface, go to **Monitor > Fabric Watch > Port Fencing**, as shown in Figure 37.

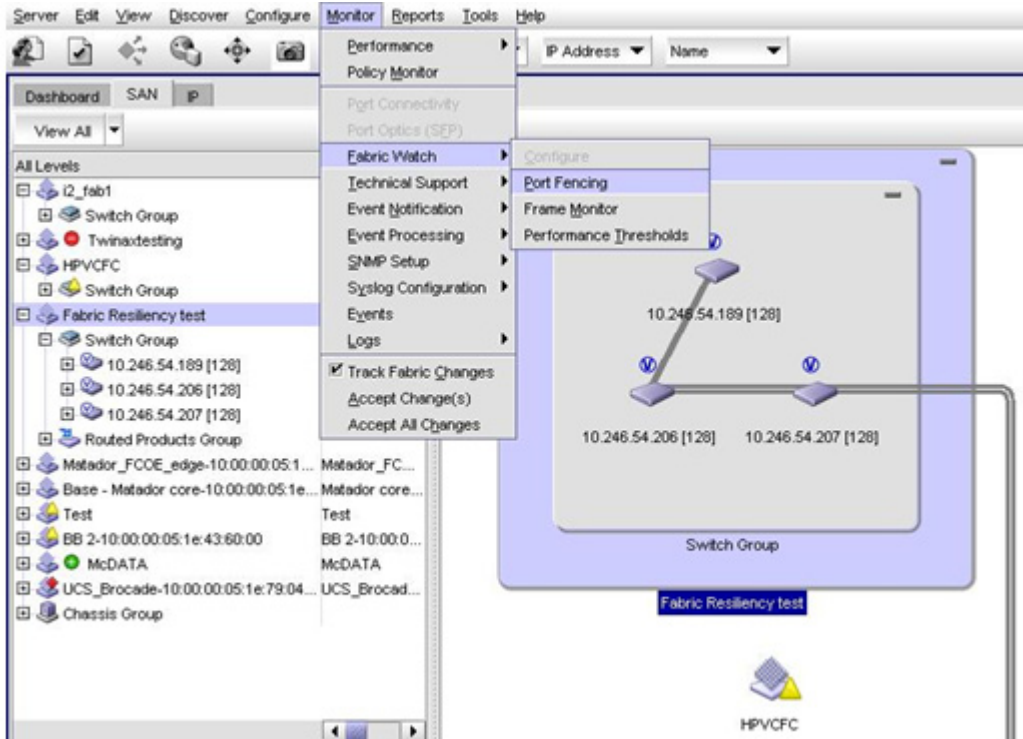


Figure 37 Port fencing dialog box

The **Port Fencing** dialog box displays, as shown in Figure 38 on page 236.

2. Select **Violation Type: C3 Discard Frames (Fabric OS only)** and click **Add** to create a port fencing threshold that can be applied to all E\_Ports or F\_Ports on the desired switch, as shown in Figure 38.

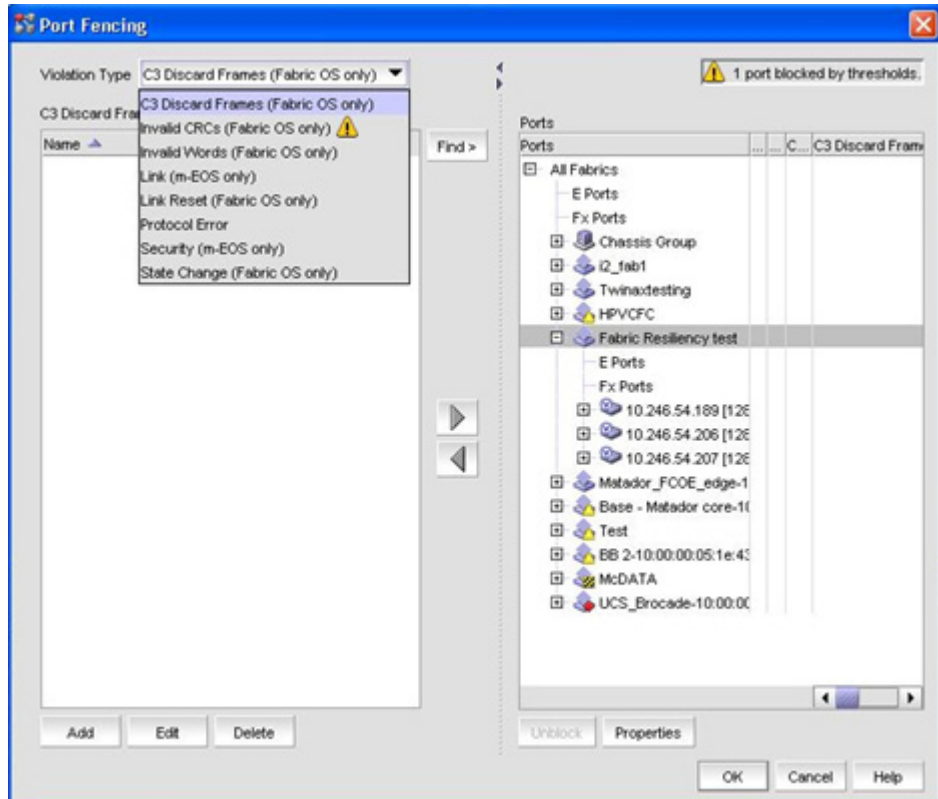


Figure 38 Port Fencing dialog box

3. Create a customized threshold for the C3 discard errors as shown Figure 39.



Figure 39 Create a customized threshold

4. Apply them to the F\_Ports on Switch\_A as shown in Figure 40.

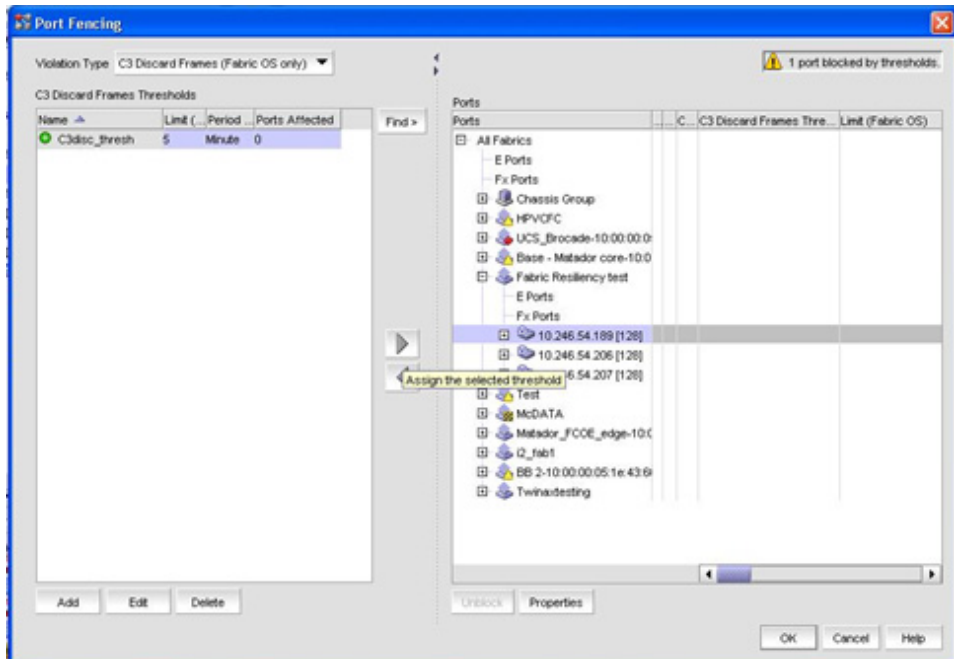


Figure 40 Apply customized threshold

The total number of fenced ports are displayed on the top right corner of the port fencing dialog as shown in [Figure 40 on page 237](#). Not only will the master log specify the specific port that has been fenced, but a notification mechanism selected such as raslogs, portlogs, email, etc., can be used to alert the user about a fenced port.

5. If the port has been fenced due to an error, based on the error type, the user needs to intervene, fix the condition causing the error, clear the errors by running **statsclear**, and re-enable the disabled or fenced port.

---

## Configuring Edge Hold Time

The slow drain condition described in this case study cannot persist indefinitely since there is a Hold Timer on the switches. The Hold Timer is the amount of time that a switch will allow a frame to sit in a queue without being transmitted. When the timer expires, the frame is discarded.

The HoldTimer for the Connectrix B Series switches is 500ms. However, this does not provide a solution to the problem with the slow drain since average frame latency time through a switch is between 600 ns to 20 usec. A Hold Timer of 500ms is 25,000 times greater than the longest average latency. Therefore, this Hold Timer will have limited value for this particular scenario.

As described in [“Edge Hold Time” on page 211](#), the Hold Time on an edge switch that has the high latency device attached can be reduced by 5 times to 100ms to relieve the backpressure effect on the core serving other edge switches, preventing a fabric-wide impact.

A user can configure the Edge Hold Time on edge switch A using the following command. The switch does not need to be disabled to modify the Hold Time. Use the **Configure edge hold time** option to turn this feature on or off.

```
Switch_A:admin> configure
Not all options will be available on an enabled switch.
To disable the switch, use the "switchDisable" command.
Configure...

Fabric parameters (yes, y, no, n): [no] yes
Configure edge hold time (yes, y, no, n): [yes]
Edge hold time: (100..500) [100]
```

The Edge Hold Time value is persistently stored in the configuration file. All configuration file operations, such as **configupload** and **configdownload**, are supported for this feature.

Note: This setting is available only in Brocade FOS v6.3.1b and later.

---

## Summary

SAN performance monitoring tools are gaining more attention with some of the newer switch firmware releases. Just as proper SAN administering and design are considered best practices to manage your SAN, proper SAN monitoring best practices are also vital.

Effective SAN monitoring not only assists in detecting any existing error conditions in a SAN, but also makes performance adjustments and aids in decisions for future capacity planning. This chapter provided recommended best practices for monitoring your SAN. These best practices are essential for large and complex SANs that manage critical data and have optimal performance requirements.



## Brocade Virtual Fabrics Case Study

This chapter provides a case study for Brocade Virtual Fabrics.

- ◆ Brocade Virtual Fabrics case study overview ..... 242
- ◆ Objectives of Virtual Fabrics architecture..... 243
- ◆ Logical Switch capability ..... 244
- ◆ Virtual Fabrics and inter-switch links ..... 247
- ◆ How to configure Brocade Virtual Fabrics case study..... 250
- ◆ Brocade Virtual Fabrics versus traditional Cisco Virtual SANs.. 268

## Brocade Virtual Fabrics case study overview

Brocade Virtual Fabrics (VFs) allows a physical Brocade or Connectrix B series switch to be partitioned into multiple Logical Switches, which can in turn be interconnected to form Logical Fabrics. Each Logical Switch acts as an independent fabric component in terms of protocol and management. Each Logical Switch has its own fabric services (name server, zoning, etc.), configuration (port, switch, fabric, etc.), and fabric characteristics (operating mode, addressing, etc.).

As the number and size of SANs continue to grow, the complexity of managing these SANs has become a main concern. Large MetaSANs pose problems including fragmented SAN islands, lack of isolation in meta SANs, and limited scalability. For more information and examples showing MetaSANS, refer to "SAN routing concepts" *Networked Storage Concepts and Protocols TechBook*, available through on the [E-Lab Navigator](#), **Documents > Topology Resource Center**. The Virtual Fabrics feature introduced with Brocade Fabric OS v6.2.x provides solutions to these problems.

The Virtual Fabrics feature is supported on the Connectrix ED-DCX-B, ED-DCX-4S-B, DS-5300B, and DS-5100B switch platforms beginning with Brocade Fabric OS v6.2.0e.

The Virtual Fabric concepts discussed in this section will become clearer as we examine a case study for creating and configuring logical or virtual fabrics in "How to configure Brocade Virtual Fabrics case study" on page 250.

---

**Note:** If you are used to working with Cisco VSANs, please review "Brocade Virtual Fabrics versus traditional Cisco Virtual SANs" on page 268 before proceeding. This section outlines specific details that the user must be aware of before configuring Brocade VFs.

---

The following information is provided in this section:

- ◆ "Objectives of Virtual Fabrics architecture" on page 243
- ◆ "Logical Switch capability" on page 244
- ◆ "Virtual Fabrics and inter-switch links" on page 247
- ◆ "How to configure Brocade Virtual Fabrics case study" on page 250
- ◆ "Brocade Virtual Fabrics versus traditional Cisco Virtual SANs" on page 268

## Objectives of Virtual Fabrics architecture

The objective of the Virtual Fabrics architecture is to provide the following:

- ◆ Virtualize hardware boundaries

Traditionally, the SAN design/management is performed at the granularity of a physical switch. The Virtual Fabrics feature allows SAN design to be made at the granularity of a port.

- ◆ Isolation

The Virtual Fabrics feature provides isolation across different SANs that are part of the same metaSAN.

The following levels of isolation are provided:

- Management isolation: The different SANs are part of different management domains and are isolated from each other.
- Protocol isolation: Protocol events do not cross SAN boundaries.
- Data path isolation: The physical links used by the different SANs can be isolated such that no two SANs share the same physical link.

- ◆ Scalability

The Virtual Fabrics architecture segments the SANs in a metaSAN such that each of the SANs can scale independently. The scalability of the meta-SAN is decoupled from the scalability of the individual SANs.

## Logical Switch capability

Logical Switch capability allows partitioning of a physical chassis, such as the ED-DCX-B, ED-DCX-4S-B, DS-5300B, or DS-5100B, into multiple Logical Switches. Each such Logical Switch forms an independent L2 switch, as shown in Figure 41. Such independent Logical Switches can be connected to form a Logical or Virtual Fabric.




Figure 41 Logical Switches

In Figure 41, (A) shows a VF-capable default L2 switch; (B) shows the same L2 switch in (A) with VF enabled. In (B), the partitions configured on the L2 switch are the Default Logical Switch (DEF LS), Logical Switches (LS1, LS2) and the Base Switch (BS).

The following switches are discussed in this section:

- ◆ “Default Logical Switch” on page 244
- ◆ “Logical Switch” on page 245
- ◆ “Base Switch” on page 246

### Default Logical Switch

When Virtual Fabrics is enabled on a switch that supports the VF feature, the switch transforms into a Default Logical Switch, which is like a regular L2 switch. All the ports in the switch now belong to the *Default Logical Switch*.

In Figure 41, L2 is a physical switch (with ports P1 to P6) on which VF can be enabled. Once VF is enabled, a Default Logical Switch (DEF LS) will be created with ports P1 to P6 showing up as members of the DEF LS.

## Logical Switch

A switch/chassis can be divided into multiple Logical Switches (LS). Each Logical Switch acts as a L2 switch.

The user must configure each port with a Fabric ID (FID) that uniquely maps a port to a Logical Switch. Any given port can only be in *one* Logical Switch. All ports with the same FID are a part of the same Logical Switch. Thus, the FID is the attribute that distinguishes one Logical Switch from the other.

As the user initially allocates ports to new Logical Switches, those ports are removed from the Default Logical Switch and assigned to the specific Logical Switch that is being created.

**Note:** Some types of ports and blades cannot be removed from the Default Logical Switch. This information is provided in the *EMC Connectrix B Series Administrator's Guide* for FOS v6.2.x.

A port is automatically disabled when being assigned to a Logical Switch. User can also move ports from one Logical Switch to another. Each Logical Switch can have as many ports as available in the chassis. For FOS v6.2, the limits for the number of Logical Switches per product are shown in [Table 8](#).

**Table 8** Number of Logical Switches supported per product

Product	Maximum number of Logical Switches per chassis (includes default switch)
ED-DCX-8510-8B	8
ED-DCX-8510-4B	8
DS-6510B	4
DS-6520B	4

In [Figure 41 on page 244](#), two Logical Switches, LS1 and LS2, each with unique FIDs, were created by the user. Note that the DEF LS is also assigned a default Fabric ID of 128. Ports P2 and P3 were added to LS1 from the DEF LS, while ports P4 and P5 were added to LS2 from the DEF LS.

Each Logical Switch can be configured to have its own preferred Domain ID and other fabric parameters. During Virtual Fabric (or Logical Fabric) formation, any conflict will be resolved as it would be for a regular L2 switch fabric.

**Base Switch**

A Base Switch, also known as Base Logical Switch, provides a common address space for communication between different logical fabrics. A Base Switch can be created with the same CLI commands used to create a Logical Switch and the user can choose if the Logical Switch is a Base Switch or not. Just like the Logical Switch, a Base Switch can be configured like an L2 switch, with the preferred Domain ID.

In [Figure 41 on page 244](#), the user created a Base Switch (BS) and allocated port P6 to the base switch. Base Switch ports on different chassis can be connected together to form a Base Fabric. By default, E\_Port links between Base Switches would be a shared ISL (XISL).

Once a Base Fabric is formed (out of Base Switches) and has become stable, logical Fabric formation will begin. If Base Switches have different FIDs, the base switches with conflicting IDs that are inconsistent with the other base switches' ID will be disabled, or the link between them will be disabled. This applies to links between all Logical Switches. The FIDs of the two Logical Switches or Base switches connecting to each other must match.

Base Switches are also used for FCR support/connectivity. All EX\_Ports on a switch must be a part of the Base Switch. Base Switches do not support direct device connectivity; therefore, a Base Switch must have only E\_Ports or EX\_Ports.

## Virtual Fabrics and inter-switch links

This section briefly discusses the following:

- ◆ “Dedicated ISLs (DISLs),” next
- ◆ “Conventional ISLs” on page 247
- ◆ “Extended ISL (XISLs) and Logical ISLs (LISLs)” on page 248

### Dedicated ISLs (DISLs)

An interswitch link (ISL) connected between two Logical Switches (LS) is called a *DISL* (Dedicated ISL). The user does not need to explicitly configure the port to be a DISL. A DISL can carry L2 frames associated with the local FID only. If a DISL is connected between two Logical Switches with different FIDs, the DISL will be segmented. [Figure 42](#) shows DISL connections between Logical Switches L1, LS2, and Default LS (DEF LS).




Figure 42 DISL connections between Logical Switches

### Conventional ISLs

As shown in [Figure 43 on page 248](#), if the DISL from P3 of a Logical Switch (LS 1) is connected to P4 of (non-VF-capable) L2 Switch, it would now be termed as a regular ISL. The user does not have to change configuration of the port to convert from a DISL to an L2 ISL, or vice versa. The terminologies are introduced to differentiate the connection points at the ends.

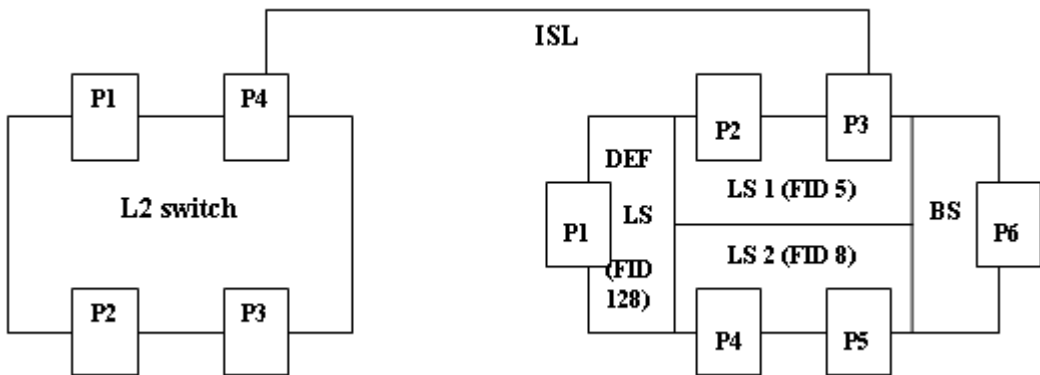


Figure 43 ISL connection between Logical Switches and non-VF-capable switch

**Extended ISL (XISLs) and Logical ISLs (LISLs)**

An ISL connecting two Base Switches is called an *XISL*. In [Figure 44 on page 249](#), there is no *DISL* connecting the two Logical Switches of the two different chassis. The Base Switches (BS) on each chassis form a Base Fabric through the *XISL* connection. Logical Switches LS 1 and LS 2 are configured to share the *XISL* to communicate. Once the Base Fabric is stable and the Logical Switch configuration exchange is complete, a logical link (*LISL*) will be formed between the two Logical Switches, as shown by the dotted line in [Figure 44](#). This *LISL* is part of the physical *XISL* connecting the Base Switches.

**Note:** Only the *DISL* and *XISL* are actual physical connections. The *LISL* does not represent a physical ISL.



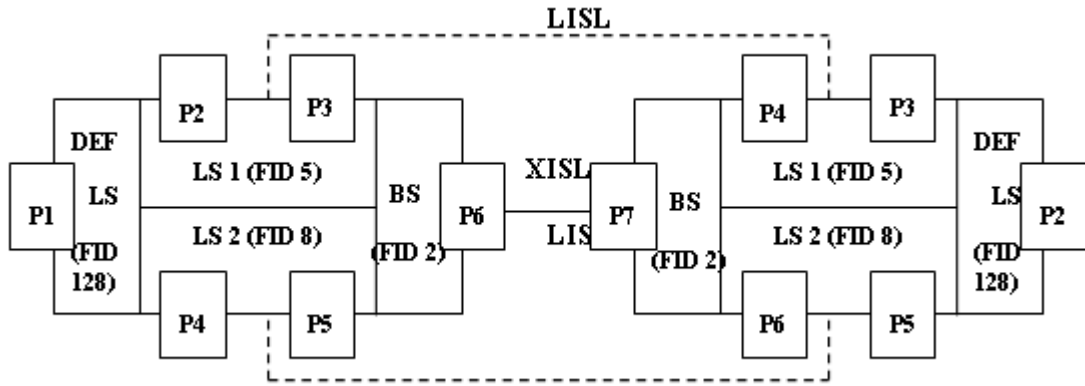


Figure 44 XISL connection between Base Switches (BS)

As shown in [Figure 44](#), LISL connections are automatically formed between Logical Switches LS 1 and LS2.

## How to configure Brocade Virtual Fabrics case study

This case study converts topology (A), as shown in Figure 45 on page 250 to topology (B), as shown in Figure 46 on page 251.

### General layout

The example for this case study on Virtual Fabrics comprises of two ED-DCX-4S-B Director class switches (IPs: 172.23.199.4, 172.23.199.5) at the core, with an ED-DCX-B switch (IP: 172.23.199.6), a DS-5000B (IP:172.23.199.7), a DS-5300B (IP:172.23.199.8) and a DS-4900B (IP: 172.23.199.9) switch at the edge, as shown in Figure 45.




Figure 45 Topology A example

This case study will convert the topology shown in Figure 45 to enable and configure VFs, as shown in Figure 46 on page 251.

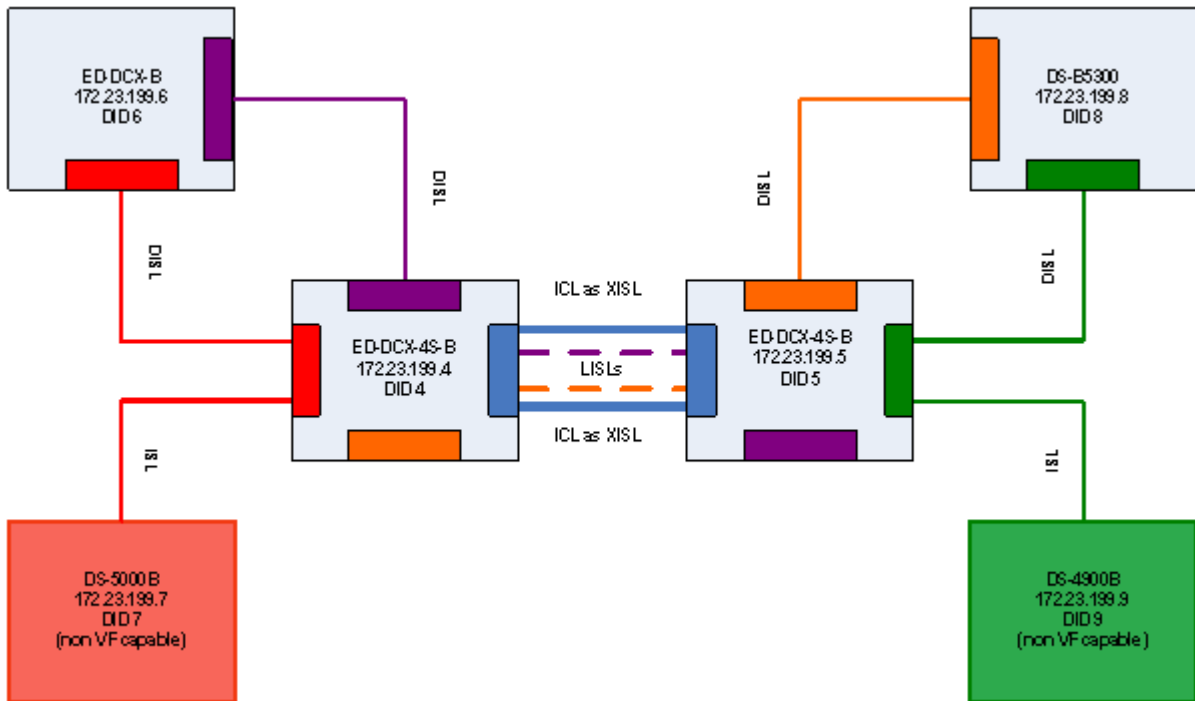


Figure 46 Topology B example

Note the following:

- ◆ The Virtual Fabrics (VF) feature is only supported on the ED-DCX-4S-B, ED-DCX-B, and DS-5300B switches in this configuration.
- ◆ The colored blocks within the ED-DCX-4S-B, ED-DCX-B, and DS-5300B in Figure 46 represent the Logical Switches on those respective switches. We can ignore the color coding in the figure for now. This will be clearer as you read through the case study and refer to Figure 47 on page 261.

**Assumptions specific to this case study**

For this case study, we will assume the following:

- ◆ All these switches have been powered up, have been assigned Domain IDs as specified in [Figure 46 on page 251](#) and listed below:

Switch	IP	Domain ID (DID)
ED-DCX-4S-B	172.23.199.4	4
ED-DCX-4S-B	172.23.199.5	5
ED-DCX-B	172.23.199.6	6
DS-5000B	172.23.199.7	7
DS-5300B	172.23.199.8	8
DS-4900B	172.23.199.9	9

- ◆ All switches are isolated to begin with (i.e., they are not connected to any other switches).
- ◆ The VF-capable switches (ED-DCX-4S-B, ED-DCX-B, DS-5300B) are running EMC-supported FOS v6.2.x and up and the other switches are running compatible EMC-supported FOS versions.

---

**Note:** Refer to the [EMC Support Matrix](#) (ESM) for the most up-to-date information on EMC-supported switches.

---

- ◆ All switches are operating in native Brocade mode (interopmode 0)
- ◆ No zoning configuration is present on any of the switches

---

**Note:** When VF is enabled on a switch, if there is an existing zoning configuration it now becomes part of the Default Logical Switch that was automatically created once VF was enabled. When new Logical Switches are partitioned from the Default Logical Switch, the zoning configuration is cleared and is no longer part of the Logical Switch; it remains only with the Default Logical Switch.

---

### Objectives of this case study

This case study can be used as a reference to meet the following objectives:

- ◆ To enable VFs on the VF-capable switches
- ◆ To partition a physical switch into multiple Logical Switches
- ◆ To create base switches that can be used to carry the traffic coming from multiple Logical Switches
- ◆ To differentiate between DISLs, XISLs, and LISLs
- ◆ To configure four Virtual fabrics within an existing fabric topology
- ◆ To create zoning configurations within the individual VFs

### Configuring Virtual Fabrics

To configure the Virtual Fabric, complete the following steps:

1. Access the individual switches with IP-based management and log in using an account assigned to the admin role.
2. On the ED-DCX-4S-B, ED-DCX-B, and DS-5300B switches, enable the Virtual Fabrics feature using CLI, if it is not already enabled. The steps that need to be executed on each of these switches are as follows:

---

**Note:** VF is disabled by default on a new Connectrix B series switch and on switches that are upgraded to Fabric OS 6.2.0 or later. Before using the VF features, such as Logical Switch and logical fabric, VFs must be enabled. When enabling VFs, the CPs are rebooted and all existing EX\_Ports are disabled after the reboot. All Admin Domains must be deleted.

---

- a. Enter the following command at the prompt to check whether Virtual Fabrics is enabled:

```
fosconfig --show
```

Expected output for this command is shown in Step c.

- b. Enter the following command to enable Virtual Fabrics. You will receive a Warning.

```
fosconfig --enable vf
```

**Warning!**

**This is a disruptive operation that requires a reboot to take effect.**

**All EX ports will be disabled upon reboot.  
Would you like to continue [Y/N] y**

- c. Once the switch is up again, the following command can be run to verify that the Virtual Fabrics is enabled:

```
switch:admin> fosconfig --show
FC Routing service: disabled
iSCSI service: Service not supported on this Platform
iSNS client service: Service not supported on this Platform
Virtual Fabric: enabled
```

In the case of our example, verify that all the VF-capable switches display “Virtual fabric: enabled” in the command output.

On enabling VFs, a default Logical Switch with FID 128 will be created. All ports on the switch will be assigned to the default Logical Switch, and the **switchname:admin>** prompt will change to **switchname:FID128:admin>**.

3. Create base switches on the ED-DCX-4S-B switches and assign fabric IDs that will become the FID of the base fabric.

---

**Note:** In order to create a base switch, a Logical Switch must be created and defined as a base switch. When the Logical Switch is created, it is automatically enabled and empty — that is, it does not have any ports in it. After creating the logical/base switch, the user must disable the switch to configure it and set the domain ID. Then the user must assign ports to the logical/base switch. Each switch can have only one base switch. The logical/base switches are created with an empty zoning configuration, independent of whether the parent switch had a configuration present.

---

- a. Enter the following command to create a base switch on each of the ED-DCX-4S-B:

```
lscfg --create 127 -base
```

where *127* is the fabric ID that is to be associated with the base switch and the *-base* option is specified to define the Logical Switch as a base switch.

- b. Log in to the newly created Base Switch (FID = 127) by running the following command:

```
setcontext 127
```

where *127* is the fabric ID of the Base Switch that was just created.

- c. Disable the base switch.

**switchdisable**

- d. Configure the switch attributes, including assigning a unique domain ID by running the following command

**configure**

- Enter **y** at the **Fabric Parameters** prompt.

---

**Note:** As a best practice and for ease of management, EMC recommends that the user enters the domain ID of the parent switch from which the Logical Switch is being created. That way, if VFs are being created within an existing fabric, there is no possibility of any domain conflicts occurring within a Virtual Fabric.

---

- Enter **4** for ED-DCX-4S-B with DID:4 and enter **5** for ED-DCX-4S-B with DID:5 at the **Domain ID** prompt
- All other attributes can stay at the default values, including "Allow XISL use," which is set to "yes" by default.
- **Ctrl D** can be used to save changes and exit from this menu.

- e. Enable the base switch by running the following command:

**switchenable**

- f. Both base switches on the different physical switches will come up with the same default name. The user can change the name of these base switches by running the following command:

**switchname** *bsname*

where *bsname* specifies the name assigned to the base switch by the user.

Base switches with FID:127 and user defined Domain IDs and switch names have now been created and enabled on the ED-DCX-4S-B switches in this configuration.

4. Ports can now be assigned to the base switches that were created in [Step 3](#). For this example, we intend to use ICL connections between the base switches, so we will add the ICL ports on each of the ED-DCX-4S-B chassis to the respective base switches.

---

**Note:** All ports on the chassis, including the ICL ports, are initially assigned to the default logical switch (FID 128). If ICLs are being deployed in the base switch, then all ports associated with those ICLs must be assigned to the base switch.

---

- a. On each of the two ED-DCX-4S-B prompts the context must be set to the default Logical Switch to which the ICL ports are currently assigned. For our example, we will run the following command:

```
setcontext 128
```

where *128* is the fabric ID of the default Logical Switch where the ICL ports are currently present.

- b. On each of the two switches, enter the following command to assign ports to the base switches:

```
lscfg --config 127 -slot 3 -port 0-15 (press enter and enter y at the prompt)
```

```
lscfg --config 127 -slot 6 -port 0-15
```

where *127* is the fabric ID of the base switch to which the ports are to be assigned, *3* and *6* are the slot numbers with the ICL ports, and *0-15* is the port range of ICL ports to be assigned to the base switch. If the *-port* option is omitted, all ports on the specified slot are automatically assigned to the logical/base switch.

The ICL ports are automatically disabled, then removed, from their current Logical Switches and assigned to the base switches specified by *FID 127*.

5. Physically connect the ICL ports between the base switches in the ED-DCX-4S-B chassis. These ICL connections form XISLs, which are capable of routing traffic coming from different Logical Switches that will be created on the ED-DCX-4S-B chassis.
6. Enable all of the base switches by enabling the disabled ICL ports that are in the base switch.
  - a. On both the ED-DCX-4S-B prompts the context must be set to the base switch to which the ICL ports have been assigned. For our example we will run the following command:

```
setcontext 127
```



where 127 is the fabric ID of the base switch where the ICL ports are currently present.

- b. All the ICL ports that were disabled by default must now be enabled by running the following commands:

```
iclcfg --enable 3/0 (press enter)
iclcfg --enable 3/1 (press enter)
iclcfg --enable 6/0 (press enter)
iclcfg --enable 6/1
```

3/0, 3/1, 6/0, 6/1 specify the ICL port groups.

- c. Once all the ports are up and connected, run the **fabricshow** command to verify that the two base switches are now showing up as members of the fabric.

This forms the base fabric.

- 7. Configure the Logical Switches on all the VF-capable switches. Logical Switches with the following FIDs and DIDs will be created on the switches specified next:

Switch	Switch DID	FID	Logical Switch ID
ED-DCX-4S-B	4	2	4
ED-DCX-4S-B	4	6	4
ED-DCX-4S-B	4	8	4
ED-DCX-4S-B	5	4	5
ED-DCX-4S-B	5	6	5
ED-DCX-4S-B	5	8	5
ED-DCX-B	6	2	6
ED-DCX-B	6	6	6
DS-5300B	8	4	8
DS-5300B	8	6	8

---

**Note:** When the Logical Switch is created, it is automatically enabled and empty — that is, it does not have any ports in it. After creating the logical/base switch, the user must disable the switch to configure it and set the domain ID. Then the user must assign ports to the logical/base switch. Each switch can have only one base switch. The logical/base switches are created with an empty zoning configuration, independent of whether the parent switch had a configuration present.

---

Use the following steps, [Step a](#) to [Step f](#), to configure the first row in the above table; that is, to create a Logical Switch with FID 2 and DID 4 on the ED-DCX-4S-B (172.23.199.4) with switch DID 4.

- a. Enter the following command to create a base switch on each ED-DCX-4S-B:

**lscfg --create 2**

where 2 is the fabric ID that is to be associated with the Logical Switch.

- b. Set the context to the new Logical Switch.

**setcontext 2**

where 2 is the fabric ID of the Logical Switch that was just created.

- c. Disable the base switch.

**switchdisable**

- d. Configure the switch attributes, including assigning a unique domain ID by running the following command

**configure**

- Enter y at the **Fabric Parameters** prompt.

---

**Note:** As a best practice and for ease of management, EMC recommends that the user enters the domain ID of the parent switch from which the Logical Switch is being created. That way, if VFs are being created within an existing fabric, there is no possibility of any domain conflicts occurring within a Virtual Fabric.

---

- Based on the above note, enter 4 for ED-DCX-4S-B with DID:4.
- All other attributes can stay at the default values, including "Allow XISL use," which is set to "yes" by default.

- **Ctrl D** can be used to save changes and exit from this menu.
- e. Enable the base switch by running the following command:
- switchenable**
- f. The Logical Switches get a default switch name assigned to them. The user can change the name of these base switches by running the following command:
- switchname *lsname***
- where *lsname* specifies the name assigned to the Logical Switch by the user.
- Logical Switches with FID:2 and a user defined Domain ID and switchname have now been created and enabled on the ED- DCX-4S-B (DID:4) switch in this configuration.
- g. [Step a](#) to [Step f](#) above can now be repeated to create the Logical Switches with FIDs and Logical Switch DIDs specified in the table shown in [Step 7](#), starting on [page 257](#).
8. The next step is to assign ports to the Logical Switches created above. The following ports have to be added to the respective Logical Switches with the FIDs and Logical Switch DIDs that have been created on the switches specified in the following table:

Switch	FID	Logical Switch ID	Slot number	Port number(s)
ED-DCX-4S-B	2	4	1	45, 46, 47
ED-DCX-4S-B	6	4	1	48
ED-DCX-4S-B	8	4	1	45
ED-DCX-4S-B	4	5	8	4, 6, 7
ED-DCX-4S-B	6	5	8	5
ED-DCX-4S-B	8	5	8	8
ED-DCX-B	2	6	1	1
ED-DCX-B	6	6	1	2,3
DS-5300B	4	8	-	0, 2
DS-5300B	6	8	-	1, 3

Use the following steps to configure the first row in the above table above; that is, to add ports 1/44, 1/46, 1/47 to the Logical Switch with FID 2 and DID 4 on the ED-DCX-4S-B (172.23.199.4) with switch DID 4.

- a. On the ED-DCX-4S-B (DID 4) prompt, the context must be set to the default Logical Switch to which the ports 1/44, 1/46, 1/47 are currently assigned. For our example, we will run the following command:

```
setcontext 128
```

where 128 is the fabric ID of the default Logical Switch where the ICL ports are currently present.

- b. Enter the following command to assign ports to the base switch:

```
lscfg --config 2 -slot 1 -port 44 (press Enter and enter y at the prompt)
```

```
lscfg --config 2 -slot 1 -port 46-47
```

where 2 is the fabric ID of the base switch to which the ports are to be assigned, 1 if for the slot numbers with the desired ports, and 44, 45, and 46 is the port range of the ports to be assigned to the Logical Switch. If the *-port* option is omitted, all ports on the specified slot are automatically assigned to the logical/base switch.

The specified ports are automatically disabled, then removed, from their current Logical Switches and assigned to the base switches specified by *FID 2*.

- c. [Step a](#) to [Step b](#) above can now be repeated to add the desired ports to the Logical Switches with FIDs and Logical Switch DIDs specified in the table shown in [Step 8](#), starting on [page 259](#).

9. Figure 47 shows the individual Logical Switches and Virtual Fabric configurations.

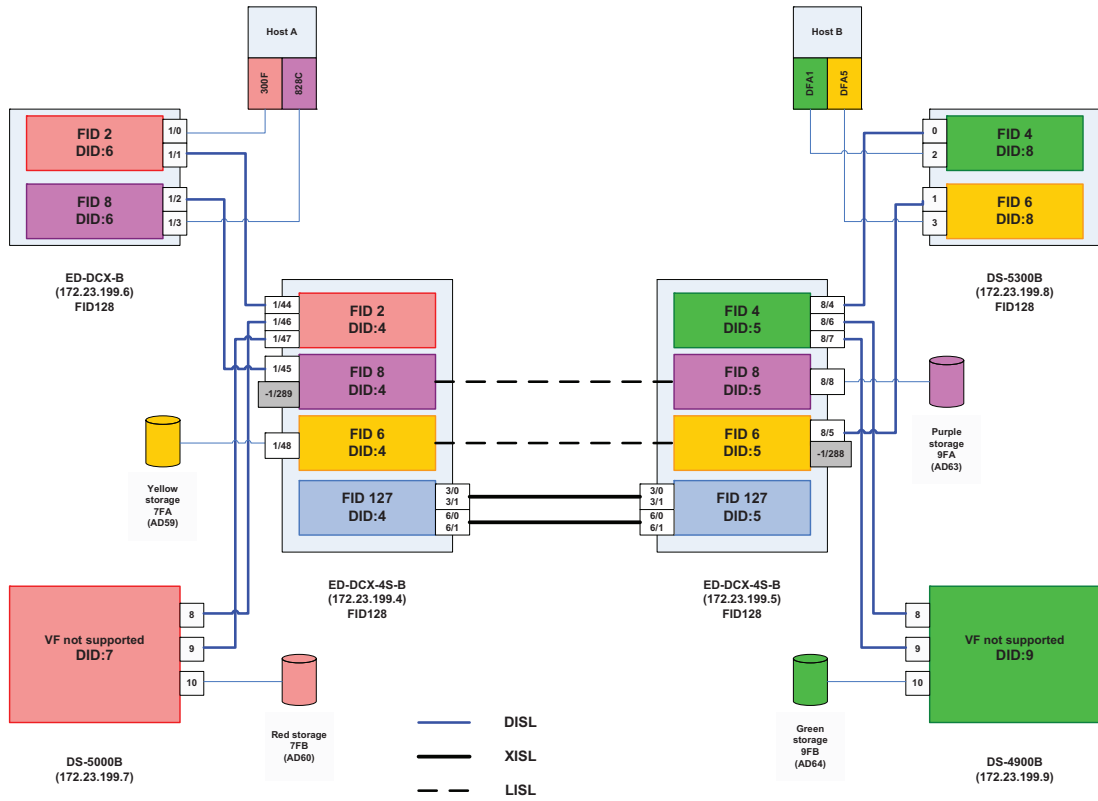


Figure 47 Block diagram of fabric topology

Physically connect devices and ISLs as shown in Figure 47. The ISL connections between two Logical Switches form the DISLs. Also note that ISLs can only be connected between two Logical Switches with the same FIDs.

**Note:** As a best practice, EMC recommends that a minimum of 2 DISLs are connected between any two Logical Switches in a production environment.

10. Enable all the Logical Switches by enabling the disabled ports in each Logical Switch.

The following is an example for one of the Logical Switches, that is, the Logical Switch (LS) with FID 2, LS DID 4 in the ED-DCX-4S-B (DID 4).

- a. On the ED-DCX-4S-B (DID 4) prompt the context must be set to the Logical Switch to that has to be enabled. For our example we will run the following command:

**setcontext 2**

where 2 is the fabric ID of the Logical Switch which has the ports to be enabled.

- b. All the ports that were disabled by default must now be enabled by running the following commands:

**portenable 1/44** (press Enter)

**portenable 1/46** (press Enter)

**portenable 1/47** (press Enter)

where 1/44, 1/46, and 1/47 specify the ports assigned to Logical Switch 2 that have to be enabled.

- c. Repeat [Step a](#) to [Step b](#) to enable all the ports that were assigned to the different logical switches.
- d. Once all the ports are up and connected, run the **fabricshow** command from each unique Logical Switch FID to verify that all the Logical Switches with the same FID are now showing up as members of the fabric. For example, set the context on the ED-DCX-4S-B to 2:

**setcontext 2**

and then run the **fabricshow** command.

This should display all Logical Switches that are a part of this fabric. All these Logical Switches will have the same FIDs. Thus, all Logical Switches with the same FID form a Virtual Fabric, as shown in the following table.

Virtual Fabrics ID	Switch	Switch ID	FID	Logical Switch ID	Color
2	ED-DCX-4S-B	4	2	4	Red
	ED-DCX-B	6	2	6	
	DS-5000B	7	-	-	

Virtual Fabrics ID	Switch	Switch ID	FID	Logical Switch ID	Color
4	ED-DCX-4S-B	5	4	5	Green
	DS-5300B	8	4	8	
	DS-4900B	9	-	-	
6	ED-DCX-4S-B	5	6	5	Yellow
	DS-5300B	8	6	8	
	ED-DCX-4S-B	4	6	4	
8	ED-DCX-4S-B	5	8	5	Purple
	ED-DCX-B	6	8	6	
	ED-DCX-4S-B	4	8	4	

Switches displayed in the basic fabric is shown in the following table.

Basic Fabric ID	Switch	Switch ID	FID	Logical Switch ID	Color
127	ED-DCX-4S-B	4	127	4	Blue
	ED-DCX-4S-B	5	127	5	

- We will now configure some of the Logical Switches to use XISLs.

From [Figure 47 on page 261](#), we can see that there are no DISLs between the Logical Switches with FID 6, DID4 and FID 6, DID 5 in the ED-DCX-4S-B switches, and between the Logical Switches with FID 8, DID 4 and FID 8, DID 5 in the ED-DCX-4S-B switches. These Logical Switches belonging to Virtual Fabrics 6, 8 will be configured to use XISLs.

Actually, the user does not have to do anything special in the case of this example because all Logical Switches that were created are configured to allow XISL use by default. An exception to this are Base Switches and the Default Logical Switch that get created when VF is enabled on the switch. The "Allow XISL use" is set to OFF for these Logical Switches.

This can be verified by running the **switchshow** command on every Logical Switch and validating that the "Allow XISL use" attribute is set to **ON**.

---

**Note:** If Logical Switches have the option to use either a DISL or the XISL, they will use the DISL to route traffic, since it is a lower cost path as compared to an XISL. The Logical Switch will always look for DISLs first to route the traffic to another Logical Switch in the same Virtual Fabric.

---

If the "Allow XISL use" attribute is set to **OFF**, complete the following steps:

- a. Disable the base switch.  
**switchdisable**
- b. Configure the switch attributes by running the following command  
**configure**
- c. Enter **y** at the **Fabric Parameters** prompt.
- d. Press **Enter** at the **Domain ID** prompt and enter **y** at **Allow XISL use** prompt.
- e. All other attributes can stay at their default values.
- f. **Ctrl D** can be used to save changes and exit from this menu.

---

**Note:** Every Logical Switch in the Virtual Fabric where XISLs are being used must have XISL use enabled. This must be set on every individual switch in the fabric. It is not a fabric-wide enabled feature.

---

For this example, since only Logical Switches with FID 6 and FID 8 will actually be using the XISLs, a **switchshow** on the Logical Switch with FID 8, DID 4 will display a logical E\_Port with -1 for the slot and a virtual port number 289 for the slot port. Similarly, a **switchshow** on the Logical Switch with FID 6, DID 5 will display a logical E\_Port with -1 for the slot and a virtual port number 288 for the slot port. This is an implication that this Logical Switch is using XISLs to route traffic.



## Zoning with Virtual Fabrics

Each Virtual Fabric has its own zoning configuration. The following table specifies the devices that are attached to each Virtual Fabric, as can also be seen from [Figure 47 on page 261](#).

Virtual Fabric ID	Initiator color/WWN	Target color/WWN
2	Red / 10:00:00:00:c9:38:30:0f	Red / 50:00:09:72:08:13:AD:60
8	Purple / 10:00:00:00:c9:38:82:8C	Purple / 50:06:04:82:cc:19:AD:63
4	Green / 21:01:00:e0:c8:b8:DF:A1	Green / 50:06:04:82:cc:19:AD:64
6	Yellow / 21:01:00:e0:c8:b8:DF:A5	Yellow / 50:06:04:82:cc:19:AD:59

The following steps are used to configure a zone for VF\_ID 2 with end devices specified in the first row of the above table.

1. Select any Logical Switch in Virtual Fabric 2. For this example, we will select the Logical Switch with FID 2, DID 6 in ED-DCX-B. This switch can be accessed by setting the context on the ED-DCX-B switch to FID 2 as shown next:

```
setcontext 2
```

where 2 is the fabric ID of the Logical Switch where a zone configuration has to be created.

2. Create zones using the **zonecreate** command:

```
zonecreate "HostA_RedHBA_Symm_7FB",
"10:00:00:00:c9:38:30:0f; 50:00:09:72:08:13:ad:60"
```

3. Create the zone configuration using the **cfgcreate** command:

```
cfgcreate "VF2_cfg" , "HostA_RedHBA_Symm_7FB"
```

4. Enable the zone configuration using the **cfgenable** command.

```
cfgenable "VF2_cfg"
```

5. Enter **y** at the confirmation prompt.

6. Enter **cfgshow** to display zoning information.

The zone information should be similar to what is shown next.

### Defined configuration:

```
cfg: VF2_cfg
HostA_RedHBA_Symm_7FB
zone: HostA_RedHBA_Symm_7FB
10000000c938300f; 500009720813ad60
```

**Effective configuration:**

```
CFG VF2_cfg
Zone: HostA_RedHBA_Symm_7FB
10000000c938300f
500009720813ad60
```

7. One Logical Switch from each of the other FIDs: 8, 4, 6, can be selected and [Step 1](#) through [Step 6](#) must be executed. The zone information on Logical Switches on the switches in the FIDs should appear as follows:

**On FID 8:****Defined configuration:**

```
cfg: VF8_cfg
HostA_PurpleHBA_Symm_9FA
zone: HostA_PurpleHBA_Symm_9FA
10000000c938828c; 500009720813ad63
```

**Effective configuration:**

```
CFG VF8_cfg
Zone: HostA_PurpleHBA_Symm_9FA
10000000c938828c
500009720813ad63
```

**On FID 4:****Defined configuration:**

```
cfg: VF4_cfg
HostA_GreenHBA_Symm_7FA
zone: HostA_GreenHBA_Symm_7FA
210100e0c8b8dfa1; 500009720813ad63
```

**Effective configuration:**

```
CFG VF4_cfg
Zone: HostA_GreenHBA_Symm_7FA
210100e0c8b8dfa1
500009720813ad63
```

**On FID 6:****Defined configuration:**

```
cfg: VF6_cfg
HostA_YellowHBA_Symm_9FB
zone: HostA_YellowHBA_Symm_9FB
210100e0c8b8dfa5; 500009720813ad64
```

**Effective configuration:**

```
CFG VF6_cfg
Zone: HostA_YellowHBA_Symm_9FB
210100e0c8b8dfa5
500009720813ad64
```

---

**Note:** To verify the VF state on the switch, use the following command:  
**lfcfg --show**

To verify the zoning configuration on the switch, use the following command: **cfgshow**

For information on the following **show** commands and for details on the supported platforms for Virtual Fabrics, Virtual Fabrics interaction with other FOS features, and limitations and restrictions of Virtual Fabrics, refer to the *EMC Connectrix Administrator's Guide for FOS v6.2.0e*.

---

## Brocade Virtual Fabrics versus traditional Cisco Virtual SANs

This section explains why Brocade Virtual Fabrics (VF) configuration needs more planning on the user-front than Cisco's Virtual SANs (VSAN) configuration. For more basic information on VSANs, refer to "VSANs" in the *Networked Storage Concepts and Protocols TechBook*, available on the [E-Lab Navigator](#), **Documents > Topology Resource Center**.

It also highlights the design-based differences that a user, familiar with the traditional Cisco VSANs but deploying Brocade Virtual Fabrics, must be aware of. [Table 9](#) maps existing Cisco VSAN terminology with the new Brocade Virtual Fabrics terminology.

**Table 9 Cisco VSAN versus Brocade VF terminology**

Cisco	Brocade
VSAN	VF
VSAN id	FID
ISLs within a VSAN	DISL (between Logical Switches)
EISL	XISL (between base switches)
IVR	FCR between VFs

The following information are the design-based differences that a Brocade VF user, familiar with Cisco VSANs configurations, must be aware of.

### Domain IDs must be allocated to Logical Switches while partitioning an existing fabric into logical/virtual fabrics

- VF** The user needs to know what switches within an existing fabric will need to be a part of a VF before creating the VF. When Logical Switches that can participate in the same VF are created on the independent physical switches, their default Domain IDs get set to **1**. All Logical Switches within the same VF *must* have the *same* FID, but definitely *not* the same Domain ID or they will segment (due to a Domain ID conflict). Thus, the user needs to go into every Logical Switch within the VF and reconfigure them with unique Domain IDs.

**VSAN** This is not the case with Cisco VSANs. When VSANs get created, the switches that are participating in a VSAN retain their original Domain ID so the user does not have to reconfigure them.

---

### Zoning information may be shared between Logical Switches, even when FIDs may not match

**VF** While creating a VF within an operational fabric, it is *essential* to verify that the E\_Ports between Logical Switches (on different physical switches) being connected to each other have the *same* FID. If there is a mismatch between the FIDs, the management application will notify the user that the ISL was segmented due to an FID conflict. In case one of the segmented switches does not have an active zoning configuration, it will import the configuration from the switch it was linked to through ISLs. The reason for this is that when two Logical Switches are linked through ISLs, the following sequence of checks take place:

- ◆ Domain ID check.
- ◆ If the Domain ID check passes, then a zone check, operating mode, and FID check takes place. In most cases, the zone check takes place before the FID check. If one of the Logical Switches has no zone configuration, it pulls it from the other Logical Switch to which it is attached. It later segments due to an FID conflict (if the FIDs do not match). The user now needs to go into, and delete, the Logical Switch that imported the undesired zone configuration.

**VSAN** This is not the case with Cisco VSANs, which work by prepending a Virtual Fabric header tag onto each frame as it enters the switch, which indicates on which VSAN the frame arrived. This tag is subsequently used when the switch makes distributed hardware forwarding decisions. Cisco frame forwarding ASICs are Virtual Fabric-aware and use the industry-standard Virtual Fabric Trunking (VFT) extended headers in forwarding decisions, therefore would not allow any transactions between two switches that have ports belonging to different VSANs.

---

## Switch name must be configured on the Logical Switches

**VF** When a Logical Switch (LS) is created on a switch, the new LS uses a default switch name. Therefore, in a given Virtual Fabric all Logical Switches will end up with the same switch name if the user does not redefine them. This has been purposely designed to provide flexibility to administrators (who may be using the different VFs), to use their specific naming conventions within their VFs, rather than having to use a pre-defined switch name.

**VSAN** For Cisco switches, the default name is its IP address and this name, or the user-defined name, is retained even when its ports become a part of different VSANs. The user then has the ability to edit it, if desired.

---

**Note:** For users that may not expect the Logical Switch settings to go to a default value, (which addresses the configuration specifics discussed previously in this section), Brocade has added the following CLI warning message after a Logical Switch is created:

“Logical Switch has been created with default configurations. Please configure the Logical Switch with appropriate switch and protocol settings before activating the Logical Switch.”

---

---

## Deleting a VF

**VF** If the user wants to delete a VF using Brocade CLI, the following two steps must be executed:

1. All the ports that are a part of the Logical Switches in that VF need to be transferred to the default Logical Switch (FID 128), or any other Logical Switch.
2. All the Logical Switches within the VF need to be individually deleted.

**VSAN** With the Cisco CLI, a VSAN can be deleted, which pulls out all the ports on that switch that were participating in the VSAN and automatically transfers them into the default VSAN (VSAN ID: 4094).

With Cisco Fabric Manager, once a VSAN is deleted all the ports under the participating switches in that VSAN get automatically transferred to the respective default VSANs on those physical switches.

---

## Brocade Base Switches concept

**VF** Brocade's VF introduces the concept of a *Base Switch*. A Base Switch is a Logical Switch (LS) that comprises of E\_Ports that can connect to ports in other base switches only, forming XISLs, which are capable of carrying traffic from multiple VFs. Therefore, for Brocade switches, the user has to create a base switch with the same FID on two switches that are a part of the same VF, rather than just configuring a port that is in the Default Logical Switch.

**VSAN** For Cisco switches, ports that are not a part of any VSAN (on a switch with more than one VSAN) can be configured as TE\_Ports (trunk E\_Ports). When two TE\_Ports are connected, they form an EISL. An EISL can carry traffic coming from multiple VSANs. The TE\_Port does not need to be in any VSAN.

---

## FCR with Brocade VFs

**VF** Brocade's FCR technology can be used to route traffic between different VFs. A Base Switch in the fabric will be used for legacy FCR support. All EX\_Ports need to be part of the Base Switch. The EX\_Ports will be disabled if they are in non-base switch. "[Extended ISL \(XISLs\) and Logical ISLs \(LISLs\)](#)" on page 248 explained how a Base switch is created and can be connected to other Base Switches through a special XISL to form a Base Fabric. For FCR between VFs, users can view the base switch as a FCR-backbone switch and Base Fabric as FCR-backbone fabric (in Brocade FCR concept). The Base switches can have EX\_Ports attached to E\_Ports on the Logical Switches that are a part of different VFs with the end devices that have to communicate. Thus, the VFs form the edge fabrics. LSAN zones have to be created on the edge VFs, (VFs with the end devices) so that the initiator on one VF can access the target in another VF.

**VSAN** Cisco IVRs allows data traffic to be transported between specific initiators and targets on different VSANs without merging the VSANs into a single logical fabric. IVR is not limited to VSANs present on a common switch. Routes that traverse one or more VSANs across multiple switches can be established, if necessary, to establish proper interconnections.





This chapter provides the following information on FICON topologies.

◆ Overview .....	274
◆ Topology support .....	276
◆ Zoning practices .....	277
◆ Cascading .....	278
◆ Terminology .....	279
◆ IOCP considerations .....	281
◆ FICON and EMC ControlCenter .....	283
◆ CUP (fabric management server) .....	284
◆ Connectrix B series .....	285
◆ Connectrix MDS series .....	288

## Overview

FICON (Fibre Connectivity) is an I/O channel designed to support low-latency, high-bandwidth connections between a mainframe and a storage controller. FICON builds on the technology of Fibre Channel, sharing the lower levels of Fibre Channel, including FC-0 (Physical Interface), FC-1 (8b/10b Encode/Decode), FC-2 (Framing and Signaling), and FC-3 (Common Services).

FICON is an FC-4 type. It was designed as a replacement for ESCON to support mainframe attached CKD (Count Key Data) formatted storage systems. FICON transports ESCON architecture frames over Fibre Channel. This is analogous to SCSI FCP transport of SCSI commands over Fibre Channel.

The FICON standards are developed and maintained by the T11 Technical Committee of the International Committee for Information Technology Standards. Drafts and other technical documentation can be found at [www.T11.org](http://www.T11.org) under FC-SB.

Table 10 lists some attributes of FICON and ESCON.

**Table 10 FICON compared to ESCON**

Attribute	FICON	ESCON
Link rate	212 MB/s	20 MB/s
Effective max data rate	200 MB/s	17 MB/s
Duplex	Full duplex	Half duplex
Type of switching	Packet switching with frame-by-frame routing using FSPF and Classes 2 and 3	Virtual Circuits (Class 1)
Multiple concurrent I/O	Yes	No
Maximum distance without droop	Depends on BB_Credit and link speed (typically 100 km at 100 MB/s).	9 km

The FICON technologies described in this Topology Guide are organized as follows:

- ◆ General FICON connectivity information — Covers the topic with switch vendor-neutral terminology and concepts.
- ◆ Vendor-specific information — Covers the topic by switch vendor, applying the vendor's terminology, product names, support specifics, and unique considerations.

Refer to:

- “Connectrix B series” on page 285
- “Connectrix MDS series” on page 288
- “Connectrix MDS series” on page 288

## Reference material

Refer to the following for more information:

- ◆ Connectrix documentation on EMC Online Support at <https://support.emc.com>:
- ◆ HCD/IOCP: <http://www.ibm.com/servers/resourcelink>  
Select **Library**, then the appropriate Mainframe CPU
- ◆ IBM Redbooks, at <http://www.ibm.com/redbooks>:
  - *FICON Native Implementation and Reference Guide*, PN SG24-6266-01
  - *Getting Started with the Brocade M Series Intrepid FICON Director*, PN SG24-6857-00
  - *Getting Started with the Inrange FC/9000 FICON Director*, PN SG24-6858-00
  - *Cisco FICON Basic Implementation*, PN REDP-4392-00

## Topology support

Topology support for FICON covers:

- ◆ Switch support for FICON-to-Symmetrix connectivity by switch vendor, model number, and recommended firmware levels
- ◆ Intermixing FICON and FCP on the same switch
- ◆ Intermixing FICON and SRDF on the same switch
- ◆ Cascading (multiswitch fabric support)
- ◆ Size of fabric
- ◆ Interoperability among switch vendors
- ◆ Compatibility between a FICON environment and EMC ControlCenter

Each of these topology capabilities is determined on a per-vendor, per-model basis. Refer to the following for more information:

Connectrix switch series	Topology Guide reference	Vendor website
M series	"Connectrix MDS series" on page 288	<a href="http://www.Brocade M Series.com">http://www.Brocade M Series.com</a>
B series	"Connectrix B series" on page 285	<a href="http://www.Brocade.com">http://www.Brocade.com</a>
MDS series	"Connectrix MDS series" on page 288	<a href="http://www.Cisco.com">http://www.Cisco.com</a>

---

## Zoning practices

The recommended zoning practice for FICON environments is to build a single FICON zone and include all FICON N\_Ports, channels, and control units in that zone. FICON channels depend on State Change Notifications (SCNs) from the switch to perform error recovery.

A single zone is administratively simple, and insures that all FICON N\_Ports receive their SCNs. Zoning based on World Wide Port Name (WWPN) is recommended, but port-based zoning is also supported. This practice varies from the Single Initiator zoning of Open Systems, because mainframe channels do not depend on the name server for device discovery. Instead, mainframes use the device address information in the IOCDS for discovery. (“IOCP considerations” on page 281 provides more information.)

If the switch or fabric includes Open Systems (intermix) or SRDF ports, these ports should be zoned per practices. For Open Systems, this is single-initiator zoning based on WWPN. For SRDF, place all SRDF ports into a single SRDF zone.

---

**Note:** Refer to “Use single initiator zoning” on page 20 for information on SRDF zoning.

---

## Cascading

Cascading is a FICON topology where the channel and the control unit are on different switches, connected through an interswitch link (ISL). There are additional security and port addressing considerations in a cascaded switch environment. There is a limitation of one 'hop' between switches for cascaded FICON.

Refer to the following for more information:

Connectrix switch series	Topology Guide reference	Vendor website
M series	"Connectrix MDS series" on page 288	<a href="http://www.Brocade M Series.com">http://www.Brocade M Series.com</a>
B series	"Connectrix B series" on page 285	<a href="http://www.Brocade.com">http://www.Brocade.com</a>
MDS series	"Connectrix MDS series" on page 288	<a href="http://www.Cisco.com">http://www.Cisco.com</a>

## Terminology

- ◆ **Entry switch** — A FICON switch that is connected to the processor's FICON channel(s) and to a cascaded switch. An entry switch can also be a cascaded switch.
- ◆ **Cascaded switch** — A FICON director that connects to a control unit (, for example, a storage array) and to an entry switch. A cascaded switch can also be an entry switch.
- ◆ **Switch ID** and **switch address** (1 byte) — Ways to address a FICON director. The Switch ID and Switch address describe fields used in the IOCDS. Both values should be set equal to the Domain ID of the switch.

The IOCDS expects physical values in hex. In this context, the physical value refers to the value that would be seen in a Finisar/I-Tech trace. The logical value would be the value as seen from a management application. Since the switch vendors may use logical values (a physical value plus an offset) and display the Domain ID in decimal, care must be taken when transferring this information from the switch's management application to the IOCDS. Refer the specific switch vendor information:

- [“Connectrix B series” on page 285](#)
- [“Connectrix MDS series” on page 288](#)
- [“Connectrix MDS series” on page 288](#)

The switch ID is assigned by the customer, and must be unique within the scope of the IOCP (Input/Output Configuration Program) or HCD (Hardware Configuration Definition). The switch ID is an arbitrary unique number given to identify the switch. It is highly recommended that the switch ID be set equal to the Domain ID.

The link statement in the IOCDS is where the actual switch Domain ID and port address are defined. For example, the link statement **LINK= (6104,6104)** defines a connection to port 4 on Domain ID 1.

The switch address is the Domain ID of the switch. These addresses can be customized to a preplanned value. They must be unique within a fabric.

- ◆ **Port address** — Address (with a one-byte value) of the physical FICON director port.

- ◆ **Insistent Domain ID** — A switch using the Insistent Domain ID feature ensures that it will join a fabric if (and only if) its administratively assigned Domain ID is granted during the fabric initialization procedure.

### Notes on cascading

Switch addresses for FICON directors must be unique across all of the processor's CHPIDs (Channel Path Identifiers).

With SYSPLEX installations, the switch address must be unique across the SYSPLEX complex.

Cascading requires the use of two-byte link addresses that consist of a high-order byte defining the switch address (Domain ID) and a low-order byte that defines the FICON director port address.

Two-byte link addresses require that the FICON director have the Fabric Binding and Insistent Domain ID features installed. The channel checks this during initialization by sending the Security Attributes Extended Link Services query to the switch. In the response from the switch, the channel checks that Fabric Binding Enforcement and Insistent Domain ID bits are set. If they are not set, the link does not complete initialization.



## IOCP considerations

Figure 48 shows an example of cascaded IOCP.

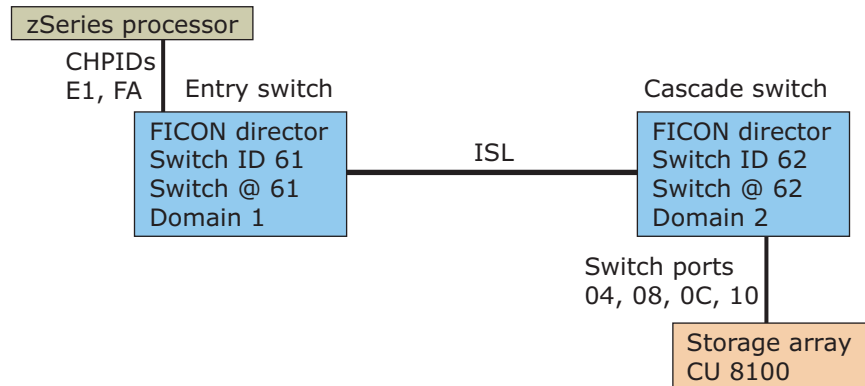


Figure 48 Cascaded IOCP

A cascaded FICON IOCP based on this figure would look like this:

### Type keyword

CHPID will operate in FICON native mode, required for cascading

### Switch keyword

Logical switch number  
Channel entry switch.id, required for FICON (FC mode)

```
CHPID PATH=(E1), SHARED, TYPE=FC, SWITCH=61
```

```
CHPID PATH=(FA), SHARED, TYPE=FC, SWITCH=61
```

```
CNTLUNIT CUNUMBR=8100, PATH=(E1, FA), UNIT=2032,  
UNITADD=(00, 1), LINK=(6204, 6208, 620C, 6210)
```

### Link keyword

Two bytes  
Destination port address, one for each switched path  
Switch address in high byte plus port address in low byte

```
IODEVICE (Same as for non FICON)
```

Figure 49 Cascaded FICON IOCP

Once a two-byte link address has been specified for a channel path:

- ◆ It is recommended that you specify two-byte link addresses for all paths from the same CEC that have the same entry switch.
  - This allows IOCP to perform better checking of the switch configuration.
  - HCD requires additional information and performs a different checking method; therefore, the above is an IOCP recommendation only.
- ◆ All link addresses defined for access from the same channel path *must* be two-byte link addresses.

---

## FICON and EMC ControlCenter

The EMC ControlCenter switch management functions (including zoning) are qualified on a switch vendor basis. Refer to the specific switch vendor information for the latest support.

Any time there is more than one management control point, care must be taken to insure that administrative conflicts do not arise. With ControlCenter there can be three separate and potentially conflicting control points: ControlCenter, Connectrix Manager, and System Automation OS/390. It is largely the administrator's responsibility to insure that conflicting controls are avoided.

As a result, having a single administrator in an intermix environment for all Open Systems and FICON switch management is the preferred practice.

## CUP (fabric management server)

Control Unit Port (CUP) is an inband communications protocol that allows an OS/390 to manage a switch with the same level of control and security that S/390 has on an ESCON switch. Control functions include blocking and unblocking ports, as well as error reporting and monitoring.

CUP is also required for other mainframe management functions, such as:

- ◆ HCD — Port activation/deactivation
- ◆ HCM — Performance display
- ◆ SMF — Performance statistics gathering for RMF
- ◆ ZOS — Error handling, problem determination
- ◆ DCM — Dynamic CHPID management

The services provided by CUP are used by IBM management software System Automation for OS/390.

CUP is not a requirement for mainframe/storage FICON connectivity. Typically, large mainframe data centers utilize its services.

---

## Connectrix B series

This section contains the following information:

- ◆ “Supported products” on page 285
- ◆ “Configuring” on page 285
- ◆ “OCP considerations” on page 286
- ◆ “CUP support” on page 286
- ◆ “Switch node identifier” on page 287
- ◆ “EMC documentation” on page 287

---

### Supported products

These B series products support FICON connectivity:

- ◆ DS-32B2
- ◆ DS-4100B
- ◆ ED-12000B
- ◆ ED-24000B
- ◆ ED-48000B
- ◆ DS-4900B
- ◆ PB-48K-18i blade (FC-IP)
- ◆ Silkworm 7500 (FC-IP)

---

### Configuring

#### Topology support

- ◆ Single switch is supported.
- ◆ Cascading is supported with FOS 5.2.1 or higher
  - Cascading is *not* supported with ED-12000B.
- ◆ Intermixing FICON and FCP on the same switch is supported.

In a FCP/FICON intermix environment, as long as the FICON N\_Ports, Channel, and Control Unit reside on the same domain, and the IOCDs uses single byte addressing, FICON I/O in a cascaded Open Systems Brocade fabric is supported.

- ◆ Intermixing FICON and SRDF on the same switch is supported.
- ◆ B series/M series interoperability is not supported in a FICON environment.

### Recommended FICON environment configuration settings

- ◆ The older FICON products may not connect to 8 Gb ports. The 8 Gb SFP negotiates only to 2, 4 or 8 Gb link speeds.
  - Older FICON products may only support a 1 Gb link speed.
  - A switch port with a 4 Gb SFP will be required for connectivity. The 4 Gb SFP negotiates to 1, 2, or 4 Gb link speeds.
- ◆ The fillword needs to be set to ARB(FF) for 8 Gb to 8 Gb connections. The IDLE fillword is still used for slower link speeds.
- ◆ Enable in-order delivery (**iodset** command).

---

## OCP considerations

Switch ID Definition — No offsets on switch ID or port address, but every value must be in hex for the mainframe.

---

## CUP support

CUP is a separately licensed inband management service installed on each switch.

You can manage CUP using Web Tools or Fabric Manager. Limited support for CUP is provided through the Fabric OS CLI.

CUP provides the following advantages:

- ◆ Single point of control and monitoring for channel, director (switch), and control unit.
- ◆ Automated tools on the mainframe can take advantage of the statistics to move channels where they are needed. This cannot be done from the switch alone.
- ◆ Seamless integration into existing management tools that are also used to manage ESCON directors (switches). This makes migration from ESCON to FICON smoother.

You can monitor the FICON director (switch) using CUP to obtain the following port statistics:

- ◆ Number of words transmitted
- ◆ Number of words received
- ◆ Frame Pacing Time (the number of 2.5 ms units that frame transmission is blocked due to zero credit).

Refer to the IBM document *FICON Director Programming Interface with Cascading Support*.

---

## Switch node identifier

You can find switch node information such as the serial number and manufacturer name. This information is the same as the Switch Node ID in the RNID ELS:

- ◆ Configuration file information — Provides a list of configuration files on the switch. You can also obtain the actual file content, including port address name and port connectivity.
- ◆ History summary (director history buffer) — The history buffer logs each change in status or configuration of the ports. You can retrieve the history buffer using CUP.
- ◆ Switch configuration data — Provides switch configuration data such as time-out values and number of ports per card.

You can find information on CUP functions and commands in the IBM-proprietary document *FICON Director Programming Interface With Cascading Support*.

---

## EMC documentation

Refer to these documents for more information on FICON connectivity:

- ◆ *EMC Connectrix B Series Version 4.4 Features Guide*
- ◆ *EMC Connectrix B Series Fabric OS Version 4.4 Procedures Guide*

---

## Connectrix MDS series

This section contains the following information:

- ◆ “Supported products” on page 288
- ◆ “Requirements” on page 288
- ◆ “Configuring” on page 289
- ◆ “OCP considerations” on page 290
- ◆ “CUP support” on page 290
- ◆ “FICON configuration files” on page 290
- ◆ “Switch node identifier” on page 291
- ◆ “FICON port numbering” on page 292
- ◆ “References” on page 293

---

### Supported products

These MDS series products support FICON connectivity:

- ◆ MDS 9513
- ◆ MDS 9509
- ◆ MDS 9506
- ◆ MDS 9216
- ◆ MDS 9216A
- ◆ MDS 9216i

---

### Requirements

Requires the purchase and installation of the MDS mainframe license package. Note that MDS 9500 and 9200 series have different model numbers for licenses.



## Configuring

### Topology support

Note the following:

- ◆ You must create a FICON VSAN.
- ◆ A multiswitch environment is supported.
- ◆ Cascading is supported with one ISL hop.
- ◆ Intermixing FICON and FCP on the same switch is supported using separate VSANs.

Use standard or interop mode VSANs for FCP traffic, and FICON VSANs for FICON traffic.

- ◆ Intermixing FICON, SRDF, MirrorView™, Open Replicator, and SAN Copy™ on the same switch is supported.
- ◆ Connectrix M series/B series interoperability is supported if those switches exist in another VSAN set for *interop-1* mode.
- ◆ FICON over IP is not supported.
- ◆ FC Write Acceleration is not supported for FICON.
- ◆ TE\_Ports/EISLs may support FICON and FCP traffic. If dedicated bandwidth for FICON is required, create multiple ISLs dedicated to the FICON VSAN.
- ◆ Port Channels for FICON or mixed FCP/FICON traffic is supported.

FICON ports must be bound to a Port Channel.

### Recommended FICON environment configuration settings

- ◆ Configure ports that are connected to 1 GB/s channels for fixed 1 GB/s speed. Otherwise, when using fixed 1 GB/s channels (both G5 and FICON Express), the FICON host might generate erroneous link incidents when the channels are coming online. These link incidents will result in a call home. Other than the generated link incident, the channel will come online and function normally.
- ◆ Enable in-order delivery.

---

## OCP considerations

Switch ID Definition — No offsets on switch ID or port address, but every value must be in hex for the mainframe.

---

## CUP support

Control Unit Port (CUP) is supported by switches and directors in the Cisco MDS 9000 Family. The CUP function allows the mainframe to manage the Cisco MDS switches.

CUP comes with the Mainframe license package required for FICON support.

CUP provides the following advantages:

- ◆ Single point of control and monitoring for channel, director (switch), and control unit.
- ◆ Automated tools on the mainframe can take advantage of the statistics to move channels where they are needed. This cannot be done from the switch alone.
- ◆ Seamless integration into existing management tools that are also used to manage ESCON directors (switches). This makes migration from ESCON to FICON smoother.

You can monitor the FICON director (switch) using CUP to obtain the following port statistics:

- ◆ Number of words transmitted
- ◆ Number of words received
- ◆ Frame Pacing Time (the number of 2.5 ms units that frame transmission is blocked due to zero credit).

The IBM document *FICON Director Programming Interface with Cascading Support* contains more information on CUP.

---

## FICON configuration files

You can save up to 16 FICON configuration files on each FICON-enabled VSAN (in persistent storage).

FICON configuration files contain the following configuration for each implemented port address:

- ◆ Block
- ◆ Prohibit mask
- ◆ Port address name

---

## Switch node identifier

You can find switch node information, such as the serial number and manufacturer name. This information is the same as the switch node ID in the RNID ELS:

- ◆ Configuration file information — You can get a list of configuration files on the switch. You can also obtain the actual file content, including port address name and port connectivity.
- ◆ History summary (director history buffer) — The history buffer logs each change in status or configuration of the ports. You can retrieve the history buffer using CUP.
- ◆ Switch configuration data — Provides switch configuration data as time-out values and number of ports per card.

You can find more information on CUP functions and commands in the IBM-proprietary document *FICON Director Programming Interface With Cascading Support*.

## FICON port numbering

Table 11 represents the logical port number for the MDS 9216, MDS 9506, and MDS 9509.

**Table 11 FICON port numbering in the MDS 9000 family**

Product	Slot number	Implemented port allocation		Unimplemented ports	Notes
		To ports	To portchannel/FCIP		
MDS 9200 series	Slot 1	0 through 31	64 through 89	90 through 253 and port 255	Similar to a switching module
	Slot 2	32 through 63			The first 16 port numbers in a 16-port module are used and the rest remain unused.
MDS 9506 Director	Slot 1	0 through 31	128 through 153	154 through 253 and port 255	The first 16 port numbers in a 16-port module are used and the rest remain unused.
	Slot 2	32 through 63			
	Slot 3	64 through 95			
	Slot 4	96 through 127			
	Slot 5	None			
	Slot 6	None			Supervisor modules are not allocated port numbers.
MDS 9509 Director	Slot 1	0 through 31	224 through 249	250 through 253 and port 255	The first 16 port numbers in a 16-port module are used and the rest remain unused.
	Slot 2	32 through 63			
	Slot 3	64 through 95			
	Slot 4	96 through 127			
	Slot 5	None			
	Slot 6	None			
	Slot 7	128 through 159			The first 16 port numbers in a 16-port module are used and the rest remain unused.
	Slot 8	160 through 191			
	Slot 9	192 through 223			

---

## References

These documents and resources contain for more information on FICON connectivity:

- ◆ [E-Lab Navigator](#), for supported MDS firmware for FICON usage
- ◆ [Connectrix MDS Release Notes](#), for specific information related to new firmware
- ◆ *Cisco MDS 9000 Family Configuration Guide*, available at <http://www.Cisco.com>
- ◆ *Implementing the Cisco MDS 9000 in an Intermix FCP, FCIP, and FICON Environment* (part number SG24-6397-00), available at <http://www.redbooks.ibm.com>



**B**

blade switch with direct attached storage 97  
Bottleneck Detection 318, 330

**C**

Cisco Inter VSAN Routing (IVR)  
    in a heterogeneous environment 285  
complex Fibre Channel SAN topologies 116  
compound core edge switches 153  
core-edge topology, fabric design 464  
core-edge topology, implementing 464  
core-edge topology, overview 464  
CUP (Control Unit Port) 500

**E**

E\_Port interoperability 189  
Edge Hold Time 319

**F**

fabric resiliency  
    concepts 326  
    conditions 326  
    features 318  
    thresholds 320

**FICON**

    and EMC ControlCenter 499  
    cascading 494  
    terminology 495  
    topology support 492  
    zoning 493

FICON connectivity 487, 490  
    for Connectrix B series 501

    for Connectrix MDS series 510

**H**

Heterogeneous interoperability  
    in EMC context 189  
    switch 188  
    test information 297  
Heterogeneous SAN design 191

**I**

Interoperability modes  
    EMC-tested 295  
interoperable switched fabric topology  
    set up 191  
IOCP considerations 497

**L**

latency  
    detection 330  
    severity 329

**M**

multi-vendor switch configuration 190

**S**

SAN, monitor 309  
scalable core-edge SAN topology, deploying 468  
scalable core-edge topology, fabric design 464  
scalable core-edge topology, implementing 464  
scalable core-edge topology, overview 464  
Simple Fibre Channel SAN topologies 39, 42

- single switch fabrics 42
- switch
  - node identifier 503
- switch configuration
  - multi-vendor 190

**T**

- two switch fabrics 62