



New and Changed Information

As of Cisco DCNM Release 5.2, Cisco Fabric Manager and Cisco Data Center Network Manager for LAN are merged into one unified product called Cisco Data Center Network Manager (DCNM) that can manage both LAN and SAN environments. As a part of this product merger, the name Cisco DCNM for SAN replaces the name Cisco Fabric Manager.

The following documentation changes support the merged Cisco DCNM product:

- Cisco DCNM product documentation for Cisco DCNM Release 5.2 is retitled with the name Cisco DCNM for LAN.
- Cisco Fabric Manager product documentation for Cisco DCNM Release 5.2 is retitled with the name Cisco DCNM for SAN.
- Cisco DCNM for SAN product documentation is now published to the Data Center Network Manager listing page on Cisco.com:
http://www.cisco.com/en/US/products/ps9369/tsd_products_support_configure.html
This URL is also the listing page for Cisco DCNM for LAN product documentation.
- Cisco Fabric Manager documentation for software releases earlier than Cisco DCNM Release 5.2, retains the name Cisco Fabric Manager and remains available at its current Cisco.com listing page:
http://www.cisco.com/en/US/products/ps10495/tsd_products_support_configure.html
You should continue to use the Cisco Fabric Manager documentation if you are using a release of Cisco Fabric Manager software that is earlier than Cisco DCNM Release 5.2.
- The name DCNM-SAN is used in place of Cisco DCNM for SAN in the user interface of Cisco Data Center Network Manager; likewise, the name DCNM-LAN is used in place of Cisco DCNM for LAN in the user interface. To match the user interface, the product documentation also uses the names DCNM-SAN and DCNM-LAN.
- The following new publications support both Cisco DCNM for LAN and DCNM for SAN, and address the new licensing model, the new installation process, and the new features of Cisco DCNM:
 - *Cisco DCNM Installation and Licensing Guide*
 - *Cisco DCNM Release Notes*
- For a complete list of Cisco DCNM documentation, see the “Related Documentation” section in the Preface.

As of Cisco MDS NX-OS Release 4.2(1), software configuration information is available in new feature-specific configuration guides for the following information:

- System management
- Interfaces
- Fabric

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Quality of service
- Security
- IP services
- High availability and redundancy

The information in these new guides previously existed in the *Cisco MDS 9000 Family CLI Configuration Guide* and in the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*. Those configuration guides remain available on Cisco.com and should be used for all software releases prior to MDS NX-OS Release 4.2(1). Each guide addresses the features introduced in or available in a particular release. Select and view the configuration guide that pertains to the software installed in your switch.

For a complete list of document titles, see the list of Related Documentation in the “Preface.”

To find additional information about Cisco MDS NX-OS Release 4.2(x), see the *Cisco MDS 9000 Family Release Notes* available at the following Cisco Systems website:

http://www.cisco.com/en/US/products/ps5989/prod_release_notes_list.htm

About This Guide

The information in the new *Fabric Configuration Guide, Cisco DCNM for SAN* previously existed in Part 4: Fabric of the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*.

Table 1-1 lists the New and Changed features for this guide, starting with MDS NX-OS Release 4.2(1).

Table 1-1 New and Changed Features

Feature	GUI Change	Description	Changed in Release	Where Documented
Smart Zoning	Zone and Zoneset configuration windows	Smart zoning supports zoning among more devices by reducing the number of zoning entries that needs to be programmed by considering device type information without increasing the size of the zone set.	6.1(1)	Chapter 1, “Configuring and Managing Zones.”
FCoE Configuration	FC Services > FCoE	Added information about discovering Cisco Nexus 7000 and Cisco MDS 9000 family switches using the FCoE wizard.	5.2(1)	Chapter 1, “Configuring FCoE”
Host Provision Wizard	Host Provision Wizard	Added information about Host Provision Wizard.	4.2(1)	Chapter 1, “Configuring and Managing VSANs”
Zones and Zonesets	Zone and Zoneset configuration windows	Added information about adding multiple end devices to zones and multiple zones to zone sets.	4.2(1)	Chapter 1, “Configuring and Managing Zones”
Device Alias	Interfaces configuration window	Added information about populating device alias to interface description.	4.2(1)	Chapter 1, “Distributing Device Alias Services”

Send documentation comments to

Table 1-1 New and Changed Features (continued)

Feature	GUI Change	Description	Changed in Release	Where Documented
X2 DWDM	Module configuration window	Added information about configuring X2 DWDM transceiver frequency.	4.2(1)	Chapter 1, “Configuring Dense Wavelength Division Multiplexing”
XRC Acceleration	Interfaces configuration window	Added information about configuring XRC Acceleration.	4.2(1)	Chapter 1, “Configuring FICON”

For a complete list of Cisco DCNM documentation, see the “Related Documentation” in the Preface.

Send documentation comments to dcnm-san-docfeedback@cisco.com



Preface

This preface describes the audience, organization, and conventions of the *Fabric Configuration Guide, Cisco DCNM for SAN*. It also provides information on how to obtain related documentation.

Audience

This guide is for experienced network administrators who are responsible for configuring and maintaining the Cisco MDS 9000 Family of multilayer directors and fabric switches.

Organization

The *Fabric Configuration Guide, Cisco DCNM for SAN* is organized as follows:

Chapter	Title	Description
Chapter 1	Fabric Overview	Provides an overview of features described in this guide.
Chapter 1	Configuring and Managing VSANs	Describes how virtual SANs (VSANs) work, explains the concept of default VSANs, isolated VSANs, VSAN IDs, and attributes, and provides details on how to create, delete, and view VSANs.
Chapter 1	Configuring SAN Device Virtualization	Describes how to configure virtual devices to represent physical end devices for switches running Cisco MDS SAN-OS Release 3.1(2) and NX-OS Release 4.1(1a).
Chapter 1	Creating Dynamic VSANs	Defines the Dynamic Port VSAN Membership (DPVM) feature that is used to maintain fabric topology when a host or storage device connection is moved between two Cisco MDS Family switches.
Chapter 1	Configuring and Managing Zones	Defines various zoning concepts and provides details on configuring a zone set and zone management features.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Chapter	Title	Description
Chapter 1	Distributing Device Alias Services	Describes the use of the Distributed Device Alias Services (device alias) to distribute device alias names on a fabric-wide basis.
Chapter 1	Configuring FCoE	Describes how to configure Fibre Channel over Ethernet (FCoE) on a Cisco Nexus 5000 Series, Nexus 7000 Series, and MDS 9000 Family switch.
Chapter 1	Configuring Fibre Channel Routing Services and Protocols	Provides details and configuration information on Fibre Channel routing services and protocols.
Chapter 1	Configuring Dense Wavelength Division Multiplexing	Dense Wavelength-Division Multiplexing (DWDM) multiplexes multiple optical carrier signals on a single optical fiber. DWDM uses different wavelengths to carry various signals.
Chapter 1	Managing FLOGI, Name Server, FDMI, and RSCN Databases	Provides name server and fabric login details required to manage storage devices and display registered state change notification (RSCN) databases.
Chapter 1	Discovering SCSI Targets	Describes how the SCSI LUN discovery feature is started and displayed.
Chapter 1	Configuring FICON	Provides details on the Fibre Connection (FICON) interface, fabric binding, and the Registered Link Incident Report (RLIR) capabilities in Cisco MDS 9000 Family switches.
Chapter 1	Configuring Advanced Fabric Features	Describes the advanced configuration features—time out values, fctrace, fabric analyzer, world wide names, flat FC IDs, loop monitoring, and interoperating switches.

Document Conventions

Command descriptions use these conventions:

boldface font	Commands and keywords are in boldface.
<i>italic font</i>	Arguments for which you supply values are in italics.
[]	Elements in square brackets are optional.
[x y z]	Optional alternative keywords are grouped in brackets and separated by vertical bars.

Screen examples use these conventions:

<code>screen font</code>	Terminal sessions and information the switch displays are in screen font.
<code>boldface screen font</code>	Information you must enter is in boldface screen font.
<i><code>italic screen font</code></i>	Arguments for which you supply values are in italic screen font.

Send documentation comments to dcnm-san-docfeedback@cisco.com

< >	Nonprinting characters, such as passwords, are in angle brackets.
[]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

This document uses the following conventions:



Note

Means reader *take note*. Notes contain helpful suggestions or references to material not covered in the manual.



Caution

Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

Related Documentation

The documentation set for the Cisco MDS 9000 Family includes the following documents. To find a document online, use the Cisco MDS NX-OS Documentation Locator at:

http://www.cisco.com/en/US/docs/storage/san_switches/mds9000/roadmaps/doclocator.htm

Release Notes

- *Cisco MDS 9000 Family Release Notes for Cisco MDS NX-OS Releases*
- *Cisco MDS 9000 Family Release Notes for MDS SAN-OS Releases*
- *Cisco MDS 9000 Family Release Notes for Cisco MDS 9000 EPLD Images*
- *Cisco DCNM Release Notes*

Regulatory Compliance and Safety Information

- *Regulatory Compliance and Safety Information for the Cisco MDS 9000 Family*

Compatibility Information

- *Cisco Data Center Interoperability Support Matrix*
- *Cisco MDS 9000 NX-OS Hardware and Software Compatibility Information and Feature Lists*
- *Cisco MDS 9000 Family Switch-to-Switch Interoperability Configuration Guide*

Hardware Installation

- *Cisco MDS 9500 Series Hardware Installation Guide*

Send documentation comments to dcnm-san-docfeedback@cisco.com

- *Cisco MDS 9200 Series Hardware Installation Guide*
- *Cisco MDS 9100 Series Hardware Installation Guide*
- *Cisco MDS 9124 and Cisco MDS 9134 Multilayer Fabric Switch Quick Start Guide*

Software Installation and Upgrade

- *Cisco MDS 9000 NX-OS Software Upgrade and Downgrade Guide*

Cisco NX-OS

- *Cisco MDS 9000 Family NX-OS Licensing Guide*
- *Cisco MDS 9000 Family NX-OS Fundamentals Configuration Guide*
- *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*
- *Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide*
- *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide*
- *Cisco MDS 9000 Family NX-OS Quality of Service Configuration Guide*
- *Cisco MDS 9000 Family NX-OS Security Configuration Guide*
- *Cisco MDS 9000 Family NX-OS IP Services Configuration Guide*
- *Cisco MDS 9000 Family NX-OS Intelligent Storage Services Configuration Guide*
- *Cisco MDS 9000 Family NX-OS High Availability and Redundancy Configuration Guide*
- *Cisco MDS 9000 Family NX-OS Inter-VSAN Routing Configuration Guide*
- *Cisco MDS 9000 Family Cookbook for Cisco MDS SAN-OS*

Cisco DCNM-SAN

- *Cisco DCNM Fundamentals Guide, Release 6.x*
- *System Management Configuration Guide, Cisco DCNM for SAN, Release 6.x*
- *Interfaces Configuration Guide, Cisco DCNM for SAN, Release 6.x*
- *Fabric Configuration Guide, Cisco DCNM for SAN, Release 6.x*
- *Quality of Service Configuration Guide, Cisco DCNM for SAN, Release 6.x*
- *Security Configuration Guide, Cisco DCNM for SAN, Release 6.x*
- *IP Services Configuration Guide, Cisco DCNM for SAN, Release 6.x*
- *Intelligent Storage Services Configuration Guide, Cisco DCNM for SAN, Release 6.x*
- *High Availability and Redundancy Configuration Guide, Cisco DCNM for SAN, Release 6.x*
- *Inter-VSAN Routing Configuration Guide, Cisco DCNM for SAN, Release 6.x*
- *SMI-S and Web Services Programming Guide, Cisco DCNM for SAN, Release 5.x*

Send documentation comments to dcnm-san-docfeedback@cisco.com

Command-Line Interface

- *Cisco MDS 9000 Family Command Reference*

Intelligent Storage Networking Services Configuration Guides

- *Cisco MDS 9000 Family I/O Acceleration Configuration Guide*
- *Cisco MDS 9000 Family SANTap Deployment Guide*
- *Cisco MDS 9000 Family Data Mobility Manager Configuration Guide*
- *Cisco MDS 9000 Family Storage Media Encryption Configuration Guide*

Troubleshooting and Reference

- *Cisco MDS 9000 Family and Nexus 7000 Series System Messages Reference*
- *Cisco MDS 9000 Family SAN-OS Troubleshooting Guide*
- *Cisco MDS 9000 Family NX-OS MIB Quick Reference*
- *Cisco DCNM for SAN Database Schema Reference*

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation, at:

<http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html>

- Subscribe to the *What's New in Cisco Product Documentation* as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS version 2.0.

Send documentation comments to dcnm-san-docfeedback@cisco.com



CHAPTER 1

Fabric Overview

The Cisco MDS 9000 Family NX-OS command-line interface (CLI) can configure and manage features such as VSANs, SAN device virtualization, dynamic VSANs, zones, distributed device alias services, Fibre Channel routing services and protocols, FLOGI, name server, FDMI, RSCN database, SCSI targets, FICON, and other advanced features.

This chapter describes some of these features and includes the following topics:

- [Virtual SANs, page 1-1](#)
- [Dynamic Port VSAN Membership, page 1-2](#)
- [SAN Device Virtualization, page 1-2](#)
- [Zoning, page 1-2](#)
- [Distributed Device Alias Services, page 1-3](#)
- [Fibre Channel Routing Services and Protocols, page 1-3](#)
- [Multiprotocol Support, page 1-3](#)

Virtual SANs

Virtual SAN (VSAN) technology partitions a single physical SAN into multiple VSANs. VSAN capabilities allow Cisco NX-OS software to logically divide a large physical fabric into separate, isolated environments to improve Fibre Channel SAN scalability, availability, manageability, and network security. For FICON, VSANs facilitate hardware-based separation of FICON and open systems.

Each VSAN is a logically and functionally separate SAN with its own set of Fibre Channel fabric services. This partitioning of fabric services greatly reduces network instability by containing fabric reconfigurations and error conditions within an individual VSAN. The strict traffic segregation provided by VSANs helps ensure that the control and data traffic of a specified VSAN are confined within the VSAN's own domain, increasing SAN security. VSANs help reduce costs by facilitating consolidation of isolated SAN islands into a common infrastructure without compromising availability.

Users can create administrator roles that are limited in scope to certain VSANs. For example, a network administrator role can be set up to allow configuration of all platform-specific capabilities, while other roles can be set up to allow configuration and management only within specific VSANs. This approach improves the manageability of large SANs and reduces disruptions due to human error by isolating the effect of a user action to a specific VSAN whose membership can be assigned based on switch ports or the worldwide name (WWN) of attached devices.

Send documentation comments to dcnm-san-docfeedback@cisco.com

VSANs are supported across FCIP links between SANs, which extends VSANs to include devices at a remote location. The Cisco MDS 9000 Family switches also implement trunking for VSANs. Trunking allows Inter-Switch Links (ISLs) to carry traffic for multiple VSANs on the same physical link.

Dynamic Port VSAN Membership

Port VSAN membership on the switch is assigned on a port-by-port basis. By default each port belongs to the default VSAN. You can dynamically assign VSAN membership to ports by assigning VSANs based on the device WWN. This method is referred to as Dynamic Port VSAN Membership (DPVM). DPVM offers flexibility and eliminates the need to reconfigure the port VSAN membership to maintain fabric topology when a host or storage device connection is moved between two Cisco MDS 9000 family switches or two ports within a switch. DPVM retains the configured VSAN regardless of where a device is connected or moved.

SAN Device Virtualization

Cisco SAN device virtualization (SDV) allows virtual devices representing physical end devices to be used for SAN configuration. Virtualization of SAN devices significantly reduces the time needed to swap out hardware. For example, if a storage array was replaced without using SDV, server downtime would be required for SAN zoning changes and host operating system configuration updates. With SDV, only the mapping between virtual and physical devices needs to change after hardware is swapped, insulating the SAN and end devices from extensive configuration changes.

Zoning

Zoning provides access control for devices within a SAN. Cisco NX-OS software supports the following types of zoning:

- N port zoning—Defines zone members based on the end-device (host and storage) port.
 - WWN
 - Fibre Channel identifier (FC-ID)
- Fx port zoning—Defines zone members based on the switch port.
 - WWN
 - WWN plus interface index, or domain ID plus interface index
- Domain ID and port number (for Brocade interoperability)
- iSCSI zoning—Defines zone members based on the host zone.
 - iSCSI name
 - IP address
- LUN zoning—When combined with N port zoning, LUN zoning helps ensure that LUNs are accessible only by specific hosts, providing a single point of control for managing heterogeneous storage-subsystem access.
- Read-only zones—An attribute can be set to restrict I/O operations in any zone type to SCSI read-only commands. This feature is especially useful for sharing volumes across servers for backup, data warehousing, etc.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Broadcast zones—An attribute can be set for any zone type to restrict broadcast frames to members of the specific zone.

To provide strict network security, zoning is always enforced per frame using access control lists (ACLs) that are applied at the ingress switch. All zoning policies are enforced in hardware, and none of them cause performance degradation. Enhanced zoning session-management capabilities further enhance security by allowing only one user at a time to modify zones.

Distributed Device Alias Services

All switches in the Cisco MDS 9000 Family support Distributed Device Alias Services (device alias) on a per-VSAN basis and on a fabric-wide basis. Device alias distribution allows you to move host bus adapters (HBAs) between VSANs without manually reentering alias names.

Fibre Channel Routing Services and Protocols

Fabric Shortest Path First (FSPF) is the standard path selection protocol used by Fibre Channel fabrics. The FSPF feature is enabled by default on all Fibre Channel switches. You do not need to configure any FSPF services except in configurations that require special consideration. FSPF automatically calculates the best path between any two switches in a fabric. Specifically, FSPF is used to perform these functions:

- Dynamically compute routes throughout a fabric by establishing the shortest and quickest path between any two switches.
- Select an alternative path in the event of the failure of a given path. FSPF supports multiple paths and automatically computes an alternative path around a failed link. FSPF provides a preferred route when two equal paths are available.

Multiprotocol Support

In addition to supporting Fibre Channel Protocol (FCP), Cisco NX-OS software supports IBM Fibre Connection (FICON), Small Computer System Interface over IP (iSCSI), and Fibre Channel over IP (FCIP) in a single platform. Native iSCSI support in the Cisco MDS 9000 Family switches helps customers consolidate storage for a wide range of servers into a common pool on the SAN.

Send documentation comments to dcnm-san-docfeedback@cisco.com



CHAPTER 1

Configuring and Managing VSANs

You can achieve higher security and greater stability in Fibre Channel fabrics by using virtual SANs (VSANs) on Cisco MDS 9000 Family switches and Cisco Nexus 5000 Series switches. VSANs provide isolation among devices that are physically connected to the same fabric. With VSANs you can create multiple logical SANs over a common physical infrastructure. Each VSAN can contain up to 239 switches and has an independent address space that allows identical Fibre Channel IDs (FC IDs) to be used simultaneously in different VSANs. This chapter includes the following sections:

- [Information About VSANs, page 1-1](#)
- [Licensing Requirements for VSAN, page 1-9](#)
- [Default Settings, page 1-9](#)
- [Configuring VSANs, page 1-10](#)
- [Configuring Load Balancing, page 1-12](#)
- [Verifying VSAN Configuration, page 1-14](#)
- [Field Descriptions for VSAN, page 1-14](#)
- [Additional References, page 1-16](#)

Information About VSANs

A VSAN is a virtual storage area network (SAN). A SAN is a dedicated network that interconnects hosts and storage devices primarily to exchange SCSI traffic. In SANs, you use the physical links to make these interconnections. A set of protocols run over the SAN to handle routing, naming, and zoning. You can design multiple SANs with different topologies.

With the introduction of VSANs, the network administrator can build a single topology containing switches, links, and one or more VSANs. Each VSAN in this topology has the same behavior and property of a SAN. A VSAN has the following additional features:

- Multiple VSANs can share the same physical topology.
- The same Fibre Channel IDs (FC IDs) can be assigned to a host in another VSAN, thus increasing VSAN scalability.
- Every instance of a VSAN runs all required protocols such as FSPF, domain manager, and zoning.
- Fabric-related configurations in one VSAN do not affect the associated traffic in another VSAN.
- Events causing traffic disruptions in one VSAN are contained within that VSAN and are not propagated to other VSANs.

Send documentation comments to dcnm-san-docfeedback@cisco.com

This section describes VSANs and includes the following topics:

- [VSANs Topologies, page 1-2](#)
- [VSAN Advantages, page 1-4](#)
- [VSANs Versus Zones, page 1-5](#)
- [VSAN Configuration, page 1-6](#)
- [About VSAN Creation, page 1-7](#)
- [About Port VSAN Membership, page 1-7](#)
- [About the Default VSAN, page 1-7](#)
- [About the Isolated VSAN, page 1-7](#)
- [Operational State of a VSAN, page 1-8](#)
- [About Static VSAN Deletion, page 1-8](#)
- [About Load Balancing, page 1-9](#)
- [About Interop Mode, page 1-9](#)
- [About FICON VSANs, page 1-9](#)

VSANs Topologies

The switch icons shown in both [Figure 1-1](#) and [Figure 1-2](#) indicate that these features apply to any switch in the Cisco MDS 9000 Family.

[Figure 1-1](#) shows a fabric with three switches, one on each floor. The geographic location of the switches and the attached devices is independent of their segmentation into logical VSANs. No communication between VSANs is possible. Within each VSAN, all members can talk to one another.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 1-1 Logical VSAN Segmentation

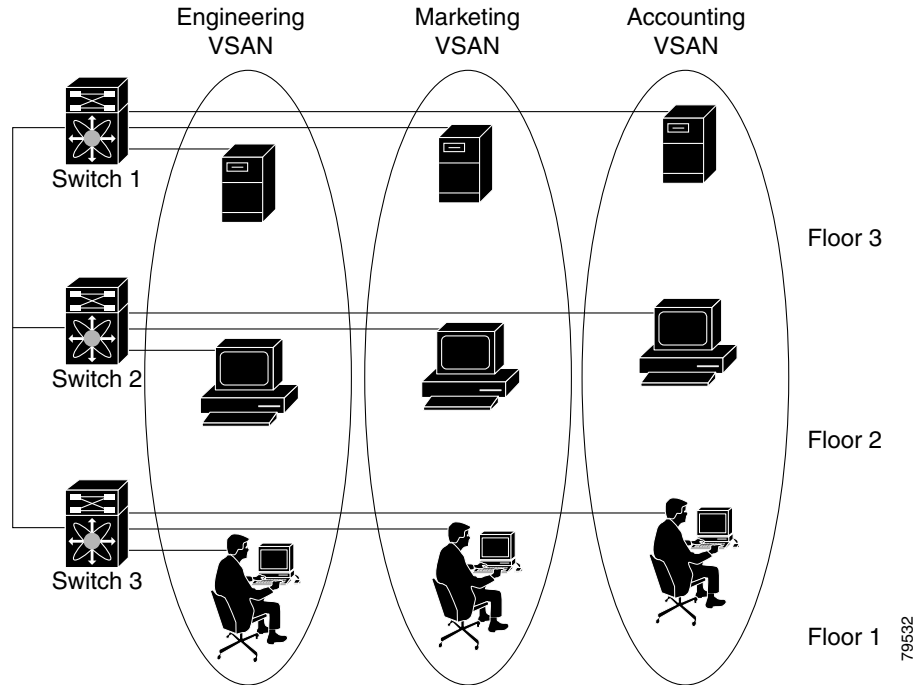



Figure 1-2 shows a physical Fibre Channel switching infrastructure with two defined VSANs: VSAN 2 (dashed) and VSAN 7 (solid). VSAN 2 includes hosts H1 and H2, application servers AS2 and AS3, and storage arrays SA1 and SA4. VSAN 7 connects H3, AS1, SA2, and SA3.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 1-2 Example of Two VSANs



The four switches in this network are interconnected by trunk links that carry both VSAN 2 and VSAN 7 traffic. The inter-switch topology of both VSAN 2 and VSAN 7 are identical. This is not a requirement and a network administrator can enable certain VSANs on certain links to create different VSAN topologies.

Without VSANs, a network administrator would need separate switches and links for separate SANs. By enabling VSANs, the same switches and links may be shared by multiple VSANs. VSANs allow SANs to be built on port granularity instead of switch granularity. [Figure 1-2](#) illustrates that a VSAN is a group of hosts or storage devices that communicate with each other using a virtual topology defined on the physical SAN.

The criteria for creating such groups differ based on the VSAN topology:

- VSANs can separate traffic based on the following requirements:
 - Different customers in storage provider data centers
 - Production or test in an enterprise network
 - Low and high security requirements
 - Backup traffic on separate VSANs
 - Replicating data from user traffic
- VSANs can meet the needs of a particular department or application.

VSAN Advantages

VSANs offer the following advantages:

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Traffic isolation—Traffic is contained within VSAN boundaries and devices reside only in one VSAN ensuring absolute separation between user groups, if desired.
- Scalability—VSANs are overlaid on top of a single physical fabric. The ability to create several logical VSAN layers increases the scalability of the SAN.
- Per VSAN fabric services—Replication of fabric services on a per VSAN basis provides increased scalability and availability.
- Redundancy—Several VSANs created on the same physical SAN ensure redundancy. If one VSAN fails, redundant protection (to another VSAN in the same physical SAN) is configured using a backup path between the host and the device.
- Ease of configuration—Users can be added, moved, or changed between VSANs without changing the physical structure of a SAN. Moving a device from one VSAN to another only requires configuration at the port level, not at a physical level.

Up to 256 VSANs can be configured in a switch. Of these, one is a default VSAN (VSAN 1), and another is an isolated VSAN (VSAN 4094). User-specified VSAN IDs range from 2 to 4093.

VSANs Versus Zones

You can define multiple zones in a VSAN. Because two VSANs are equivalent to two unconnected SANs, zone A on VSAN 1 is different and separate from zone A in VSAN 2. [Table 1-1](#) lists the differences between VSANs and zones.

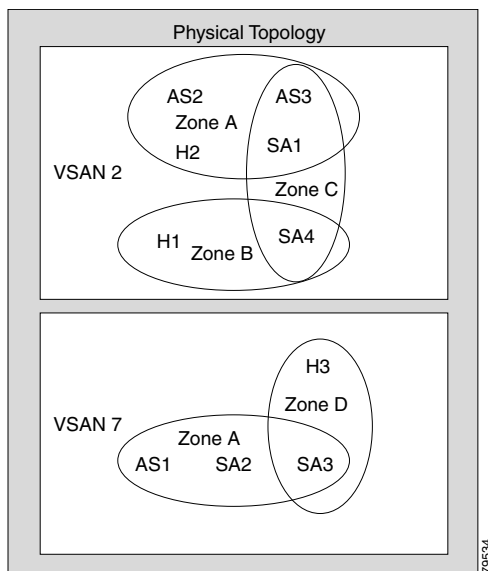
Table 1-1 VSAN and Zone Comparison

VSAN Characteristic	Zone Characteristic
VSANs equal SANs with routing, naming, and zoning protocols.	Routing, naming, and zoning protocols are not available on a per-zone basis.
—	Zones are always contained within a VSAN. Zones never span two VSANs.
VSANs limit unicast, multicast, and broadcast traffic.	Zones limit unicast traffic.
Membership is typically defined using the VSAN ID to Fx ports.	Membership is typically defined by the pWWN.
An HBA or a storage device can belong only to a single VSAN—the VSAN associated with the Fx port.	An HBA or storage device can belong to multiple zones.
VSANs enforce membership at each E port, source port, and destination port.	Zones enforce membership only at the source and destination ports.
VSANs are defined for larger environments (storage service providers).	Zones are defined for a set of initiators and targets not visible outside the zone.
VSANs encompass the entire fabric.	Zones are configured at the fabric edge.

[Figure 1-3](#) shows the possible relationships between VSANs and zones. In VSAN 2, three zones are defined: zone A, zone B, and zone C. Zone C overlaps both zone A and zone B as permitted by Fibre Channel standards. In VSAN 7, two zones are defined: zone A and zone D. No zone crosses the VSAN boundary—they are completely contained within the VSAN. Zone A defined in VSAN 2 is different and separate from zone A defined in VSAN 7.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 1-3 VSANS with Zoning



VSAN Configuration

VSANs have the following attributes:

- **VSAN ID**—The VSAN ID identifies the VSAN as the default VSAN (VSAN 1), user-defined VSANs (VSAN 2 to 4093), and the isolated VSAN (VSAN 4094).
- **State**—The administrative state of a VSAN can be configured to an active (default) or suspended state. Once VSANs are created, they may exist in various conditions or states.
 - The active state of a VSAN indicates that the VSAN is configured and enabled. By enabling a VSAN, you activate the services for that VSAN.
 - The suspended state of a VSAN indicates that the VSAN is configured but not enabled. If a port is configured in this VSAN, it is disabled. Use this state to deactivate a VSAN without losing the VSAN's configuration. All ports in a suspended VSAN are disabled. By suspending a VSAN, you can preconfigure all the VSAN parameters for the whole fabric and activate the VSAN immediately.
- **VSAN name**—This text string identifies the VSAN for management purposes. The name can be from 1 to 32 characters long and it must be unique across all VSANs. By default, the VSAN name is a concatenation of VSAN and a four-digit string representing the VSAN ID. For example, the default name for VSAN 3 is VSAN0003.



Note A VSAN name must be unique.

- **Load balancing attributes**—These attributes indicate the use of the source-destination ID (src-dst-id) or the originator exchange OX ID (src-dst-ox-id, the default) for load balancing path selection.

Send documentation comments to dcnm-san-docfeedback@cisco.com

**Note**

OX ID based load balancing of IVR traffic from IVR-enabled switches is not supported on Generation 1 switching modules. OX ID based load balancing of IVR traffic from a non-IVR MDS 9000 Family switch should work. Generation 2 switching modules support OX ID based load balancing of IVR traffic from IVR-enabled switches.

About VSAN Creation

A VSAN is in the operational state if the VSAN is active and at least one port is up. This state indicates that traffic can pass through this VSAN. This state cannot be configured.

About Port VSAN Membership

Port VSAN membership on the switch is assigned on a port-by-port basis. By default, each port belongs to the default VSAN. You can assign VSAN membership to ports using one of two methods:

- Statically—By assigning VSANs to ports.

See the [“Assigning Static Port VSAN Membership” section on page 1-11](#).

- Dynamically—By assigning VSANs based on the device WWN. This method is referred to as dynamic port VSAN membership (DPVM).

See [Chapter 1, “Creating Dynamic VSANs.”](#)

Trunking ports have an associated list of VSANs that are part of an allowed list (refer to the *Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide*).

About the Default VSAN

The factory settings for switches in the Cisco MDS 9000 Family have only the default VSAN 1 enabled. We recommend that you do not use VSAN 1 as your production environment VSAN. If no VSANs are configured, all devices in the fabric are considered part of the default VSAN. By default, all ports are assigned to the default VSAN.

**Note**

VSAN 1 cannot be deleted, but it can be suspended.

**Note**

Up to 256 VSANs can be configured in a switch. Of these, one is a default VSAN (VSAN 1), and another is an isolated VSAN (VSAN 4094). User-specified VSAN IDs range from 2 to 4093.

About the Isolated VSAN

VSAN 4094 is an isolated VSAN. All non-trunking ports are transferred to this VSAN when the VSAN to which they belong is deleted. This avoids an implicit transfer of ports to the default VSAN or to another configured VSAN. All ports in the deleted VSAN are isolated (disabled).

Send documentation comments to dcnm-san-docfeedback@cisco.com

**Note**

When you configure a port in VSAN 4094 or move a port to VSAN 4094, that port is immediately isolated.

**Caution**

Do not use an isolated VSAN to configure ports.

**Note**

Up to 256 VSANs can be configured in a switch. Of these, one is a default VSAN (VSAN 1), and another is an isolated VSAN (VSAN 4094). User-specified VSAN IDs range from 2 to 4093.

Operational State of a VSAN

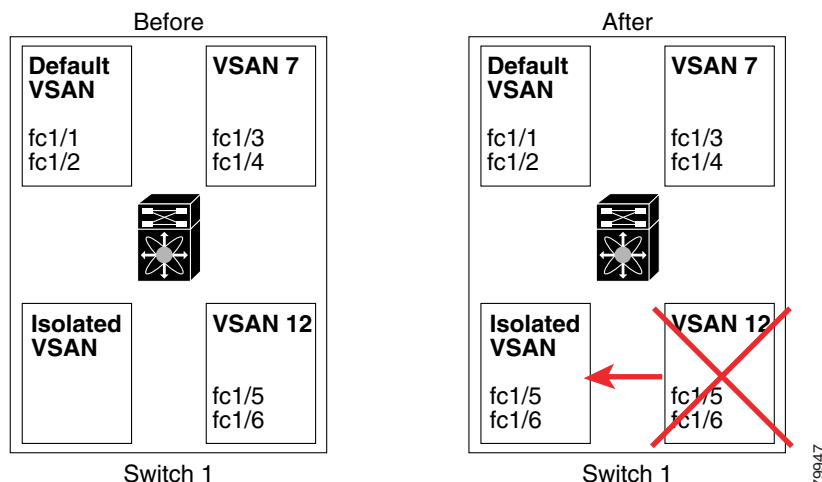
A VSAN is in the operational state if the VSAN is active and at least one port is up. This state indicates that traffic can pass through this VSAN. This state cannot be configured.

About Static VSAN Deletion

When an active VSAN is deleted, all of its attributes are removed from the running configuration. VSAN-related information is maintained by the system software as follows:

- VSAN attributes and port membership details are maintained by the VSAN manager. This feature is affected when you delete a VSAN from the configuration. When a VSAN is deleted, all the ports in that VSAN are made inactive and the ports are moved to the isolated VSAN. If the same VSAN is recreated, the ports do not automatically get assigned to that VSAN. You must explicitly reconfigure the port VSAN membership (see [Figure 1-4](#)).

Figure 1-4 VSAN Port Membership Details



- VSAN-based runtime (name server), zoning, and configuration (static routes) information is removed when the VSAN is deleted.
- Configured VSAN interface information is removed when the VSAN is deleted.

Send documentation comments to dcnm-san-docfeedback@cisco.com

**Note**

The allowed VSAN list is not affected when a VSAN is deleted (refer to the *Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide*).

Any commands for a nonconfigured VSAN are rejected. For example, if VSAN 10 is not configured in the system, then a command request to move a port to VSAN 10 is rejected.

About Load Balancing

Load balancing attributes indicate the use of the source-destination ID (src-dst-id) or the originator exchange OX ID (src-dst-ox-id, the default) for load balancing path selection.

About Interop Mode

Interoperability enables the products of multiple vendors to come into contact with each other. Fibre Channel standards guide vendors towards common external Fibre Channel interfaces. See the [“Switch Interoperability” section on page 1-5](#).

About FICON VSANs

You can enable FICON in up to eight VSANs. See the [“FICON VSAN Prerequisites” section on page 1-7](#).

Host Provisioning Wizard

The Host Provisioning wizard provides an intuitive way to commission a new host or decommission an existing host without requiring the use of multiple tools and features. The wizard allows you to create a device alias, and configure DPVM, zoning, and flow creation.

Licensing Requirements for VSAN

The following table shows the licensing requirements for this feature:

License	License Description
ENTERPRISE _PKG	The enterprise license is required to enable VSAN. For a complete explanation of the licensing scheme, see the <i>Cisco MDS 9000 Family NX-OS Licensing Guide</i> .

Default Settings

[Table 1-2](#) lists the default settings for all configured VSANs.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Table 1-2 **Default VSAN Parameters**

Parameters	Default
Default VSAN	VSAN 1.
State	Active state.
Name	Concatenation of VSAN and a four-digit string representing the VSAN ID. For example, VSAN 3 is VSAN0003.
Load-balancing attribute	OX ID (src-dst-ox-id).

Configuring VSANs

This section includes the following topics:

- [Creating VSANs, page 1-10](#)
- [Assigning Static Port VSAN Membership, page 1-11](#)
- [Deleting Static VSANs, page 1-11](#)
- [Commissioning a Host, page 1-12](#)
- [Decommissioning a Host, page 1-13](#)

Creating VSANs

Restrictions


You cannot configure any application-specific parameters for a VSAN before creating the VSAN.

Detailed Steps

To create and configure VSANs, follow these steps:

- Step 1** Click the **Create VSAN** icon (see [Figure 1-5](#)).

Figure 1-5 **Create VSAN Icon**



Note As of Cisco SAN-OS Release 3.1(2) and later, if you check the Static Domain IDs check box, DCNM-SAN creates the VSAN in suspended mode and then automatically activates the VSAN.

- Step 2** Check the switches that you want in this VSAN.
- Step 3** Fill in the VSAN Name and VSAN ID fields.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Step 4** Set the **LoadBalancing** value and the **InterOperValue**.
 - Step 5** Set the Admin State to **active** or **suspended**.
 - Step 6** Check the **Static Domain Ids** check box to assign an unused static domain ID to the VSAN.
 - Step 7** (Optional) Select the **FICON** and **Enable Fabric Binding for Selected Switches** options if you want these features enabled.

See the “[Configuring FICON](#)” section on page 1-20 and refer to the *Cisco MDS 9000 Family NX-OS Security Configuration Guide* for details.
 - Step 8** Complete the fields in this dialog box and click **Create** to add the VSAN or click **Close**.
-

Assigning Static Port VSAN Membership

Detailed Steps

To statically assign VSAN membership for an interface, follow these steps:

- Step 1** Choose **FC Interfaces > Physical** from the Physical Attributes pane. You see the interface configuration in the Information pane.
 - Step 2** Click the **General** tab.

You see the Fibre Channel general physical information. Double-click and complete the PortVSAN field.
 - Step 3** Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.
-

Deleting Static VSANs

Detailed Steps

To delete a VSAN and its attributes, follow these steps:

- Step 1** Select **All VSANs** from the Logical Domains pane.

The VSANs in the fabric are listed in the Information pane.
 - Step 2** Right-click the VSAN that you want to delete and select **Delete Row** from the drop-down menu.

You see a confirmation dialog box.
 - Step 3** Click **Yes** to confirm the deletion or **No** to close the dialog box without deleting the VSAN.
-

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuring Load Balancing

Detailed Steps

To configure load balancing on an existing VSAN, follow these steps:

-
- Step 1** Choose **Fabricxx > All VSANs** from the Logical Domains pane.
You see the VSAN configuration in the Information pane.
 - Step 2** Select a VSAN and complete the LoadBalancing field.
 - Step 3** Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.
-


Commissioning a Host

Detailed Steps

To commission a new host, follow these steps:

-
- Step 1** From the DCNM-SAN window, select **Tools > Host Provisioning**.
The Host Provisioning wizard window is displayed.
 - Step 2** Click the **Commission** radio button.
 - Step 3** Click [...] and select the host from the existing configurations or VSAN (see [Figure 1-6](#)), or enter the WWN of a host that is not in VSAN or not configured yet.

Figure 1-6 *Select a Host*



If the host configuration already exists, the switch, device alias, and VSAN information are populated in the window.

If the configuration does not exist already, enter a device alias for the WWN, enter a switch where the configuration will be initiated, and select a VSAN to which the host should belong. The entries are created and saved when you click Next in the Host Provisioning wizard window.

- Step 4** Uncheck the **Skip Zoning** check box.
- Step 5** Click **Next**. The Select Targets and the Select Zone windows appear.
- Step 6** Uncheck the **Skip DPVM** check box.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 7 Click **Next**. The DPVM entries are created.

Step 8 Click **Next**.

The Select Targets window appears.



Note The Host Provisioning wizard requires that basic and enhanced device alias, DPVM, and CFS to be enabled in all switches in the selected VSAN.

Step 9 Select the target with which the host needs to communicate, and click **Add**.

The target entry is moved to the bottom of the window.

Step 10 Click **Next**.

The Select Zone window appears.

Step 11 Select a zone and check the **Create Flow after Activation** check box.

The host and storage are added to a zone and the zone is activated, and a flow between host and storage is created when you click Finish.

Step 12 Click **Finish**.

The device alias and DPVM entries are created, a zone is created and activated, and the flow is created based on the check boxes you checked.

Decommissioning a Host

Detailed Steps

To decommission an existing host, follow these steps:

Step 1 From the DCNM-SAN window, select **Tools > Host Provisioning**.

The Select Host window appears.

Step 2 Click the **Decommission** radio button.

Step 3 Click [...] and select the host from the existing configurations or VSAN, or enter the WWN of a host that is not in VSAN.

The device alias and DPVM state from all of the switches in the selected VSAN are populated if device alias with CFS and CFS DPVM are enabled and if the WWN is an eight-byte number.

Step 4 Click **Finish**. The device aliases are removed.

Step 5 Uncheck the **Skip Zoning** check box.

The WWN zone member is removed from all zones. If the zones without a WWN member become single member zones, these zones also are removed.

Step 6 Click **Finish**. If there is a local active zone set change due to the removal of zones, the appropriate zone set is activated.

Step 7 Uncheck the **Skip DPVM** check box.

Step 8 Click **Finish**. The DPVM entry is removed.

Step 9 Click **Next**.

Send documentation comments to dcnm-san-docfeedback@cisco.com

The Decommission Zones window appears.

Step 10 Check the **Remove Flow after Deactivation** check box.

The flow entry associated with the host is removed when you click Finish.

Step 11 Click **Finish**.

The device alias and DPVM entries are deleted, the zone is deactivated and deleted (if it has only one member after removing the host), and the flow is deleted depending on the check boxes you checked.

Verifying VSAN Configuration

- [Displaying Isolated VSAN Membership, page 1-14](#)

Displaying Isolated VSAN Membership

To display interfaces that exist in the isolated VSAN, follow these steps:

Step 1 Expand **Fabricxx**, and then select **All VSANs** in the Logical Domains pane.

You see the VSAN configuration in the Information pane.

Step 2 Click the **Isolated Interfaces** tab.

You see the interfaces that are in the isolated VSAN.

Field Descriptions for VSAN

The following are the field descriptions for VSAN.

VSAN General

Field	Description
Name	The name of the VSAN. Note that default value will be the string <code>VSANxxxx</code> where <code>xxxx</code> is value of <code>vsanIndex</code> expressed as 4 digits. For example, if <code>vsanIndex</code> is 23, the default value is <code>VSAN0023</code> .
Mtu	The MTU of the VSAN. Normally, this is 2112.
LoadBalancing	The type of load balancing used on this VSAN. <ul style="list-style-type: none"> • <code>srcdst</code>— use source and destination ID for path selection • <code>srcdst 0xld</code>— use source, destination, and exchange IDs

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
InterOp	The interoperability mode of the local switch on this VSAN. <ul style="list-style-type: none"> • standard • interop-1 • interop-2 • interop-3
AdminState	The state of this VSAN.
OperState	The operational state of the VSAN.
InOrderDelivery	The InorderDelivery guarantee flag of device. If true, then the inorder delivery is guaranteed. If false, it is not guaranteed.
DomainId	Specifies an insistent domain ID.
FICON	True if the VSAN is FICON-enabled.
Network Latency	Network latency of this switch on this VSAN. This is the time interval after which the frames are dropped if they are not delivered in the order they were transmitted.

VSAN Membership

Field	Description
Switch	Name of the switch
Ports	FC ports in VSAN
Channels	PortChannels in VSAN
FCIP	FCIP Interfaces in VSAN
iSCSI	iSCSI Interfaces in VSAN
FICON	Interfaces in VSAN by FICON
FC Virtual Interface	Virtual FC interfaces in VSAN

VSAN Interop-4 WWN

Field	Description
VSAN ID	The ID of the VSAN containing the McData switch.
WWN	The WWN of the McData switch.

Send documentation comments to dcnm-san-docfeedback@cisco.com

VSAN Timers

Field	Description
VSAN Id	The ID of the VSAN.
R_A_TOV	The Resource_Allocation_Timeout Value used for FxPorts as the timeout value for determining when to reuse an NxPort resource such as a Recovery_Qualifier. It represents E_D_TOV plus twice the maximum time that a frame may be delayed within the fabric and still be delivered. Note that all switches in a fabric should be configured with the same value of this timeout.
D_S_TOV	The Distributed_Services_Timeout Value which indicates that how long a distributed services requestor will wait for a response.
E_D_TOV	The Error_Detect_Timeout Value used for FxPorts as the timeout value for detecting an error condition. Note that all switches in a fabric should be configured with the same value of this timeout. Note that value must be less than value of D_S_TOV.
NetworkDropLatency	Network latency of this switch on this VSAN.

VSAN Default Zone Policies

Field	Description
Zone Behavior	Represents the initial value for default zone behavior on a VSAN when it is created. If a VSAN were to be deleted and re-created again, the default zone behavior will be set to the value specified for this object.
Propagation Mode	Represents the initial value for zone set propagation mode on a VSAN when it is created. If a VSAN were to be deleted and re-created again, the zone set propagation mode will be set to the value specified for this object.

Additional References

For additional information related to implementing VSANs, see the following section:

- [Related Document, page 1-17](#)
- [Standards, page 1-17](#)
- [RFCs, page 1-17](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- [MIBs, page 1-17](#)

Related Document

Related Topic	Document Title
Cisco MDS 9000 Family Command Reference	<i>Cisco MDS 9000 Family Command Reference</i>

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified.	—

MIBs

MIBs	MIBs Link
<ul style="list-style-type: none">• CISCO-VSAN-MIB• CISCO-VSAN-CAPABILITY	To locate and download MIBs, go to the following URL: http://www.cisco.com/en/US/products/ps5989/prod_technical_reference_list.html

Send documentation comments to dcnm-san-docfeedback@cisco.com



CHAPTER 1

Configuring SAN Device Virtualization

This chapter describes how to configure virtual devices to represent physical end devices for switches running Cisco MDS SAN-OS Release 3.1(2) and later, or NX-OS Release 4.1(1a) and later.

Cisco SAN device virtualization (SDV) is a licensed feature included in the Cisco MDS 9000 Family Enterprise package (ENTERPRISE_PKG). Refer to the *Cisco NX-OS Family Licensing Guide* for details about acquiring licenses.

This chapter includes the following topics:

- [Information About SDV, page 1-1](#)
- [Licensing Requirements for SAN Device Virtualization, page 1-5](#)
- [Guidelines and Limitations, page 1-5](#)
- [Default Settings, page 1-7](#)
- [Configuring SDV, page 1-7](#)
- [Field Descriptions for SDV, page 1-12](#)
- [Additional References, page 1-12](#)

Information About SDV

As of Cisco SAN-OS Release 3.1(2) and later, you can use Cisco SAN device virtualization to create virtual devices that represent physical end-devices. Virtualization of SAN devices accelerates swayout or failover to a replacement disk array, and it also minimizes downtime when replacing host bus adapters (HBAs) or when rehosting an application on a different server.


SAN device virtualization enables you to:

- Reduce the amount of time it takes for data migration, and ultimately the overall amount of downtime.
- Improve ease-of-use and reduce the possibility of user-introduced errors during the failover by performing the operation in a single step.
- Easily scale to larger numbers of targets.

SAN devices that are virtualized can be either initiators or targets. You can virtualize targets to create a *virtual target* and also virtualize initiators to create a *virtual initiator*. SAN device configurations do not distinguish between virtual initiators and virtual targets (see [Figure 1-1](#) and [Figure 1-2](#)).


Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 1-1 Target Virtualization



183017

Figure 1-2 Initiator Virtualization



183018




Note

While most of the examples in this chapter describe target virtualization, the initiator virtualization functions similarly.

Typically, today's deployments for handling device failures are designed for high availability (HA), with redundancy being a key part of this design. Consider the situation where a target is designed to be redundant. Two arrays are deployed—a primary and secondary in this situation. Enterprises often use some type of consistency technology (such as EMF SRDF) between the primary and secondary arrays to ensure that the secondary is a mirrored copy of the production LUN. However, if the primary array fails, it must be replaced by the secondary because all I/O must occur on the secondary array. Problems can occur because the time required to bring the secondary array up and have it working often takes longer than most can afford (Figure 1-3 illustrates this dilemma).

Figure 1-3 Typical Deployment for Handling Device Failures Before SDV



182097

Send documentation comments to dcnm-san-docfeedback@cisco.com

If a storage array is replaced *without* using Cisco SDV, then it may require the following actions:

- Taking down a server to modify zoning and account for the new array.
- Changing the Cisco NX-OS configuration to accommodate Fibre Channel IDs (FC IDs) and pWWNs of the new array.
- Changing a server configuration to accommodate the new FC IDs and pWWNs.

More specifically, without SDV you might experience the following conditions:

- It can take a considerable amount of time to configure a secondary device for a typical production environment.
- In the zoning configuration, all the initiators must be rezoned with the secondary device, and certain initiators must also be reconfigured. For example, the WWN and FC ID of the secondary device are different, so driver files must be changed and the server must be rebooted.
- Clustering (multiple initiators) compounds the problem, and the failover procedure must be repeated for each server of the cluster. Think of a server cluster as a set of HBAs—any storage array FC ID changes must be performed for each HBA.

SDV enables you to achieve the following performance targets:

- Reduce the amount of time it takes for data migration, and ultimately the overall amount of downtime.
- Easily scale to larger numbers of devices.


Figure 1-4 illustrates the benefits of SDV. In this configuration, disk array Y replaces disk array X. When disk array X was deployed, the user created virtual devices for all the Fibre Channel interfaces using SDV. After data replication from disk array X was completed, the user briefly pauses activity on the application server and relinked disk array Y to the virtual devices used by the server, completing the swapout of disk array X. No zoning changes or host operating system configuration changes were required during the time-critical period when the swap was performed; this significantly minimized application downtime.



Note

The array administrator will likely have to perform actions on array Y for it to become a primary device and accept server logins before linking the virtual device to the array Y pWWN.

Figure 1-4 SDV Example



Send documentation comments to dcnm-san-docfeedback@cisco.com

This section includes the following topics:

- [Key Concepts, page 1-4](#)
- [Automatic Failover and Fallback, page 1-4](#)
- [Resolving Fabric Merge Conflicts, page 1-4](#)

Key Concepts

The following terms are used throughout this chapter:

- **Virtual device**—The virtualized or proxy representation of the real device, which is registered with the name server and has a pWWN and FC ID. A virtual device exists as long as its real (physical) counterpart is online. The virtual device pWWN and FC ID must be unique and cannot clash with any real device pWWNs and FC IDs.
- **Virtual domain**—Reserved by SDV to assign FC IDs to virtual devices. If the switch that reserved the domain goes down, another switch takes over its role using the same domain.
- **Primary device**—The device that is configured as primary. By default, the primary device becomes the active device if it is online.
- **Secondary device**—The additional device that is configured. By default, the secondary device is standby.
- **Active device**—The device that is currently virtualized is called the active device. By default, the primary device becomes the active device if it is online. The active device is indicated by a (*) symbol.

Automatic Failover and Fallback

As of Cisco MDS NX-OS Release 4.1(1a), SAN device virtualization supports automatic failover and fallback configurations for the virtual devices. In all of the earlier releases, when there was a failure, you needed to manually configure the device as primary to make it active. With the introduction of automatic failover and fallback configurations, the active device is distinguished from the primary device indicated by a (*) symbol.

- **Auto failover**—When there is a failure, the failover auto attribute automatically shuts down the primary device and brings up the secondary device to active state. When the primary device comes back online, it requires user intervention to switchover.
- **Auto failover with fallback**—In addition to automatic failover, when the primary device comes back online after a failover, the primary device is brought to active state and the secondary device moved to standby state.

Resolving Fabric Merge Conflicts

Whenever two fabrics merge, SDV merges its database. A merge conflict can occur when there is a run-time information conflict or configuration mismatch. Run-time conflicts can occur due to:

- Identical pWWNs have been assigned to different virtual devices.
- The same virtual devices are assigned different pWWNs.
- The virtual device and virtual FC ID are mismatched.

Send documentation comments to dcnm-san-docfeedback@cisco.com

A *blank commit* is a commit operation that does not contain configuration changes, and enforces the SDV configuration of the committing switch fabric-wide. A blank commit operation resolves merge conflicts by pushing the configuration from the committing switch throughout the fabric, which reinitializes the conflicting virtual devices. Exercise caution while performing this operation, as it can easily take some virtual devices offline.

Merge failures resulting from a pWWN conflict can cause a failure with the device alias as well. A blank commit operation on a merge-failed VSAN within SDV should resolve the merge failure in the device alias.

You can avoid merge conflicts due to configuration mismatch by ensuring that:

- The pWWN and device alias entries for a virtual device are identical (in terms of primary and secondary).
- There are no virtual device name conflicts across VSANs in fabrics.

Licensing Requirements for SAN Device Virtualization

The following table shows the licensing requirements for this feature:

License	License Description
ENTERPRISE_PKG	The enterprise license is required for SAN device virtualization. For a complete explanation of the licensing scheme, see the <i>Cisco MDS 9000 Family NX-OS Licensing Guide</i> .

Guidelines and Limitations

As of MDS NX-OS Release 4.1(1a), the following conditions must be considered when configuring the virtual device failover attributes:

- The attribute configuration is supported only with MDS NX-OS Release 4.1(1a) and later. In a mixed mode fabric where earlier releases are combined, the attribute configuration will fail.
- When the failover attribute is configured, if the primary device is offline then the secondary device becomes active.
- When the failover attribute is deleted after the primary device failover to the secondary device, then the primary becomes active if the primary device is online. If the primary device is not online, then the SDV virtual device is shut down.



Note

The SDV attributes configuration is supported in Cisco DCNM for SAN Release 4.1(2) and later.

This section includes the guidelines and limitations for this feature:

- [SDV Requirements and Guidelines](#), page 1-6
- [Guidelines for Downgrading SDV](#), page 1-6
- [Downgrading with Attributes Configured](#), page 1-7
- [Downgrading with Virtual Initiators Configured](#), page 1-7
- [Downgrading with SDV LUN Zoning Configured](#), page 1-7

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

SDV Requirements and Guidelines

Be aware of the following requirements and guidelines as you plan and configure SDV:

- SDV should be enabled on switches where devices that are part of SDV zones are connected.
- SDV does not work for devices connected to non-MDS switches.
- Broadcast zoning is not supported for a zone with a virtual device.
- IVR and SDV cannot be used for the same device. A SDV-virtualized device cannot be part of an IVR zone or zoneset.
- Virtual device names should be unique across VSANs because they are registered with the device alias server, which is unaware of VSANs. For example, if you have enabled SDV and have registered a name, vt1 in both VSAN 1 and VSAN 2, then the device alias server cannot store both entries because they have the same name.
- You cannot specify the same primary device for different virtual devices.
- SDV does not work with soft zoning (*Soft zoning* means that zoning restrictions are applied only during interaction between the name server and the end device). If an end device somehow knows the FC ID of a device outside its zone, it can access that device; it does not work with the **zone default-zone permit vsan** operation (which would otherwise permit or deny traffic to members in the default zone).
- If devices are not already zoned with the initiators, then you can configure SDV virtual device zones with no negative impact. If they are already zoned, then zoning changes are required.
- The real device-virtual device zone cannot coexist with the real device-real device zone. If the real devices are not already zoned together, then you can configure the real device-virtual device zone with no negative impact. If these devices are already zoned, then adding the real device-virtual device zone may cause the zone activation to fail. If this occurs, then you must delete one of the zones before activation.

For example, a user attempts to create a configuration with zone A, which consists of I, the initiator, and T, the target (I,T), and zone B, which consists of a virtual initiator, VI, and real target, T (zone VI, T). Such a configuration would fail. Likewise, an attempt to configure zone C, which consists of an initiator, I, and target T, with zone D, which consists of an initiator, I, and virtual target, VT (zone I, VT), would also fail.



Caution

There must be at least one SDV-enabled switch that is *not* a Cisco MDS 9124 Switch between the server and the device that are being virtualized. SDV does not work when initiators and primary devices are connected to the same Cisco MDS 9124 Switch.

Guidelines for Downgrading SDV

As of MDS NX-OS Release 4.1(1a), SDV supports failover and fallback attribute configuration. Downgrading to an earlier release requires you to remove the attribute configurations before downgrading.

As of SAN-OS Release 3.1(3), SDV supports virtual initiators and LUN zoning. Consequently, in SAN-OS Releases 3.1(3) and later, if virtual initiators are configured or SDV devices are configured as LUN-based members of a zone, a configuration check will indicate that downgrading to SAN-OS Release 3.1(2) may be disruptive and is therefore not recommended.

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

Downgrading with Attributes Configured

As of MDS NX-OS Release 4.1(1a), SDV supports failover and fallback attribute configuration. To successfully downgrade to an earlier release, you must remove the attribute configurations before downgrading.

Downgrading with Virtual Initiators Configured

If SDV virtual initiators are configured, you will be unable to downgrade to SAN-OS Release 3.1(2).

This incompatibility only warns before a downgrade. We recommend that you remove the virtual initiator configuration or shut down the initiator port so that there are no inconsistencies in the downgraded version.

Downgrading with SDV LUN Zoning Configured

The following are downgrade scenarios when SDV LUN zoning is configured:

- Real initiator and SDV virtual target with LUN
- SDV virtual initiator and real target with LUN
- SDV virtual initiator and SDV virtual target with LUN

In each of these cases, a configuration check is registered to prevent users from downgrading to SAN-OS Release 3.1(2). This incompatibility will be disruptive if you proceed with the downgrade.

To avoid the configuration check, delete all the LUN zone members from SDV zones, and then activate the zone set before the downgrade.

Default Settings

Table 1-1 lists the default settings for SDV parameters.

Table 1-1 Default SDV Configuration Parameters

Parameters	Default
enable	disabled

Configuring SDV

SDV is a distributed service and uses Cisco Fabric Services (CFS) distribution to synchronize the databases. When you configure SDV, it starts a CFS session and locks the fabric. When a fabric is locked, Cisco NX-OS software does not allow any configuration changes from a switch other than the switch holding the lock and issues a message to inform users about the locked status. Configuration changes are held in a pending database for the application. You must perform a commit operation to make the configuration active and to release the lock for all switches.

Refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide* for more details about CFS.

Send documentation comments to dcnm-san-docfeedback@cisco.com

**Note**

When you enable SDV, CFS distribution is also enabled; CFS distribution cannot be disabled for SDV.

This section includes the following topics:

- [Configuring a Virtual Device, page 1-8](#)
- [Configuring a Zone for a Virtual Device, page 1-9](#)
- [Linking a Virtual Device with a Physical Device, page 1-10](#)
- [Discarding Changes, page 1-11](#)
- [Clearing SDV Changes, page 1-11](#)

Configuring a Virtual Device

A virtual device is identified by an alphanumeric name of up to 32 characters and defines all the real devices (one primary and one or more secondary) that it represents. Upon the successful creation of a virtual device, the virtual device name is internally registered as the device alias name with the device alias database; the pWWN is automatically assigned by the system using Cisco Organizational Unique Identifier (OUI). A virtual device appears as a real, physical device. You can enumerate up to 128 devices for a virtual device. There is a limit of 4095 on the number of virtual devices that you can create in a single VSAN.

**Note**

As of Cisco MDS SAN-OS Release 3.1(2) and NX-OS Release 4.1(1a), SDV supports up to 1024 virtual devices per VSAN.

Detailed Steps

To configure a virtual target and commit it to the fabric configuration, follow these steps:

- Step 1** Expand SAN in the Logical Domains pane, and then expand the fabric in which your VSAN resides.
- Step 2** Expand the VSAN in which you want to create the virtual target and select SDV. You see the switches in the VSAN that you selected listed in the Information pane.
- Step 3** In the **Control** tab, select **enable** from the drop-down menu in the Command column to enable SAN device virtualization for a particular switch in the VSAN.
- Step 4** Click the **Apply Changes** icon to commit the configuration change.
- Step 5** Click the **CFS** tab. Confirm that the SAN device virtualization feature is enabled for the switch.
- Step 6** Click the **Virtual Devices** tab and then click the **Create Row** icon.
You see the Create Virtual Devices dialog box.
- Step 7** Select the Virtual Device ID from the drop-down list (ranges from 1 to 4096).
- Step 8** Enter a Name for the Virtual Device. Select the Virtual Domain and enter a Virtual FC ID for the virtual target.
- Step 9** Check only the **autoFailover** check box or check the **autoFailover** and **primFallback** check boxes. For more information, see the [“Automatic Failover and Fallback” section on page 1-4](#). You can also change the option in the Option column of the Virtual Devices tab.
- Step 10** Click **Create** to create the virtual target.


Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 11 Click the CFS icon to commit and distribute the configuration changes.

Examples

Figure 1-5 shows a configuration that includes a new virtual device, vt1.

Figure 1-5 Creating a Virtual Device



The pWWN of the virtual target does not appear in the zoning end devices database in DCNM-SAN. If you want to zone the virtual device with a pWWN, you must enter it in the Add Member to Zone dialog box when creating a zone. However, if the device alias is in enhanced mode, the virtual device names appear in the device alias database in the DCNM-SAN zoning window. In this case, users can choose to select either the device alias name or enter the pWWN in the Add Member to Zone dialog box.

For more information, see the “Adding Zone Members” section on page 1-16.

Configuring a Zone for a Virtual Device

After configuring a virtual device, you must create a zone that includes all the other real devices and the virtual device as members, and add this zone to a zone set, which you can activate. You can add the virtual device to the zone using the configured name and member type as the device alias.



Note

This configuration process does not support interoperability. If you are working in interop-VSANs, we recommend that you configure the zone directly using the system-assigned pWWN of the virtual device.


Set the device alias mode to enhanced when using SDV (because the pWWN of a virtual device could change).

Send documentation comments to dcnm-san-docfeedback@cisco.com

Examples

Figure 1-6 shows a virtual device-name device alias (vt1) zoned with the real devices activated; the primary device is online.

Figure 1-6 Zoning the Virtual Device with Real Devices



SDV is enabled on a switch and a virtual device is defined. SDV assigns a pWWN for the virtual device, and it is zoned based on the pWWN in a zone. If you later disable SDV, this configuration is lost. If you reenabling SDV and create the virtual device using the same name, there is no guarantee that it will get the same pWWN again. You would have to rezone the pWWN-based zone. However, if you perform zoning based on the device-alias name, there are no configuration changes required if or when the pWWN changes.

Be sure you understand how device alias modes work before enabling them. Refer to [Chapter 1, "Distributing Device Alias Services"](#) for details and requirements about device alias modes.

Linking a Virtual Device with a Physical Device

After creating a virtual device and configuring it as part of a zone, you can define the primary device for it using the **link** command, which is also used to fail over to the secondary device.



Note

When a link operation fails over to the secondary device, the virtual device is taken offline, and then brought online.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Prerequisites

As of MDS NX-OS Release 4.1(1a), the following conditions must be considered before linking a device:

- If you link to the secondary device which is currently active because of failover, the primary tag is moved to the secondary device and the secondary device becomes the primary device.
- When the secondary device is active, if you link to a third device, and if the fallback attribute was not configured, the third device becomes the primary device but the secondary device continues to be the active device.
- When the secondary device is active, if you link to a third device, and if the fallback attribute was configured, then the third device becomes the primary device as well as the active device.

Detailed Steps

To link a virtual target with a physical target, follow these steps:

-
- | | |
|---------------|---|
| Step 1 | Click the Real Devices tab and then click the Create Row icon. |
| Step 2 | Select the Virtual Device ID from the pull-down list or enter an existing ID for the virtual target that you are linking with a physical target. |
| Step 3 | Select the Real Device ID of the physical target that you are linking with the virtual target. |
| Step 4 | Click either the pWWN or deviceAlias radio button, and select the appropriate pWWN or device alias from the pull-down menu. The Name field is automatically populated when you select the pWWN or device alias. |
| Step 5 | Click either the primary or secondary radio button for the Map Type. |
| Step 6 | Click the CFS icon to save and distribute these changes, or click Close to discard any unsaved changes. |
-

Discarding Changes

At any time, you can discard the uncommitted changes to the running configuration and release the fabric lock (prior to entering the **sdv commit** command). If you discard the pending changes, the configuration remains unaffected and the lock is released.

Clearing SDV Changes

If you have performed a SDV task and have forgotten to release the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your changes to the pending database are discarded and the fabric lock is released.



Tip

The pending database is only available in the volatile directory and is subject to being discarded if the switch is restarted.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field Descriptions for SDV

This section displays the field descriptions for this feature.

SDV Virtual Devices

Field	Description
Name	Represents the name of this virtual device.
Virtual Domain	The user preference for a persistent Domain ID for this virtual device to indicate a specific partition (domain) of the fabric that this virtual device should belong to.
Virtual FCID	The user preference for a persistent FCID for this virtual device.
Port WWN	The assigned pWWN for this virtual device. The agent assigns this value when the configuration is committed.
Node WWN	The assigned nWWN for this virtual device. The agent assigns this value when the configuration is committed.
Assigned FCID	The assigned FCID of this virtual device. The agent assigns this value when the configuration is committed and the real device that this virtual device virtualizes is online.
Real Device Map List	The set of real device(s) that this virtual device virtualizes in this VSAN.

SDV Real Devices

Field	Description
Type	The type of real device identifier represented by the value of the corresponding instance of cFcSdvVirtRealDeviceId that this virtual device virtualizes to.
Name	Represents a real device(s) identifier that this virtual device virtualizes.
Map Type	The mapping association type of the real device(s) (initiator/target).

Additional References

For additional information related to implementing VSANs, see the following section:

- [Related Document, page 1-13](#)
- [Standards, page 1-13](#)
- [RFCs, page 1-13](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- [MIBs, page 1-13](#)

Related Document

Related Topic	Document Title
Cisco MDS 9000 Family Command Reference	<i>Cisco MDS 9000 Family Command Reference</i>

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified.	—

MIBs

MIBs	MIBs Link
<ul style="list-style-type: none">• CISCO-FC-SDV-MIB	To locate and download MIBs, go to the following URL: http://www.cisco.com/en/US/products/ps5989/prod_technical_reference_list.html

Send documentation comments to dcnm-san-docfeedback@cisco.com



CHAPTER 1

Creating Dynamic VSANs

This chapter includes the following topics:

- [Information About DPVM, page 1-1](#)
- [Licensing Requirements for VSANs, page 1-4](#)
- [Guidelines and Limitations, page 1-4](#)
- [Default Settings, page 1-5](#)
- [Creating DPVM, page 1-5](#)
- [Monitoring DPVM, page 1-11](#)
- [Field Descriptions for DPVM, page 1-12](#)
- [Additional References, page 1-13](#)

Information About DPVM

Port VSAN membership on the switch is assigned on a port-by-port basis. By default each port belongs to the default VSAN.

You can dynamically assign VSAN membership to ports by assigning VSANs based on the device WWN. This method is referred to as Dynamic Port VSAN Membership (DPVM). DPVM offers flexibility and eliminates the need to reconfigure the port VSAN membership to maintain fabric topology when a host or storage device connection is moved between two Cisco MDS 9000 family switches or two ports within a switch. It retains the configured VSAN regardless of where a device is connected or moved. To assign VSANs statically, see [Chapter 1, “Configuring and Managing VSANs.”](#)

DPVM configurations are based on port world wide name (pWWN) and node world wide name (nWWN) assignments. A DPVM database contains mapping information for each device pWWN/nWWN assignment and the corresponding VSAN. The Cisco NX-OS software checks the database during a device FLOGI and obtains the required VSAN details.

The pWWN identifies the host or device and the nWWN identifies a node consisting of multiple devices. You can assign any one of these identifiers or any combination of these identifiers to configure DPVM mapping. If you assign a combination, then preference is given to the pWWN.

DPVM uses the Cisco Fabric Services (CFS) infrastructure to allow efficient database management and distribution. DPVM uses the application driven, coordinated distribution mode and the fabric-wide distribution scope (for information about CFS, refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*).

Send documentation comments to dcnm-san-docfeedback@cisco.com

**Note**

DPVM does not cause any changes to device addressing. DPVM only pertains to the VSAN membership of the device, ensuring that the host gets the same VSAN membership on any port on the switch. For example, if a port on the switch has a hardware failure, you can move the host connection to another port on the switch and you do not need to update the VSAN membership manually.

**Note**

DPVM is not supported on FL ports. DPVM is supported only on F ports.

- [About DPVM Configuration, page 1-2](#)
- [About DPVM Databases, page 1-2](#)
- [About Autolearned Entries, page 1-3](#)
- [About DPVM Database Distribution, page 1-3](#)
- [About Locking the Fabric, page 1-4](#)
- [About Copying DPVM Databases, page 1-4](#)

About DPVM Configuration

To use the DPVM feature as designed, be sure to verify the following requirements:

- The interface through which the dynamic device connects to the Cisco MDS 9000 Family switch must be configured as an F port.
- The static port VSAN of the F port should be valid (not isolated, not suspended, and in existence).
- The dynamic VSAN configured for the device in the DPVM database should be valid (not isolated, not suspended, and in existence).

**Note**

The DPVM feature overrides any existing static port VSAN membership configuration. If the VSAN corresponding to the dynamic port is deleted or suspended, the port is shut down.

About DPVM Databases

The DPVM database consists of a series of device mapping entries. Each entry consists of a device pWWN or nWWN assignment along with the dynamic VSAN to be assigned. You can configure a maximum of 16,000 DPVM entries in the DPVM database. This database is global to the whole switch (and fabric) and is not maintained for each VSAN.

The DPVM feature uses three databases to accept and implement configurations.

- Configuration (config) database—All configuration changes are stored in the configuration database when distribution is disabled.
- Active database—The database currently enforced by the fabric.
- Pending database—All configuration changes are stored in the DPVM pending database when distribution is enabled (see the [“About DPVM Database Distribution”](#) section on page 1-3).

Send documentation comments to dcnm-san-docfeedback@cisco.com

Changes to the DPVM config database are not reflected in the active DPVM database until you activate the DPVM config database. Changes to the DPVM pending database are not reflected in the config or active DPVM database until you commit the DPVM pending database. This database structure allows you to create multiple entries, review changes, and let the DPVM config and pending databases take effect.

About Autolearned Entries

The DPVM database can be configured to automatically learn (autolearn) about new devices within each VSAN. The autolearn feature can be enabled or disabled at any time. Learned entries are created by populating device pWWNs and VSANs in the active DPVM database. The active DPVM database should already be available to enable autolearn.

You can delete any learned entry from the active DPVM database when you enable autolearn. These entries only become permanent in the active DPVM database when you disable autolearn.



Note

Autolearning is only supported for devices connected to F ports. Devices connected to FL ports are not entered into the DPVM database because DPVM is not supported on FL ports.

The following conditions apply to learned entries:

- If a device logs out while autolearn is enabled, that entry is automatically deleted from the active DPVM database.
- If the same device logs multiple times into the switch through different ports, then the VSAN corresponding to last login is remembered.
- Learned entries do not override previously configured and activated entries.
- Learning is a two-part process—Enabling autolearning followed by disabling autolearning. When the **auto-learn** option is enabled, the following applies:
 - Learning currently logged-in devices—Occurs from the time learning is enabled.
 - Learning new device logins—Occurs as and when new devices log in to the switch.

About DPVM Database Distribution

Using the CFS infrastructure, each DPVM server learns the DPVM database from each of its neighboring switches during the ISL bring-up process. If you change the database locally, the DPVM server notifies its neighboring switches, and that database is updated by all switches in the fabric.

If fabric distribution is enabled, all changes to the configuration database are stored in the DPVM pending database. These changes include the following tasks:

- Adding, deleting, or modifying database entries.
- Activating, deactivating, or deleting the configuration database.
- Enabling or disabling autolearning.

These changes are distributed to all switches in a fabric when you commit the changes. You can also discard (abort) the changes at this point.

Send documentation comments to dcnm-san-docfeedback@cisco.com



Tip

See the “[Viewing the Pending Database](#)” section on page 1-11 to view the contents of the of the pending database.

If the DPVM database is available on all switches in the fabric, devices can be moved anywhere and offer the greatest flexibility. To enable database distribution to the neighboring switches, the database should be consistently administered and distributed across all switches in the fabric. The Cisco NX-OS software uses the Cisco Fabric Services (CFS) infrastructure to achieve this requirement (refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*).

About Locking the Fabric

The first action that modifies the existing configuration creates the DPVM pending database and locks the feature in the fabric. Once you lock the fabric, the following conditions apply:

- No other user can make any configuration changes to this feature.
- A copy of the configuration database becomes the DPVM pending database. Modifications from this point on are made to the DPVM pending database. The DPVM pending database remains in effect until you commit the modifications to the DPVM pending database or discard (abort) the changes to the DPVM pending database.

About Copying DPVM Databases

The following circumstances may require the active DPVM database to be copied to the DPVM config database:

- If the learned entries are only added to the active DPVM database.
- If the DPVM config database or entries in the DPVM config database are accidentally deleted.



Note

If you copy the DPVM database and fabric distribution is enabled, you must commit the changes.

Licensing Requirements for VSANs

The following table shows the licensing requirements for this feature:

License	License Description
ENTERPRISE_PKG	The enterprise license is required to enable VSAN. For a complete explanation of the licensing scheme, see the <i>Cisco MDS 9000 Family NX-OS Licensing Guide</i> .

Guidelines and Limitations

This section explains the database guidelines for this feature.

A database merge refers to a union of the configuration database and static (unlearned) entries in the active DPVM database. For information about CFS merge support, refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide* for detailed concepts.

Send documentation comments to dcnm-san-docfeedback@cisco.com

When merging the DPVM database between two fabrics, follow these guidelines:

- Verify that the activation status and the autolearn status is the same in both fabrics.
- Verify that the combined number of device entries in each database does not exceed 16 K.

**Caution**

If you do not follow these two conditions, the merge will fail. The next distribution will forcefully synchronize the databases and the activation states in the fabric.

Default Settings

Table 1-1 lists the default settings for DPVM parameters.

Table 1-1 **Default DPVM Parameters**

Parameters	Default
DPVM	Disabled.
DPVM distribution	Enabled.
Autolearning	Disabled.

Creating DPVM

This section includes the following topics:

- [Configuring DPVM with the DPVM Wizard, page 1-5](#)
- [Configuring DPVM Config and Pending Databases, page 1-6](#)
- [Activating DPVM Config Databases, page 1-7](#)
- [Enabling Autolearning, page 1-7](#)
- [Clearing Learned Entries, page 1-8](#)
- [Disabling DPVM Database Distribution, page 1-8](#)
- [Locking the Fabric, page 1-9](#)
- [Committing Changes, page 1-9](#)
- [Discarding Changes, page 1-10](#)
- [Clearing a Locked Session, page 1-10](#)
- [Copying DPVM Databases, page 1-10](#)
- [Comparing Database Differences, page 1-11](#)

Configuring DPVM with the DPVM Wizard


Detailed Steps

To use the DPVM Setup Wizard to set up dynamic port VSAN membership, follow these steps:

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Step 1** Click the **DPVM Setup Wizard** icon in the DCNM-SAN toolbar (See [Figure 1-1](#)).

Figure 1-1 DPVM Wizard Icon



You see the Select Master Switch page.


- Step 2** Click the switch you want to be the master switch. This switch controls the distribution of the DPVM database to other switches in the fabric.
- Step 3** Click **Next**.
You see the AutoLearn Current End Devices page.
- Step 4** (Optional) Click the **Create Configuration From Currently Logged In End Devices** check box if you want to turn on autolearning.
- Step 5** Click **Next**.
You see the Edit and Activate Configuration page.
- Step 6** Verify the current or autolearned configuration. Optionally, click **Insert** to add more entries into the DPVM config database.
- Step 7** Click **Finish** to update the DPVM config database, distribute the changes using CFS, and activate the database, or click **Cancel** to exit the DPVM Setup Wizard without saving changes.
- Step 8** Select the switch you want to be the master switch. This switch controls the distribution of the DPVM database to other switches in the fabric.
- Step 9** Click **Next**.
You see the AutoLearn Current End Devices page.
- Step 10** (Optional) Check the **Create Configuration From Currently Logged In End Devices** check box if you want to enable autolearning.
- Step 11** Click **Next**.
You see the Edit and Activate Configuration page.
- Step 12** Verify the current or autolearned configuration. Optionally, click **Insert** to add more entries into the DPVM config database.
- Step 13** Click **Finish** to update the DPVM config database, distribute the changes using CFS, and activate the database, or click **Cancel** to exit the DPVM Setup Wizard without saving changes.

Configuring DPVM Config and Pending Databases

Detailed Steps

To create and populate the config and pending databases, follow these steps:

Send documentation comments to dcnm-san-docfeedback@cisco.com

-
- Step 1** Expand **Fabricxx> All VSANs**, and then select **DPVM** in the Logical Attributes pane.
You see the DPVM configuration in the Information pane.
- Step 2** Click the **CFS** tab and select a master switch by checking a check box in the Master column.
-  **Note** You must click the CFS tab in order to activate the other tabs.
-
- Step 3** Click the **Config Database** tab and then click the **Create Row** to insert a new entry.
You see the Create Config Database dialog box.
- Step 4** Choose an available WWN and VSAN combination or fill in the pWWN and Login VSAN fields.
- Step 5** Click **Create** to save these changes in the config or pending database or click **Close** to discard any unsaved changes.
- Step 6** Click the **CFS** tab and select the Config Action drop-down menu for the master database.
- Step 7** Select **commit** from the drop-down menu to distribute these changes or **abort** to discard the changes.
-

Activating DPVM Config Databases

When you explicitly activate the DPVM config database, the DPVM config database becomes the active DPVM database. Activation may fail if conflicting entries are found between the DPVM config database and the currently active DPVM database. However, you can force activation to override conflicting entries.

Detailed Steps

To activate the DPVM config database, follow these steps:

-
- Step 1** Expand **Fabricxx> All VSANs**, and then select **DPVM** from the Logical Attributes pane.
You see the DPVM configuration in the Information pane.
- Step 2** Click the **Action** tab and set the Action drop-down menu to **activate** or **forceActivate** to activate the DPVM config database.
- Step 3** Click the **CFS** tab and select the Config Action drop-down menu for the master database.
- Step 4** Select **commit** from the drop-down menu to distribute these changes or **abort** to discard the changes.
-



Note To disable DPVM, you must explicitly deactivate the currently active DPVM database.

Enabling Autolearning

Detailed Steps

To enable autolearning, follow these steps:

Send documentation comments to dcnm-san-docfeedback@cisco.com

-
- Step 1** Expand **Fabricxx> All VSANs**, and then select DPVM from the Logical Attributes pane.
You see the DPVM configuration in the Information pane.
- Step 2** Click the **Actions** tab and check the **Auto Learn Enable** check box to enable autolearning.
- Step 3** Click the **CFS** tab and select **commit** to distribute these changes or **abort** to discard the changes.
-

Clearing Learned Entries

You can clear DPVM entries from the active DPVM database (if autolearn is still enabled) using one of two methods.

Restrictions

These two procedures do not start a session and can only be issued in the local switch.

Detailed Steps

To clear a single autolearn entry, follow these steps:

-
- Step 1** Expand **Fabricxx> All VSANs**, and then select **DPVM** from the Logical Attributes pane.
You see the DPVM configuration in the Information pane.
- Step 2** Click the **Actions** tab and select **clearOnWWN** from the Auto Learn Clear drop-down menu.
- Step 3** Check the **clear WWN** check box next to the WWN of the autolearned entry that you want to clear.
- Step 4** Click **CFS** and select **commit** to distribute these changes or **abort** to discard the changes.
-

To clear all autolearn entries, follow these steps:

-
- Step 1** Expand **Fabricxx> All VSANs**, and then select **DPVM** from the Logical Attributes pane.
You see the DPVM configuration in the Information pane.
- Step 2** Click the **Actions** tab.
You see the DPVM Actions menu.
- Step 3** Select **clear** from the Auto Learn Clear drop-down menu.
- Step 4** Click the **CFS** tab and select **commit** to distribute these changes or **abort** to discard the changes.
-

Disabling DPVM Database Distribution

Detailed Steps

These changes are distributed to all switches in a fabric when you commit the changes. You can also discard (abort) the changes at this point.

Send documentation comments to dcnm-san-docfeedback@cisco.com

**Tip**

See the “[Viewing the Pending Database](#)” section on page 1-11 to view the contents of the pending database.

To disable DPVM database distribution to the neighboring switches, follow these steps:

-
- Step 1** Expand **Fabricxx> All VSANs**, and then select **DPVM** from the Logical Attributes pane.
You see the DPVM configuration in the Information pane.
- Step 2** Click the **CFS** tab and select **disable** from the Admin drop-down menu.
- Step 3** Click **Apply Changes** to save this change or click **Undo Changes** to discard the change.
-

Locking the Fabric

Detailed Steps

To lock the fabric and apply changes to the DPVM pending database, follow these steps:

-
- Step 1** Expand **Fabricxx> All VSANs**, and then select **DPVM** from the Logical Attributes pane.
You see the DPVM configuration in the Information pane.
- Step 2** Click the **Config Database** tab and **Create Row**.
You see the Create Config Database dialog box.
- Step 3** Choose an available pWWN and login VSAN.
- Step 4** Click **Create** to save this change to the pending database or click **Close** to discard any unsaved change.
-

Committing Changes

If you commit the changes made to the configuration, the configuration in the DPVM pending database are distributed to other switches. On a successful commit, the configuration change is applied throughout the fabric and the lock is released.

Detailed Steps

To commit the DPVM pending database, follow these steps:

-
- Step 1** Expand **Fabricxx> All VSANs**, and then select **DPVM** from the Logical Attributes pane.
You see the DPVM configuration in the Information pane.
- Step 2** Click the **CFS** tab and select **commit from** the Config Action drop-down menu.
- Step 3** Click **Apply Changes** to save this change or click **Undo Changes** to discard the change.
-

Send documentation comments to dcnm-san-docfeedback@cisco.com

Discarding Changes

If you discard (abort) the changes made to the DPVM pending database, the configurations remain unaffected and the lock is released.

Detailed Steps

To discard the DPVM pending database, follow these steps:

-
- Step 1** Expand **Fabricxx> All VSANs**, and then select **DPVM** from the Logical Attributes pane.
You see the DPVM configuration in the Information pane.
 - Step 2** Click the **CFS** tab and select **abort** from the Config Action drop-down menu.
 - Step 3** Click **Apply Changes** to save this change or click **Undo Changes** to discard the change.
-

Clearing a Locked Session

If you have performed a DPVM task and have forgotten to release the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your changes to the DPVM pending database are discarded and the fabric lock is released.

Restrictions

The DPVM pending database is only available in the volatile directory and is subject to being discarded if the switch is restarted.

Detailed Steps

To use administrative privileges and release a locked DPVM session using DCNM-SAN, follow these steps:

-
- Step 1** Expand **Fabricxx> All VSANs**, and then select **DPVM** from the Logical Attributes pane.
You see the DPVM configuration in the Information pane.
 - Step 2** Click the **CFS** tab and select **clear** from the Config Action drop-down menu.
 - Step 3** Click **Apply Changes** to save this change or click **Undo Changes** to discard the change.
-

Copying DPVM Databases

Detailed Steps

To copy the currently active DPVM database to the DPVM config database, follow these steps:

Send documentation comments to dcnm-san-docfeedback@cisco.com

-
- Step 1** Expand **Fabricxx> All VSANs**, and then select **DPVM** in the Logical Attributes pane.
You see the DPVM configuration in the Information pane.
- Step 2** Click the **Actions** tab and check the **CopyActive to Config** check box.
- Step 3** Click the **CFS** tab and select **commit** from the Config Action drop-down menu.
-

Comparing Database Differences

Detailed Steps

To compare the currently active database entries to the DPVM config database, follow these steps:

-
- Step 1** Expand **Fabricxx> All VSANs**, and then select **DPVM** from the Logical Attributes pane.
You see the DPVM configuration in the Information pane.
- Step 2** Click the **Active Database** tab.
You see the DPVM active database in the Information pane.
- Step 3** Select **Config** from the Compare With drop-down menu.
You see the comparison dialog box.
- Step 4** Select **Close** to close the comparison dialog box.
-

Monitoring DPVM

To view the configuration information, perform the following tasks:

- [Viewing the Pending Database, page 1-11](#)

Viewing the Pending Database

To view the pending database, follow these steps:

-
- Step 1** Expand **Fabricxx> All VSANs**, and then select **DPVM** from the Logical Attributes pane.
You see the DPVM configuration in the Information pane.
- Step 2** Click the **CFS** tab and set the Config View drop-down menu to **pending**.
- Step 3** Click **Apply Changes**.
- Step 4** Click the **Config Database** tab.
You see the pending database entries.
-

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field Descriptions for DPVM

This section describes the field descriptions for this feature.

DPVM Actions

Field	Description
Action	Helps in activating the set of bindings.
Result	Indicates the outcome of the activation.
Status	Indicates the state of activation. If true, then activation has been attempted as the most recent operation. If false, then an activation has not been attempted as the most recent operation.
CopyActive to Config	When set to copy(1), results in the active (enforced) binding database to be copied on to the configuration binding database. The learned entries are also copied.
Auto Learn Enable	Helps to learn the configuration of devices logged into the local device on all its ports and the VSANs to which they are associated.
Auto Learn Clear	Assists in clearing the autolearned entries.
Clear WWN	Represents the Port WWN (pWWN) to be used for clearing its corresponding autolearned entry.

DPVM Config Database

Field	Description
Type	Specifies the type of the corresponding instance of cdpvmLoginDev object.
WWN or Name	Represents the logging-in device.
VSAN Id	Represents the VSAN to be associated to the port on the local device on which the device represented by cdpvmLoginDev logs in.
Switch Interface	Represents the device alias.

DPVM Active Database

Field	Description
Type	Specifies the type of the corresponding instance of cdpvmEnfLoginDev.
WWN or Name	Represents the logging in device address.
VSAN Id	Represents the VSAN of the port on the local device through which the device represented by cdpvmEnfLoginDev logs in.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
Interface	Represents the device alias.
IsLearnt	Indicates whether this is a learned entry or not. If true, then it is a learned entry. If false, then it is not.

Additional References

For additional information related to implementing VSANs, see the following section:

- [Related Document, page 1-13](#)
- [Standards, page 1-13](#)
- [RFCs, page 1-13](#)
- [MIBs, page 1-13](#)

Related Document

Related Topic	Document Title
Cisco MDS 9000 Family Command Reference	<i>Cisco MDS 9000 Family Command Reference</i>

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified.	—

MIBs

MIBs	MIBs Link
<ul style="list-style-type: none">• CISCO-DYNAMIC-PORT-VSAN-MIB	To locate and download MIBs, go to the following URL: http://www.cisco.com/en/US/products/ps5989/prod_technical_reference_list.html

Send documentation comments to dcnm-san-docfeedback@cisco.com



CHAPTER 1

Configuring and Managing Zones

Zoning enables you to set up access control between storage devices or user groups. If you have administrator privileges in your fabric, you can create zones to increase network security and to prevent data loss or corruption. Zoning is enforced by examining the source-destination ID field.

Advanced zoning capabilities specified in the FC-GS-4 and FC-SW-3 standards are provided. You can use either the existing basic zoning capabilities or the advanced, standards-compliant zoning capabilities.

This chapter includes the following topics:

- [Information About Zoning, page 1-1](#)
- [Guidelines and Limitations, page 1-13](#)
- [Default Settings, page 1-15](#)
- [Configuring Zones, page 1-15](#)
- [Configuring Zone Sets, page 1-20](#)
- [Verifying Zone Configuration, page 1-38](#)
- [Configuration Examples for Zoning, page 1-38](#)
- [Field Descriptions for Zones, page 1-39](#)
- [Additional References, page 1-44](#)

Information About Zoning

Zoning has the following features:

- A zone consists of multiple zone members.
 - Members in a zone can access each other; members in different zones cannot access each other.
 - If zoning is not activated, all devices are members of the default zone.
 - If zoning is activated, any device that is not in an active zone (a zone that is part of an active zone set) is a member of the default zone.
 - Zones can vary in size.
 - Devices can belong to more than one zone.
 - A physical fabric can have a maximum of 16,000 members. This includes all VSANs in the fabric.
- A zone set consists of one or more zones.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- A zone set can be activated or deactivated as a single entity across all switches in the fabric.
 - Only one zone set can be activated at any time.
 - A zone can be a member of more than one zone set.
 - A zone switch can have a maximum of 500 zone sets.
- Zoning can be administered from any switch in the fabric.
 - When you activate a zone (from any switch), all switches in the fabric receive the active zone set. Additionally, full zone sets are distributed to all switches in the fabric, if this feature is enabled in the source switch.
 - If a new switch is added to an existing fabric, zone sets are acquired by the new switch.
- Zone changes can be configured non-disruptively. New zones and zone sets can be activated without interrupting traffic on unaffected ports or devices.
- Zone membership criteria is based mainly on WWNs or FC IDs.
 - Port world wide name (pWWN)—Specifies the pWWN of an N port attached to the switch as a member of the zone.
 - Fabric pWWN—Specifies the WWN of the fabric port (switch port's WWN). This membership is also referred to as port-based zoning.
 - FC ID—Specifies the FC ID of an N port attached to the switch as a member of the zone.
 - Interface and switch WWN (sWWN)—Specifies the interface of a switch identified by the sWWN. This membership is also referred to as interface-based zoning.
 - Interface and domain ID—Specifies the interface of a switch identified by the domain ID.
 - Domain ID and port number—Specifies the domain ID of an MDS domain and additionally specifies a port belonging to a non-Cisco switch.
 - IPv4 address—Specifies the IPv4 address (and optionally the subnet mask) of an attached device.
 - IPv6 address—The IPv6 address of an attached device in 128 bits in colon(:)-separated hexadecimal format.
- Default zone membership includes all ports or WWNs that do not have a specific membership association. Access between default zone members is controlled by the default zone policy.
- You can configure up to 8000 zones per VSAN and a maximum of 8000 zones for all VSANs on the switch.

This section includes the following topics:

- [Zone Implementation, page 1-3](#)
- [About the Edit Local Full Zone Database Tool, page 1-4](#)
- [About Zone Sets, page 1-5](#)
- [About Zone Set Creation, page 1-5](#)
- [About the Default Zone, page 1-6](#)
- [About FC Alias Creation, page 1-6](#)
- [Zone Enforcement, page 1-7](#)
- [Zone Set Distribution, page 1-7](#)
- [About Recovering from Link Isolation, page 1-8](#)
- [Zone Set Duplication, page 1-8](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- [About Backing Up and Restoring Zones, page 1-9](#)
- [About Zone-Based Traffic Priority, page 1-9](#)
- [About Broadcast Zoning, page 1-9](#)
- [About LUN Zoning, page 1-10](#)
- [About Read-Only Zones, page 1-10](#)
- [About Enhanced Zoning, page 1-11](#)
- [Merging the Database, page 1-11](#)
- [Smart Zoning, page 1-12](#)

Zone Implementation

All switches in the Cisco MDS 9000 Family automatically support the following basic zone features (no additional configuration is required):

- Zones are contained in a VSAN.
- Hard zoning cannot be disabled.
- Name server queries are soft-zoned.
- Only active zone sets are distributed.
- Unzoned devices cannot access each other.
- A zone or zone set with the same name can exist in each VSAN.
- Each VSAN has a full database and an active database.
- Active zone sets cannot be changed, without activating a full zone database.
- Active zone sets are preserved across switch reboots.
- Changes to the full database must be explicitly saved.
- Zone reactivation (a zone set is active and you activate another zone set) does not disrupt existing traffic.

If required, you can additionally configure the following zone features:

- Propagate full zone sets to all switches on a per VSAN basis.
- Change the default policy for unzoned members.
- Interoperate with other vendors by configuring a VSAN in the interop mode. You can also configure one VSAN in the interop mode and another VSAN in the basic mode in the same switch without disrupting each other.
- Bring E ports out of isolation.

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)


About the Edit Local Full Zone Database Tool

You can use the Edit Full Zone Database Tool to complete the following tasks:

- You can display information by VSAN by using the pull-down menu without having to get out of the screen, selecting a VSAN, and re-entering.
- You can use the **Add to zone or alias** button to move devices up or down by alias or by zone.
- You can add zoning characteristics based on alias in different folders.
- You can triple-click to rename zone sets, zones, or aliases in the tree.

The Edit Local Full Zone Database tool allows you to zone across multiple switches and all zoning features are available through the Edit Local Full Zone Database dialog box (see [Figure 1-1](#)).

Figure 1-1 Edit Local Full Zone Database Dialog Box



1	You can display information by VSAN by using the drop-down menu without closing the dialog box, selecting a VSAN, and re-entering.	3	You can add zoning characteristics based on alias in different folders.
2	You can use the Add to zone button to move devices up or down by alias or by zone.	4	You can triple-click to rename zone sets, zones, or aliases in the tree.

Note

The Device Alias radio button is visible only if device alias is in enhanced mode. For more information, see [“Creating Device Aliases”](#) section on page 1-7.

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

About Zone Sets

Zones provide a method for specifying access control. Zone sets are a grouping of zones to enforce access control in the fabric.

Zone sets are configured with the names of the member zones and the VSAN (if the zone set is in a configured VSAN).

Zone Set Distribution—You can distribute full zone sets using one of two methods: one-time distribution or full zone set distribution.

Zone Set Duplication—You can make a copy of a zone set and then edit it without altering the original zone set. You can copy an active zone set from the bootflash: directory, volatile: directory, or slot0, to one of the following areas:


- To the full zone set
- To a remote location (using FTP, SCP, SFTP, or TFTP)

The active zone set is not part of the full zone set. You cannot make changes to an existing zone set and activate it, if the full zone set is lost or is not propagated.

About Zone Set Creation

In [Figure 1-2](#), two separate sets are created, each with its own membership hierarchy and zone members.

Figure 1-2 Hierarchy of Zone Sets, Zones, and Zone Members



Either zone set A or zone set B can be activated (but not together).



Tip

Zone sets are configured with the names of the member zones and the VSAN (if the zone set is in a configured VSAN).

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

About the Default Zone

Each member of a fabric (in effect a device attached to an Nx port) can belong to any zone. If a member is not part of any active zone, it is considered to be part of the default zone. Therefore, if no zone set is active in the fabric, all devices are considered to be in the default zone. Even though a member can belong to multiple zones, a member that is part of the default zone cannot be part of any other zone. The switch determines whether a port is a member of the default zone when the attached port comes up.



Note

Unlike configured zones, default zone information is not distributed to the other switches in the fabric.

Traffic can either be permitted or denied among members of the default zone. This information is not distributed to all switches; it must be configured in each switch.



Note

When the switch is initialized for the first time, no zones are configured and all members are considered to be part of the default zone. Members are not permitted to talk to each other.

Configure the default zone policy on each switch in the fabric. If you change the default zone policy on one switch in a fabric, be sure to change it on all the other switches in the fabric.



Note

The default settings for default zone configurations can be changed.

The default zone members are explicitly listed when the default policy is configured as permit or when a zone set is active. When the default policy is configured as deny, the members of this zone are not explicitly enumerated when you view the active zone set.



Note

The current default zoning policy in both the switches is deny. In the Cisco MDS 9222i Switch, the active zone set is `coco_isola_zoneset`. In the Cisco MDS 9513 Switch, there is no active zone set. However, because the default zoning policy is deny, the hidden active zone set is `d__default__cfg` which causes zone merge to fail. The behavior is same between two Brocade switches.

You can change the default zone policy for any VSAN by choosing **VSANxx > Default Zone** from the DCNM-SAN menu tree and clicking the **Policies** tab. It is recommended that you establish connectivity among devices by assigning them to a nondefault zone.

About FC Alias Creation

You can assign an alias name and configure an alias member using the following values:

- **pWWN**—The WWN of the N or NL port is in hex format (for example, 10:00:00:23:45:67:89:ab).
- **fWWN**—The WWN of the fabric port name is in hex format (for example, 10:00:00:23:45:67:89:ab).
- **FC ID**—The N port ID is in 0xhhhhhh format (for example, 0xce00d1).
- **Domain ID**—The domain ID is an integer from 1 to 239. A mandatory port number of a non-Cisco switch is required to complete this membership configuration.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- IPv4 address—The IPv4 address of an attached device is in 32 bits in dotted decimal format along with an optional subnet mask. If a mask is specified, any device within the subnet becomes a member of the specified zone.
- IPv6 address—The IPv6 address of an attached device is in 128 bits in colon- (:) separated) hexadecimal format.
- Interface—Interface-based zoning is similar to port-based zoning because the switch interface is used to configure the zone. You can specify a switch interface as a zone member for both local and remote switches. To specify a remote switch, enter the remote switch WWN (sWWN) or the domain ID in the particular VSAN.



Tip

The Cisco NX-OS software supports a maximum of 2048 aliases per VSAN.

Zone Enforcement

Zoning can be enforced in two ways: soft and hard. Each end device (N port or NL port) discovers other devices in the fabric by querying the name server. When a device logs in to the name server, the name server returns the list of other devices that can be accessed by the querying device. If an Nx port does not know about the FC IDs of other devices outside its zone, it cannot access those devices.

In soft zoning, zoning restrictions are applied only during interaction between the name server and the end device. If an end device somehow knows the FC ID of a device outside its zone, it can access that device.

Hard zoning is enforced by the hardware on each frame sent by an Nx port. As frames enter the switch, source-destination IDs are compared with permitted combinations to allow the frame at wirespeed. Hard zoning is applied to all forms of zoning.



Note

Hard zoning enforces zoning restrictions on every frame, and prevents unauthorized access.

Switches in the Cisco MDS 9000 Family support both hard and soft zoning.

Zone Set Distribution

You can distribute full zone sets using one of two methods: one-time distribution or full zone set distribution.

Table 1-1 lists the differences between these distribution methods.

Table 1-1 Zone Set Distribution Differences

One-Time Distribution	Full Zone Set Distribution
Distributes the full zone set immediately.	Does not distribute the full zone set immediately.
Does not distribute the full zone set information along with the active zone set during activation, deactivation, or merge process.	Remembers to distribute the full zone set information along with the active zone set during activation, deactivation, and merge processes.


Send documentation comments to dcnm-san-docfeedback@cisco.com

About Recovering from Link Isolation

When two switches in a fabric are merged using a TE or E port, these TE and E ports may become isolated when the active zone set databases are different between the two switches or fabrics. When a TE port or an E port become isolated, you can recover that port from its isolated state using one of three options:

- Import the neighboring switch's active zone set database and replace the current active zone set (see [Figure 1-3](#)).
- Export the current database to the neighboring switch.
- Manually resolve the conflict by editing the full zone set, activating the corrected zone set, and then bringing up the link.

Figure 1-3 *Importing and Exporting the Database*



Zone Set Duplication

You can make a copy and then edit it without altering the existing active zone set. You can copy an active zone set from the bootflash: directory, volatile: directory, or slot0, to one of the following areas:

- To the full zone set
- To a remote location (using FTP, SCP, SFTP, or TFTP)

The active zone set is not part of the full zone set. You cannot make changes to an existing zone set and activate it, if the full zone set is lost or is not propagated.



Caution

Copying an active zone set to a full zone set may overwrite a zone with the same name, if it already exists in the full zone set database.

Send documentation comments to dcnm-san-docfeedback@cisco.com

About Backing Up and Restoring Zones

You can back up the zone configuration to a workstation using TFTP. This zone backup file can then be used to restore the zone configuration on a switch. Restoring the zone configuration overwrites any existing zone configuration on a switch.

About Zone-Based Traffic Priority

The zoning feature provides an additional segregation method to prioritize select zones in a fabric and set up access control between devices. Using this feature, you can configure the quality of service (QoS) priority as a zone attribute. You can assign the QoS traffic priority attribute to be high, medium, or low. By default, zones with no specified priority are implicitly assigned a low priority. Refer to the *Cisco MDS 9000 NX-OS Family Quality of Service Configuration Guide* for more information.

This feature allows SAN administrators to configure QoS using a familiar data flow identification paradigm. You can configure this attribute on a zone-wide basis rather than between zone members.



Caution

If zone-based QoS is implemented in a switch, you cannot configure the interop mode in that VSAN.

About Broadcast Zoning



Note

Broadcast zoning is not supported on the Cisco Fabric Switch for HP c-Class BladeSystem and the Cisco Fabric Switch for IBM BladeCenter.

You can configure broadcast frames in the basic zoning mode. By default, broadcast zoning is disabled and broadcast frames are sent to all Nx ports in the VSAN. When enabled, broadcast frames are only sent to Nx ports in the same zone, or zones, as the sender. Enable broadcast zoning when a host or storage device uses this feature.

Table 1-2 identifies the rules for the delivery of broadcast frames.

Table 1-2 Broadcasting Requirements

Active Zoning?	Broadcast Enabled?	Frames Broadcast?	Comments
Yes	Yes	Yes	Broadcast to all Nx ports that share a broadcast zone with the source of broadcast frames.
No	Yes	Yes	Broadcast to all Nx ports.
Yes	No	No	Broadcasting is disabled.



Tip

If any NL port attached to an FL port shares a broadcast zone with the source of the broadcast frame, then the frames are broadcast to all devices in the loop.

Send documentation comments to dcnm-san-docfeedback@cisco.com

**Caution**

If broadcast zoning is enabled on a switch, you cannot configure the interop mode in that VSAN.

About LUN Zoning

Logical unit number (LUN) zoning is a feature specific to switches in the Cisco MDS 9000 Family.

**Caution**

LUN zoning can only be implemented in Cisco MDS 9000 Family switches. If LUN zoning is implemented in a switch, you cannot configure the interop mode in that switch.

A storage device can have multiple LUNs behind it. If the device port is part of a zone, a member of the zone can access any LUN in the device. With LUN zoning, you can restrict access to specific LUNs associated with a device.

**Note**

When LUN 0 is not included within a zone, control traffic to LUN 0 (for example, REPORT_LUNS, INQUIRY) is supported, but data traffic to LUN 0 (for example, READ, WRITE) is denied.


- Host H1 can access LUN 2 in S1 and LUN 0 in S2. It cannot access any other LUNs in S1 or S2.
- Host H2 can access LUNs 1 and 3 in S1 and only LUN 1 in S2. It cannot access any other LUNs in S1 or S2.

**Note**

Unzoned LUNs automatically become members of the default zone.

Figure 1-4 shows a LUN-based zone example.

Figure 1-4 LUN Zoning Access



About Read-Only Zones

By default, an initiator has both read and write access to the target's media when they are members of the same Fibre Channel zone. The read-only zone feature allows members to have only read access to the media within a read-only Fibre Channel zone.

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

You can also configure LUN zones as read-only zones. Any zone can be identified as a read-only zone. By default all zones have read-write permission unless explicitly configured as a read-only zone.

About Enhanced Zoning

[Table 1-3](#) lists the advantages of the enhanced zoning feature in all switches in the Cisco MDS 9000 Family.

Table 1-3 Advantages of Enhanced Zoning

Basic Zoning	Enhanced Zoning	Enhanced Zoning Advantages
Administrators can make simultaneous configuration changes. Upon activation, one administrator can overwrite another administrator's changes.	Performs all configurations within a single configuration session. When you begin a session, the switch locks the entire fabric to implement the change.	One configuration session for the entire fabric to ensure consistency within the fabric.
If a zone is part of multiple zone sets, you create an instance of this zone in each zone set.	References to the zone are used by the zone sets as required once you define the zone.	Reduced payload size as the zone is referenced. The size is more pronounced with bigger databases.
The default zone policy is defined per switch. To ensure smooth fabric operation, all switches in the fabric must have the same default zone setting.	Enforces and exchanges the default zone setting throughout the fabric.	Fabric-wide policy enforcement reduces troubleshooting time.
To retrieve the results of the activation on a per switch basis, the managing switch provides a combined status about the activation. It does not identify the failure switch.	Retrieves the activation results and the nature of the problem from each remote switch.	Enhanced error reporting eases the troubleshooting process.
To distribute the zoning database, you must reactivate the same zone set. The reactivation may affect hardware changes for hard zoning on the local switch and on remote switches.	Implements changes to the zoning database and distributes it without reactivation.	Distribution of zone sets without activation avoids hardware changes for hard zoning in the switches.
The MDS-specific zone member types (IPv4 address, IPv6 address, symbolic node name, and other types) may be used by other non-Cisco switches. During a merge, the MDS-specific types can be misunderstood by the non-Cisco switches.	Provides a vendor ID along with a vendor-specific type value to uniquely identify a member type.	Unique vendor type.
The fWWN-based zone membership is only supported in Cisco interop mode.	Supports fWWN-based membership in the standard interop mode (interop mode 1).	The fWWN-based member type is standardized.

Merging the Database

The merge behavior depends on the fabric-wide merge control setting:

- **Restrict**—If the two databases are not identical, the ISLs between the switches are isolated.
- **Allow**—The two databases are merged using the merge rules specified in [Table 1-4](#).

Send documentation comments to dcnm-san-docfeedback@cisco.com

Table 1-4 Database Zone Merge Status

Local Database	Adjacent Database	Merge Status	Results of the Merge
The databases contain zone sets with the same name ¹ but different zones, aliases, and attributes groups.		Successful.	The union of the local and adjacent databases.
The databases contains a zone, zone alias, or zone attribute group object with same name ¹ but different members.		Failed.	ISLs are isolated.
Empty.	Contains data.	Successful.	The adjacent database information populates the local database.
Contains data.	Empty.	Successful.	The local database information populates the adjacent database.

1. In the enhanced zoning mode, the active zone set does not have a name in interop mode 1. The zone set names are only present for full zone sets.



Caution

Remove all non-pWWN-type zone entries on all MDS switches running Cisco SAN-OS prior to merging fabrics if there is a Cisco MDS 9020 switch running FabricWare in the adjacent fabric.

The merge process operates as follows:

1. The software compares the protocol versions. If the protocol versions differ, then the ISL is isolated.
2. If the protocol versions are the same, then the zone policies are compared. If the zone policies differ, then the ISL is isolated.
3. If the zone merge options are the same, then the comparison is implemented based on the merge control setting.
 - a. If the setting is restrict, the active zone set and the full zone set should be identical. Otherwise the link is isolated.
 - b. If the setting is allow, then the merge rules are used to perform the merge.

Smart Zoning

Smart zoning supports zoning among more devices by reducing the number of zoning entries that needs to be programmed by considering device type information without increasing the size of the zone set. Smart zoning enables you to select the end device type. You can select if the end device type should be a host or a target. Smart zoning can be enabled at zone level, zone set level, member, and at VSAN level.



Note

If smart zoning is set at the VSAN level, then you cannot enable or disable smart zoning at zone set level or zone level.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Guidelines and Limitations

This section includes the guidelines and limitations for this feature:

- [Zone Member Configuration Guidelines, page 1-13](#)
- [Active and Full Zone Set Considerations, page 1-13](#)
- [Read-Only Zone Configuration Guidelines, page 1-14](#)

Zone Member Configuration Guidelines

All members of a zone can communicate with each other. For a zone with N members, $N*(N-1)$ access permissions need to be enabled. Avoid configuring large numbers of targets or large numbers of initiators in a single zone. This type of configuration wastes switch resources by provisioning and managing many communicating pairs (initiator-to-initiator or target-to-target) that will never actually communicate with each other. Configuring a single initiator with a single target is the most efficient approach to zoning.

The following guidelines must be considered when creating zone members:

- Configuring only one initiator and one target for a zone provides the most efficient use of the switch resources.
- Configuring the same initiator to multiple targets is accepted.
- Configuring multiple initiators to multiple targets is not recommended.

Active and Full Zone Set Considerations

Before configuring a zone set, consider the following guidelines:

- Each VSAN can have multiple zone sets but only one zone set can be active at any given time.
- When you create a zone set, that zone set becomes a part of the full zone set.
- When you activate a zone set, a copy of the zone set from the full zone set is used to enforce zoning, and is called the active zone set. An active zone set cannot be modified. A zone that is part of an active zone set is called an active zone.
- The administrator can modify the full zone set even if a zone set with the same name is active. However, the modification will be enforced only upon reactivation.
- When the activation is done, the active zone set is automatically stored in persistent configuration. This enables the switch to preserve the active zone set information across switch resets.
- All other switches in the fabric receive the active zone set so they can enforce zoning in their respective switches.
- Hard and soft zoning are implemented using the active zone set. Modifications take effect during zone set activation.
- An FC ID or Nx port that is not part of the active zone set belongs to the default zone and the default zone information is not distributed to other switches.

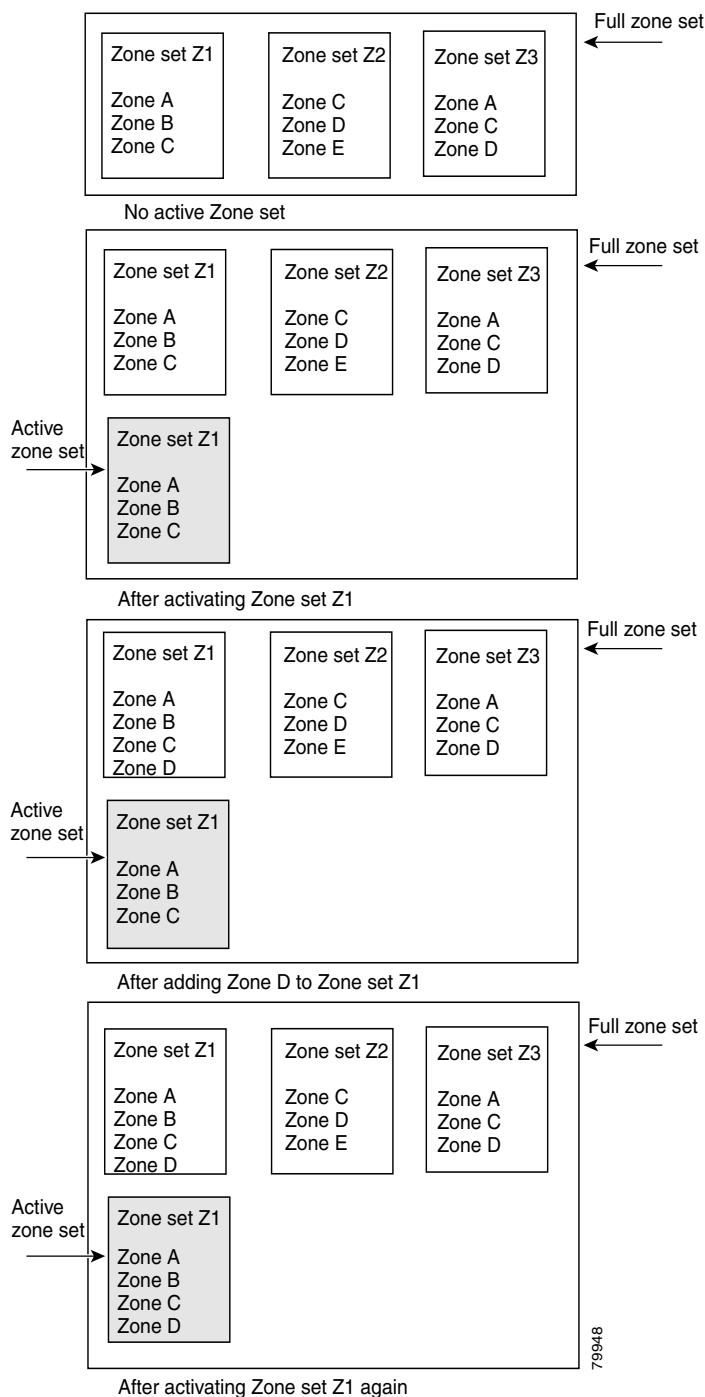
**Note**

If one zone set is active and you activate another zone set, the currently active zone set is automatically deactivated. You do not need to explicitly deactivate the currently active zone set before activating a new zone set.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 1-5 shows a zone being added to an activated zone set.

Figure 1-5 Active and Full Zone Sets



Read-Only Zone Configuration Guidelines

Follow these guidelines when configuring read-only zones:

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

- If read-only zones are implemented, the switch prevents write access to user data within the zone.
- If two members belong to a read-only zone and to a read-write zone, the read-only zone takes priority and write access is denied.
- LUN zoning can only be implemented in Cisco MDS 9000 Family switches. If LUN zoning is implemented in a switch, you cannot configure interop mode in that switch.
- Read-only volumes are not supported by some operating system and file system combinations (for example, Windows NT or Windows 2000 and NTFS file system). Volumes within read-only zones are not available to such hosts. However, if these hosts are already booted when the read-only zones are activated, then read-only volumes are available to those hosts.
- The read-only zone feature behaves as designed if either the FAT16 or FAT32 file system is used with the previously mentioned Windows operating systems.

Default Settings

Table 1-5 lists the default settings for basic zone parameters.

Table 1-5 **Default Basic Zone Parameters**

Parameters	Default
Default zone policy	Denied to all members.
Full zone set distribute	The full zone set(s) is not distributed.
Zone based traffic priority	Low.
Read-only zones	Read-write attributes for all zones.
Broadcast frames	Sent to all Nx ports.
Broadcast zoning	Disabled.
Enhanced zoning	Disabled.

Configuring Zones

This section describes how to configure zones and includes the following topics:

- [Configuring a Zone Using the Zone Configuration Tool, page 1-15](#)
- [Adding Zone Members, page 1-16](#)
- [Filtering End Devices Based on Name, WWN, or FC ID, page 1-17](#)
- [Adding Multiple End Devices to Multiple Zones, page 1-17](#)
- [Using the Quick Config Wizard, page 1-18](#)

Configuring a Zone Using the Zone Configuration Tool


Detailed Steps

To create a zone and move it into a zone set, follow these steps:

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 1 Click the **Zone** icon in the toolbar (see [Figure 1-6](#)).

Figure 1-6 Zone Icon



You see the Select VSAN dialog box.

Step 2 Select the VSAN where you want to create a zone and click **OK**.

You see the Edit Local Full Zone Database dialog box.

If you want to view zone membership information, right-click in the All Zone Membership(s) column, and then click **Show Details** for the current row or all rows from the pop-up menu.

Step 3 Click **Zones** in the left pane and click the **Insert** icon to create a zone.

You see the Create Zone dialog box.

Step 4 Enter a zone name.

Step 5 Check one of the following check boxes:

- a. **Read Only**—The zone permits read and denies write.
- b. **Permit QoS traffic with Priority**—You set the priority from the drop-down menu.
- c. **Restrict Broadcast Frames to Zone Members**

Step 6 Select the **Smart Zoning** check box to enable smart zoning.

Step 7 Click **OK** to create the zone.

If you want to move this zone into an existing zone set, skip to [Step 9](#).

Step 8 Click **Zoneset** in the left pane and click the **Insert** icon to create a zone set.

You see the Zoneset Name dialog box.

Step 9 Enter a zone set name and click **OK**.



Note One of these symbols (\$, -, ^, _) or all alphanumeric characters are supported. In interop mode 2 and 3, this symbol (_) or all alphanumeric characters are supported.

Step 10 Select the zone set where you want to add a zone and click the **Insert** icon or you can drag and drop Zone3 over Zoneset1.

You see the Select Zone dialog box.

Step 11 Click **Add** to add the zone.

Adding Zone Members

Once you create a zone, you can add members to the zone. You can add members using multiple port identification types.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Detailed Steps

To add a member to a zone, follow these steps:

-
- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Select VSAN dialog box.
- Step 2** Select a VSAN and click **OK**.
You see the Edit Local Full Zone Database dialog box for the selected VSAN.
- Step 3** Select the members you want to add from the Fabric pane and click **Add to Zone** or click the zone where you want to add members and click the **Insert** icon.
You see the Add Member to Zone dialog box.
- Step 4** Click the browse button and select a port name or check the **LUN** check box and click the browse button to configure LUNs.
- Step 5** Select the options for **Device Type** field. You can select any one of the options: **Host**, **Storage**, or **Both**.
- Step 6** Click **Add** to add the member to the zone.



Note The Device Alias radio button is visible only if device alias is in enhanced mode. For more information, see [“Creating Device Aliases” section on page 1-7](#).



Note When configuring a zone member, you can specify that a single LUN has multiple IDs depending on the operating system. You can select from six different operating systems.

Filtering End Devices Based on Name, WWN, or FC ID

Detailed Steps

To filter the end devices and device aliases, follow these steps:

-
- Step 1** Click the **Zone** icon in the toolbar.
- Step 2** Select Name, WWN, or FC ID from the With drop-down list.
- Step 3** Enter a filter condition, such as *zo1*, in the Filter text box.
- Step 4** Click **Go**.
-

Adding Multiple End Devices to Multiple Zones

Detailed Steps

To add multiple end devices to multiple zones, follow these steps:

Send documentation comments to dcnm-san-docfeedback@cisco.com

-
- Step 1** Click the **Zone** icon in the toolbar.
- Step 2** Press the **Ctrl** key to select multiple end devices.
- Step 3** Right-click the device and then select **Add to Zone**.
- Step 4** Press the **Ctrl** key to select multiple zones from the pop-up window displayed.
- Step 5** Click **Add**.
- Selected end devices are added to the selected zones.
-

Using the Quick Config Wizard



Note

The Quick Config Wizard supports only switch interface zone members.

As of Cisco SAN-OS Release 3.1(1) and NX-OS Release 4.1(2), you can use the Quick Config Wizard on the Cisco MDS 9124 Switch to add or remove zone members per VSAN. You can use the Quick Config Wizard to perform interface-based zoning and to assign zone members for multiple VSANs using Device Manager.



Note

The Quick Config Wizard is supported on the Cisco MDS 9124 Fabric Switch, the Cisco MDS 9134 Fabric Switch, the Cisco Fabric Switch for HP c-Class BladeSystem, and the Cisco Fabric Switch for IBM BladeCenter.

Restrictions

The Quick Config Wizard can only be used on standalone switches that do not have any existing zoning defined on the switch.

Detailed Steps

To add or remove ports from a zone and to zone only the devices within a specific VSAN using Device Manager on the Cisco MDS 9124 Switch, follow these steps:

-
- Step 1** Choose **FC > Quick Config** or click the **Zone** icon in the toolbar.
- You see the Quick Config Wizard (see [Figure 1-8](#)) with all controls disabled and the Discrepancies dialog box (see [Figure 1-7](#)), which shows all unsupported configurations.




Note

You will see the Discrepancies dialog box only if there are any discrepancies.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 1-7 **Discrepancies Dialog Box**



Step 2 Click **OK** to continue.

You see the Quick Config Wizard dialog box (see [Figure 1-8](#)).

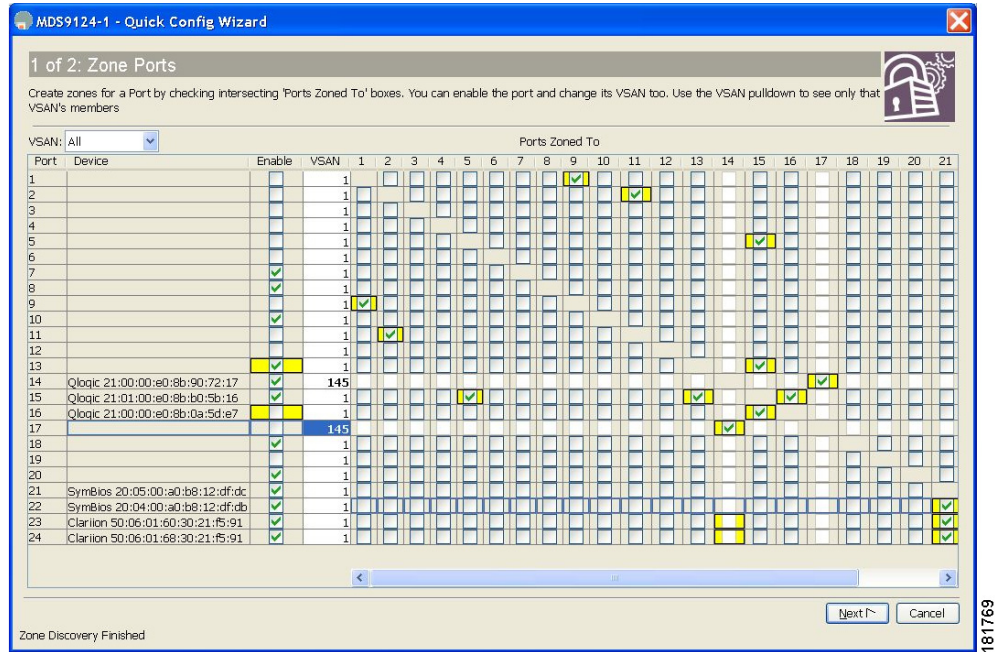


Caution

If there are discrepancies and you click **OK**, the affected VSANs in the zone databases are cleared. This might be disruptive if the switch is in use.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 1-8 Quick Config Wizard



- Step 3** Check the check box in the Ports Zoned To column for the port you want to add or remove from a zone. The check box for the matching port is similarly set. The selected port pair is added or removed from the zone, which creates a two-device zone.
- The VSAN drop-down menu provides a filter that enables you to zone only those devices within a selected VSAN.
- Step 4** Right-click any of the column names to show or hide a column.
- Step 5** Click **Next** to verify the changes.
- You see the Confirm Changes dialog box.
- Step 6** If you want to see the CLI commands, right-click in the dialog box and click **CLI Commands** from the pop-up menu.
- Step 7** Click **Finish** to save the configuration changes.

Configuring Zone Sets

This section describes how to configure zones and includes the following topics:

- [Activating a Zone Set, page 1-21](#)
- [Deactivating a Zone Set, page 1-22](#)
- [Displaying Zone Membership Information, page 1-23](#)
- [Configuring the Default Zone Access Permission, page 1-23](#)
- [Creating FC Aliases, page 1-23](#)
- [Adding Members to Aliases, page 1-24](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- [Converting Zone Members to pWWN-based Members, page 1-24](#)
- [Creating Zone Sets and Adding Member Zones, page 1-25](#)
- [Filtering Zones, Zone Sets, and Device Aliases Based on Name, page 1-26](#)
- [Adding Multiple Zones to Multiple Zone Sets, page 1-26](#)
- [Enabling Full Zone Set Distribution, page 1-26](#)
- [Enabling a One-Time Distribution, page 1-27](#)
- [Importing and Exporting Zone Sets, page 1-27](#)
- [Copying Zone Sets, page 1-28](#)
- [Backing Up Zones, page 1-28](#)
- [Restoring Zones, page 1-29](#)
- [Renaming Zones, Zone Sets, and Aliases, page 1-30](#)
- [Cloning Zones, Zone Sets, FC Aliases, and Zone Attribute Groups, page 1-30](#)
- [Migrating a Non-MDS Database, page 1-31](#)
- [Clearing the Zone Server Database, page 1-31](#)
- [Configuring Zone-Based Traffic Priority, page 1-31](#)
- [Configuring Default Zone QoS Priority Attributes, page 1-31](#)
- [Configuring the Default Zone Policy, page 1-32](#)
- [Configuring Smart Zoning, page 1-32](#)
- [Configuring Global Zone Policies, page 1-33](#)
- [Configuring Broadcast Zoning, page 1-33](#)
- [Configuring a LUN-Based Zone, page 1-34](#)
- [Assigning LUNs to Storage Subsystems, page 1-34](#)
- [Configuring Read-Only Zones, page 1-34](#)
- [Changing from Basic Zoning to Enhanced Zoning, page 1-35](#)
- [Changing from Enhanced Zoning to Basic Zoning, page 1-35](#)
- [Enabling Enhanced Zoning, page 1-36](#)
- [Modifying the Zone Database, page 1-36](#)
- [Creating Attribute Groups, page 1-36](#)
- [Analyzing a Zone Merge, page 1-37](#)
- [Configuring Zone Merge Control Policies, page 1-37](#)
- [Broadcasting a Zone, page 1-37](#)
- [Compacting the Zone Database for Downgrading, page 1-38](#)

Activating a Zone Set

Changes to a zone set do not take effect in a full zone set until you activate it.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Detailed Steps

To activate an existing zone set, follow these steps:

-
- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Select VSAN dialog box.
 - Step 2** Select a VSAN and click **OK**.
You see the Edit Local Full Zone Database dialog box for the selected VSAN.
 - Step 3** Click **Activate** to activate the zone set.
You see the pre-activation check dialog box.
 - Step 4** Click **Yes** to review the differences.
You see the Local vs. Active Differences dialog box.
 - Step 5** Click **Close** to close the dialog box.
You see the Save Configuration dialog box.
 - Step 6** Check the **Save Running to Startup Configuration** check box to save all changes to the startup configuration.
 - Step 7** Click **Continue Activation** to activate the zone set, or click **Cancel** to close the dialog box and discard any unsaved changes.
You see the Zone Log dialog box, which shows if the zone set activation was successful.
-

Deactivating a Zone Set

Detailed Steps

To deactivate an existing zone set, follow these steps:

-
- Step 1** Right-click the zone set you want to deactivate, and then click **Deactivate** from the pop-up menu.
You see the Deactivate Zoneset dialog box.
 - Step 2** Enter deactivate in the text box, and then click **OK**.
You see the Input dialog box.
 - Step 3** Enter deactivate in the text box, and then click **OK** to deactivate the zone set.



Note

To enable this option, you need to modify the server.properties file. Refer to the *Cisco DCNM Fundamentals Guide* to know more about modifying server.properties file.

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

Displaying Zone Membership Information

Detailed Steps

To display zone membership information for members assigned to zones, follow these steps:

-
- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Select VSAN dialog box.
- Step 2** Select a VSAN and click **OK**.
You see the Edit Local Full Zone Database dialog box for the selected VSAN.
- Step 3** Click **Zones** in the left pane. The right pane lists the members for each zone.



Note The default zone members are explicitly listed only when the default zone policy is configured as **permit**. When the default zone policy is configured as **deny**, the members of this zone are not shown.

Configuring the Default Zone Access Permission

Detailed Steps

To permit or deny traffic to members in the default zone, follow these steps:

-
- Step 1** Expand a **VSAN** and then select **Default Zone** in the DCNM-SAN Logical Domains pane.
- Step 2** Click the **Policies** tab in the Information pane.
You see the zone policies information in the Information pane.
The active zone set is shown in italic type. After you make changes to the active zone set and before you activate the changes, the zone set is shown in boldface italic type.
- Step 3** In the Default Zone Behaviour field, choose either **permit** or **deny** from the drop-down menu.
-

Creating FC Aliases

Detailed Steps

To create an FC alias, follow these steps:

-
- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Select VSAN dialog box.
- Step 2** Select a VSAN and click **OK**.
You see the Edit Local Full Zone Database dialog box for the selected VSAN.


Send documentation comments to dcnm-san-docfeedback@cisco.com

- Step 3** Click **Aliases** in the lower left pane. The right pane lists the existing aliases.
 - Step 4** Click the **Insert** icon to create an alias.
You see the Create Alias dialog box.
 - Step 5** Set the Alias Name and the pWWN.
 - Step 6** Click **OK** to create the alias.
-

Adding Members to Aliases

Detailed Steps

To add a member to an alias, follow these steps:

- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Select VSAN dialog box.
 - Step 2** Select a VSAN and click **OK**.
You see the Edit Local Full Zone Database dialog box for the selected VSAN.
 - Step 3** Select the member(s) you want to add from the Fabric pane and click **Add to Alias** or click the alias where you want to add members and click the **Insert** icon.
You see the Add Member to Alias dialog box.
- 

Note The Device Alias radio button is visible only if device alias is in enhanced mode. For more information, see [“Creating Device Aliases” section on page 1-7](#).
- Step 4** Click the browse button and select a port name or check the **LUN** check box and click the browse button to configure LUNs.
 - Step 5** Click **Add** to add the member to the alias.
-

Converting Zone Members to pWWN-based Members

You can convert zone and alias members from switch port or FC ID- based membership to pWWN-based membership. You can use this feature to convert to pWWN so that your zone configuration does not change if a card or switch is changed in your fabric.

Detailed Steps

To convert switch port and FC ID members to pWWN members, follow these steps:

- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Select VSAN dialog box.
- Step 2** Select a VSAN and click **OK**.

Send documentation comments to dcnm-san-docfeedback@cisco.com

You see the Edit Local Full Zone Database dialog box for the selected VSAN.

Step 3 Click the zone you want to convert.

Step 4 Choose **Tools > Convert Switch Port/FCID members to By pWWN**.

You see the conversion dialog box, which lists all members that will be converted.

Step 5 Verify the changes and click **Continue Conversion**.

Step 6 Click **Yes** in the confirmation dialog box to convert that member to pWWN-based membership.

Creating Zone Sets and Adding Member Zones

The pWWN of the virtual target does not appear in the zoning end devices database in DCNM-SAN. If you want to zone the virtual device with a pWWN, you must enter it in the Add Member to Zone dialog box when creating a zone. However, if the device alias is in enhanced mode, the virtual device names appear in the device alias database in the DCNM-SAN zoning window. In this case, users can choose to select either the device alias name or enter the pWWN in the Add Member to Zone dialog box.

For more information, see the [“Adding Zone Members” section on page 1-16](#).

Set the device alias mode to **enhanced** when using SDV (because the pWWN of a virtual device could change).

For example, SDV is enabled on a switch and a virtual device is defined. SDV assigns a pWWN for the virtual device, and it is zoned based on the pWWN in a zone. If you later disable SDV, this configuration is lost. If you reenable SDV and create the virtual device using the same name, there is no guarantee that it will get the same pWWN again. You will have to rezone the pWWN-based zone. However, if you perform zoning based on the device-alias name, there are no configuration changes required if or when the pWWN changes.

Be sure you understand how device alias modes work before enabling them. Refer to [Chapter 1, “Distributing Device Alias Services”](#) for details and requirements about device alias modes.



Note

If one zone set is active and you activate another zone set, the currently active zone set is automatically deactivated.



Tip

You do not have to copy the running configuration to the startup configuration to store the active zone set. However, you need to copy the running configuration to the startup configuration to explicitly store full zone sets. It is not available across switch resets.



Caution

If you deactivate the active zone set in a VSAN that is also configured for IVR, the active IVR zone set (IVZS) is also deactivated and all IVR traffic to and from the switch is stopped. This deactivation can disrupt traffic in more than one VSAN. Before deactivating the active zone set, check the active zone analysis for the VSAN. To reactivate the IVZS, you must reactivate the regular zone set (refer to the *Cisco MDS 9000 Family NX-OS Inter-VSAN Routing Configuration Guide*).

Send documentation comments to dcnm-san-docfeedback@cisco.com



Caution

If the currently active zone set contains IVR zones, activating the zone set from a switch where IVR is not enabled disrupts IVR traffic to and from that VSAN. We strongly recommend that you always activate the zone set from an IVR-enabled switch to avoid disrupting IVR traffic.

Filtering Zones, Zone Sets, and Device Aliases Based on Name

Detailed Steps

To filter the zones, zone sets, or device aliases, follow these steps:

-
- Step 1** Click the **Zone** icon in the toolbar (see [Figure 1-6](#)).
 - Step 2** Enter a filter condition, such as `*zo1*`, in the Filter text box.
 - Step 3** Click **Go**.
-

Adding Multiple Zones to Multiple Zone Sets

Detailed Steps

To add multiple zones to multiple zone sets, follow these steps:

-
- Step 1** Click the **Zone** icon in the toolbar (see [Figure 1-6](#)).
 - Step 2** From the tree view, select **Zoneset**.
 - Step 3** Press the **Ctrl** key to select multiple zones.
 - Step 4** Right-click and then select **Add to Zoneset**.
 - Step 5** Press the **Ctrl** key to select multiple zone sets from the pop-up window displayed.
 - Step 6** Click **Add**.
-

Selected zones are added to the selected zone sets.

Enabling Full Zone Set Distribution

All switches in the Cisco MDS 9000 Family distribute active zone sets when new E port links come up or when a new zone set is activated in a VSAN. The zone set distribution takes effect while sending merge requests to the adjacent switch or while activating a zone set.

Detailed Steps

To enable full zone set and active zone set distribution to all switches on a per-VSAN basis, follow these steps:

Send documentation comments to dcnm-san-docfeedback@cisco.com

-
- Step 1** Expand a **VSAN** and select a zone set in the Logical Domains pane.
You see the zone set configuration in the Information pane. The Active Zones tab is the default.
- Step 2** Click the **Policies** tab.
You see the configured policies for the zone.
- Step 3** In the Propagation column, choose **fullZoneset** from the drop-down menu.
- Step 4** Click **Apply Changes** to propagate the full zone set.
-

Enabling a One-Time Distribution

Detailed Steps

You can perform a one-time distribution of inactive, unmodified zone sets throughout the fabric. To propagate a one-time distribution of the full zone set, follow these steps:

-
- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Edit Local Full Zone Database dialog box.
- Step 2** Click the appropriate zone from the list in the left pane.
- Step 3** Click **Distribute** to distribute the full zone set across the fabric.
-

This procedure only distributes the full zone set information; it does not save the information to the startup configuration. You must explicitly save the running configuration to the startup configuration to save the full zone set information to the startup configuration.



Note

The one-time distribution of the full zone set is supported in **interop 2** and **interop 3** modes, not in **interop 1** mode.

Importing and Exporting Zone Sets

Detailed Steps

To import or export the zone set information from or to an adjacent switch, follow these steps:

-
- Step 1** Choose **Tools > Merge Fail Recovery**.
You see the Zone Merge Failure Recovery dialog box.
- Step 2** Click the **Import Active Zoneset** or the **Export Active Zoneset** radio button.
- Step 3** Select the switch from which to import or export the zone set information from the drop-down list.
- Step 4** Select the VSAN from which to import or export the zone set information from the drop-down list.
- Step 5** Select the interface to use for the import process.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 6 Click **OK** to import or export the active zone set.



Note

Issue the **import** and **export** from a single switch. Importing from one switch and exporting from another switch can lead to isolation again.

Copying Zone Sets

On the Cisco MDS Family switches, you cannot edit an active zone set. However, you can copy an active zone set to create a new zone set that you can edit.

Detailed Steps

To make a copy of a zone set, follow these steps:

-
- Step 1** Choose **Zone > Copy Full Zone Database**.
You see the Copy Full Zone Database dialog box.
 - Step 2** Click the **Active** or the **Full** radio button, depending on which type of database you want to copy.
 - Step 3** Select the source VSAN from the drop-down list.
 - Step 4** If you selected **Copy Full**, select the source switch and the destination VSAN from those drop-down lists.
 - Step 5** Select the destination switch from the drop-down list.
 - Step 6** Click **Copy** to copy the database.
-



Caution

If the Inter-VSAN Routing (IVR) feature is enabled and if IVR zones exist in the active zone set, then a zone set copy operation copies all the IVR zones to the full zone database. To prevent copying to the IVR zones, you must explicitly remove them from the full zone set database before performing the copy operation. Refer to the *Cisco MDS 9000 Family NX-OS Inter-VSAN Routing Configuration Guide* for more information on the IVR feature.

Backing Up Zones

Detailed Steps

To back up the full zone configuration, follow these steps:

-
- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Select VSAN dialog box.
 - Step 2** Select a VSAN and click **OK**.
You see the Edit Local Full Zone Database dialog box for the selected VSAN.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Step 3** Choose **File > Backup > This VSAN Zones** to back up the existing zone configuration to a workstation using TFTP, SFTP, SCP, or FTP.
- You see the Backup Zone Configuration dialog box.
- You can edit this configuration before backing up the data to a remote server.
- Step 4** Provide the following Remote Options information to back up data onto a remote server:
- Using**—Select the protocol.
 - Server IP Address**—Enter the IP address of the server.
 - UserName**—Enter the name of the user.
 - Password**—Enter the password for the user.
 - File Name(Root Path)**—Enter the path and the filename.
- Step 5** Click **Backup** or click **Cancel** to close the dialog box without backing up.
-

Restoring Zones

Detailed Steps

To restore the full zone configuration, follow these steps:

-
- Step 1** Choose **Zone > Edit Local Full Zone Database**.
- You see the Select VSAN dialog box.
- Step 2** Select a VSAN and click **OK**.
- You see the Edit Local Full Zone Database dialog box for the selected VSAN.
- Step 3** Choose **File > Restore** to restore a saved zone configuration using TFTP, SFTP, SCP, or FTP.
- You see the Restore Zone Configuration dialog box.
- You can edit this configuration before restoring it to the switch.
- Step 4** Provide the following **Remote Options** information to restore data from a remote server:
- Using**—Select the protocol.
 - Server IP Address**—Enter the IP address of the server.
 - UserName**—Enter the name of the user.
 - Password**—Enter the password for the user.
 - File Name**—Enter the path and the filename.
- Step 5** Click **Restore** to continue or click **Cancel** to close the dialog box without restoring.
-



Note

Click **View Config** to see information on how the zone configuration file from a remote server will be restored. When you click **Yes** in this dialog box, you are provided with the CLI commands that are executed. To close the dialog box, click **Close**.

Send documentation comments to dcnm-san-docfeedback@cisco.com

**Note**

Backup and Restore options are available to switches that run Cisco NX-OS Release 4.1(3) or later.

Renaming Zones, Zone Sets, and Aliases

Detailed Steps

To rename a zone, zone set, or alias, follow these steps:

-
- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Select VSAN dialog box.
 - Step 2** Select a VSAN and click **OK**.
You see the Edit Local Full Zone Database dialog box for the selected VSAN.
 - Step 3** Click a zone or zone set in the left pane.
 - Step 4** Choose **Edit > Rename**.
An edit box appears around the zone or zone set name.
 - Step 5** Enter a new name.
 - Step 6** Click **Activate** or **Distribute**.
-

Cloning Zones, Zone Sets, FC Aliases, and Zone Attribute Groups

Detailed Steps

To clone a zone, zone set, FC alias, or zone attribute group, follow these steps:

-
- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Select VSAN dialog box.
 - Step 2** Select a VSAN and click **OK**.
You see the Edit Local Full Zone Database dialog box for the selected VSAN.
 - Step 3** Choose **Edit > Clone**.
You see the Clone Zoneset dialog box. The default name is the word Clone followed by the original name.
 - Step 4** Change the name for the cloned entry.
 - Step 5** Click **OK** to save the new clone.
The cloned database now appears along with the original database.
-

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

Migrating a Non-MDS Database

Detailed Steps

To use the Zone Migration Wizard to migrate a non-MDS database, follow these steps:

-
- Step 1** Choose **Zone > Migrate Non-MDS Database**.
You see the Zone Migration Wizard.
- Step 2** Follow the prompts in the wizard to migrate the database.
-

Clearing the Zone Server Database

You can clear all configured information in the zone server database for the specified VSAN.

To clear the zone server database, refer to the *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide*.

**Note**

Clearing a zone set only erases the full zone database, not the active zone database.

**Note**

After clearing the zone server database, you must explicitly copy the running configuration to the startup configuration to ensure that the running configuration is used when the switch reboots.

Configuring Zone-Based Traffic Priority

Detailed Steps

To configure the zone priority, follow these steps:

-
- Step 1** Expand a **VSAN** and then select a zone set in the Logical Domains pane.
- Step 2** Click the **Policies** tab in the Information pane.
You see the Zone policy information in the Information pane.
- Step 3** Use the check boxes and drop-down menus to configure QoS on the default zone.
- Step 4** Click **Apply Changes** to save the changes.
-

Configuring Default Zone QoS Priority Attributes

QoS priority attribute configuration changes take effect when you activate the zone set of the associated zone.

Send documentation comments to dcnm-san-docfeedback@cisco.com

**Note**

If a member is part of two zones with two different QoS priority attributes, the higher QoS value is implemented. This situation does not arise in the VSAN-based QoS as the first matching entry is implemented.

Detailed Steps

To configure the QoS priority attributes for a default zone, follow these steps:

-
- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Select VSAN dialog box.
 - Step 2** Select a VSAN and click **OK**.
You see the Edit Local Full Zone Database dialog box for the selected VSAN.
 - Step 3** Choose **Edit > Edit Default Zone Attributes** to configure the default zone QoS priority attributes.
 - Step 4** Check the **Permit QoS Traffic with Priority** check box and set the QoS Priority drop-down menu to **low**, **medium**, or **high**.
 - Step 5** Click **OK** to save these changes.
-

Configuring the Default Zone Policy

Detailed Steps

To permit or deny traffic in the default zone, follow these steps:

-
- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Select VSAN dialog box.
 - Step 2** Select a VSAN and click **OK**.
You see the Edit Local Full Zone Database dialog box for the selected VSAN.
 - Step 3** Choose **Edit > Edit Default Zone Attributes** to configure the default zone QoS priority attributes.
You see the Modify Default Zone Properties dialog box.
 - Step 4** Set the Policy drop-down menu to **permit** to permit traffic in the default zone, or set it to **deny** to block traffic in the default zone.
 - Step 5** Click **OK** to save these changes.
-

Configuring Smart Zoning

Detailed Steps

To configure smart zoning, follow these steps:

Send documentation comments to dcnm-san-docfeedback@cisco.com

-
- Step 1** Expand a **VSAN** and then select a zone set in the Logical Domains pane.
- Step 2** Click the **Smart Zoning** tab in the Information pane.
- You see the smart zoning information in the Information pane.
- Step 3** You can view the details under the **Switch**, **Status**, **Command**, **Last Command**, and **Result** headings.
- Step 4** You can set the **Status**, and **Command** fields.
- Step 5** Click **Apply Changes** to save these changes.
-

Configuring Global Zone Policies

Detailed Steps

To configure global zone policy, follow these steps:

-
- Step 1** In the Logical Domains pane, select **ALL VSANs**.
- Step 2** Click the **Global Zone Policies** tab in the Information pane.
- You see the Global Zone Policy information in the Information pane.
- Step 3** Set the type of switch under the **Switch** column.
- Step 4** You either **Deny** or **Permit** the **Zone Behaviour** and set the **Propagation Mode**.
- Step 5** Select if the **Smart Zoning** feature is enabled or disabled.
- Step 6** Click **Apply Changes** to save these changes.
-

Configuring Broadcast Zoning

Detailed Steps

To broadcast frames in the basic zoning mode, follow these steps:

-
- Step 1** Expand a **VSAN** and then select a zone set in the Logical Domains pane.
- Step 2** Click the **Policies** tab in the Information pane.
- You see the Zone policy information in the Information pane.
- Step 3** Check the **Broadcast** check box to enable broadcast frames on the default zone.
- Step 4** Click **Apply Changes** to save these changes.
-

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

Configuring a LUN-Based Zone

Detailed Steps

To configure a LUN-based zone, follow these steps:

-
- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Select VSAN dialog box.
 - Step 2** Select a VSAN and click **OK**.
You see the Edit Local Full Zone Database dialog box for the selected VSAN.
 - Step 3** Click the zone where you want to add members and click the **Insert** icon.
You see the Add Member to Zone dialog box.
 - Step 4** Click either the **WWN** or **FCID** radio button from the Zone By options to create a LUN-based zone.
 - Step 5** Check the **LUN** check box and click the browse button to configure LUNs.
 - Step 6** Click **Add** to add this LUN-based zone.
-

Assigning LUNs to Storage Subsystems

LUN masking and mapping restricts server access to specific LUNs. If LUN masking is enabled on a storage subsystem and if you want to perform additional LUN zoning in a Cisco MDS 9000 Family switch, obtain the LUN number for each host bus adapter (HBA) from the storage subsystem and then configure the LUN-based zone procedure provided in the [“Configuring a LUN-Based Zone” section on page 1-34](#).



Note

Refer to the relevant user manuals to obtain the LUN number for each HBA.



Caution

If you make any errors when assigning LUNs, you might lose data.

Configuring Read-Only Zones

Detailed Steps

To configure read-only zones, follow these steps:

-
- Step 1** Choose **Zone > Edit Local Full Zone Database**.
You see the Select VSAN dialog box.
 - Step 2** Select a VSAN and click **OK**.
You see the Edit Local Full Zone Database dialog box for the selected VSAN.
 - Step 3** Click **Zones** in the left pane and click the **Insert** icon to add a zone.

Send documentation comments to dcnm-san-docfeedback@cisco.com

You see the Create Zone Dialog Box.

Step 4 Check the **Read Only** check box to create a read-only zone.

Step 5 Click **OK**.



Note

To configure the read-only option for a default zone, see [“Configuring the Default Zone Policy” section on page 1-32](#).

Changing from Basic Zoning to Enhanced Zoning

Detailed Steps

To change to the enhanced zoning mode from the basic mode, follow these steps:

Step 1 Verify that all switches in the fabric are capable of working in the enhanced mode.

If one or more switches are not capable of working in enhanced mode, then your request to move to enhanced mode is rejected.

Step 2 Set the operation mode to enhanced zoning mode.

You will be able to automatically start a session, acquire a fabric wide lock, distribute the active and full zoning database using the enhanced zoning data structures, distribute zoning policies, and then release the lock. All switches in the fabric then move to the enhanced zoning mode.



Tip

After moving from basic zoning to enhanced zoning, we recommend that you save the running configuration.

Changing from Enhanced Zoning to Basic Zoning

The standards do not allow you to move back to basic zoning. However, Cisco MDS switches allow this move to enable you to downgrade and upgrade to other Cisco SAN-OS or Cisco NX-OS releases.

Detailed Steps

To change to the basic zoning mode from the enhanced mode, follow these steps:

Step 1 Verify that the active and full zone set do not contain any configuration that is specific to the enhanced zoning mode.

If such configurations exist, delete them before proceeding with this procedure. If you do not delete the existing configuration, the Cisco NX-OS software automatically removes them.

Step 2 Set the operation mode to basic zoning mode.

Send documentation comments to dcnm-san-docfeedback@cisco.com

You will be able to automatically start a session, acquire a fabric wide lock, distribute the zoning information using the basic zoning data structure, apply the configuration changes, and release the lock from all switches in the fabric. All switches in the fabric then move to basic zoning mode.

**Note**

If a switch running Cisco SAN-OS Release 2.0(1b) and NX-OS 4(1b) or later, with enhanced zoning enabled is downgraded to Cisco SAN-OS Release 1.3(4), or earlier, the switch comes up in basic zoning mode and cannot join the fabric because all the other switches in the fabric are still in enhanced zoning mode.

Enabling Enhanced Zoning

By default, the enhanced zoning feature is disabled in all switches in the Cisco MDS 9000 Family.

Detailed Steps

To enable enhanced zoning in a VSAN, follow these steps:

-
- Step 1** Expand a VSAN and then select a zone set in the Logical Domains pane.
You see the zone set configuration in the Information pane.
 - Step 2** Click the **Enhanced** tab.
You see the current enhanced zoning configuration.
 - Step 3** From the Action drop-down menu, choose **enhanced** to enable enhanced zoning in this VSAN.
 - Step 4** Click **Apply Changes** to save these changes.
-

Modifying the Zone Database

Modifications to the zone database is done within a session. A session is created at the time of the first successful configuration command. On creation of a session, a copy of the zone database is created. Any changes done within the session are performed on this copy of the zoning database. These changes in the copy zoning database are not applied to the effective zoning database until you commit the changes. Once you apply the changes, the session is closed.

If the fabric is locked by another user and for some reason the lock is not cleared, you can force the operation and close the session. You must have permission (role) to clear the lock in this switch and perform the operation on the switch from where the session was originally created.

Creating Attribute Groups

In enhanced mode, you can directly configure attributes using attribute groups.

To configure attribute groups, refer to the *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide*.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Analyzing a Zone Merge

Detailed Steps

To perform a zone merge analysis, follow these steps:

-
- Step 1** Choose **Zone > Merge Analysis**.
You see the Zone Merge Analysis dialog box.
 - Step 2** Select the first switch to be analyzed from the Check Switch 1 drop-down list.
 - Step 3** Select the second switch to be analyzed from the And Switch 2 drop-down list.
 - Step 4** Enter the VSAN ID where the zone set merge failure occurred in the For Active Zoneset Merge Problems in VSAN Id field.
 - Step 5** Click **Analyze** to analyze the zone merge.
 - Step 6** Click **Clear** to clear the analysis data in the Zone Merge Analysis dialog box.
-

Configuring Zone Merge Control Policies

To configure merge control policies, refer to the *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide*.

Broadcasting a Zone

You can specify an enhanced zone to restrict broadcast frames generated by a member in this zone to members within that zone. Use this feature when the host or storage devices support broadcasting.

Table 1-6 identifies the rules for the delivery of broadcast frames.

Table 1-6 **Broadcasting Requirements**

Active Zoning?	Broadcast Enabled?	Frames Broadcast?	Comments
Yes	Yes	Yes	Broadcast to all Nx ports that share a broadcast zone with the source of broadcast frames.
No	Yes	Yes	Broadcast to all Nx ports.
Yes	No	No	Broadcasting is disabled.



Tip

If any NL port attached to an FL port shares a broadcast zone with the source of the broadcast frame, then the frames are broadcast to all devices in the loop.

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

Compacting the Zone Database for Downgrading

Prior to Cisco SAN-OS Release 3.0(1), only 2000 zones are supported per VSAN. If you add more than 2000 zones to a VSAN, a configuration check is registered to indicate that downgrading to a previous release could cause you to lose the zones over the limit. To avoid the configuration check, delete the excess zones and compact the zone database for the VSAN. If there are 2000 zones or fewer after deleting the excess zones, the compacting process assigns new internal zone IDs and the configuration can be supported by Cisco SAN-OS Release 2.x or earlier. Perform this procedure for every VSAN on the switch with more than 2000 zones.



Note

A merge failure occurs when a switch supports more than 2000 zones per VSAN but its neighbor does not. Also, zone set activation can fail if the switch has more than 2000 zones per VSAN and not all switches in the fabric support more than 2000 zones per VSAN.

To compact the zone database for downgrading, refer to the *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide*.

Verifying Zone Configuration

This section contains the following topic(s):

- [Displaying Zone Information, page 1-38](#)

Displaying Zone Information

To view zone information and statistics, follow these steps:


-
- | | |
|---------------|---|
| Step 1 | Expand a VSAN and select a zone set in the Logical Domains pane.
You see the zone configuration in the Information pane. |
| Step 2 | Click the Read Only Violations, Statistics tab or the LUN Zoning Statistics tab to view statistics for the selected zone. |
-

Configuration Examples for Zoning

[Figure 1-9](#) illustrates a zone set with two zones, zone 1 and zone 2, in a fabric. Zone 1 provides access from all three hosts (H1, H2, H3) to the data residing on storage systems S1 and S2. Zone 2 restricts the data on S3 to access only by H3. Note that H3 resides in both zones.


Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 1-9 Fabric with Two Zones



You can partition this fabric into zones using other methods. [Figure 1-10](#) illustrates another possibility. Assume that there is a need to isolate storage system S2 for the purpose of testing new software. To achieve this, zone 3 is configured, which contains only host H2 and storage S2. You can restrict access to just H2 and S2 in zone 3, and to H1 and S1 in zone 1.

Figure 1-10 Fabric with Three Zones



Field Descriptions for Zones

The following are the field descriptions for zoning.

Zone Set Active Zones

Field	Description
Zone	Zone name.
Type	Zone member type.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
Device Type	Specifies if the end device type is host, storage, or both.
Switch Interface	Switch interface to which the zone member is connected to.
Name	Zone member name.
WWN	Zone member WWN.
FcId	Zone member FC ID.
LUNs	Zone member LUN.
Status	<ul style="list-style-type: none"> Not in Fabric: If zone member is not in the fabric. Not in VSAN: If zone member is not present in the VSAN. n/a: Cannot determine status. Empty: Member is present in fabric and correct VSAN and can communicate with other members of the zone.

Zone Set Unzoned

Field	Description
Name	Zone member name.
WWN	Zone member WWN.
FcId	Zone member FC ID.

Zone Set Status

Field	Description
Status	Indicates the outcome of the most recent activation or deactivation.
Activation Time	When this entry was most recently activated. If this entry has been activated prior to the last reinitialization of the local network management system, then this value will be N/A.
FailureCause	The reason for the failure of the zone set activation or deactivation.
FailedSwitch	The domain ID of the device in the fabric that has caused the Change Protocol to fail.
Active == Local?	Indicates whether the enforced database is the same as the local database on this VSAN. If true, then they are the same. If false, then they are not the same.
Active Zoneset	The name of the enforced IV zone set.
Hard Zoning	Indicates whether the hard zoning is enabled on this VSAN. Hard zoning is a mechanism by which zoning is enforced in hardware. If true, then hard zoning is enabled on this VSAN. If false, then hard zoning is not enabled on this VSAN.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Zone Set Policies

Field	Description
Default Zone Behavior	Controls the behavior of the default zone on this VSAN. If it is set to permit, then the members of the default zone on this VSAN can communicate with each other. If it is set to deny, then the members of the default zone on this VSAN cannot communicate with each other.
Default Zone ReadOnly	Indicates whether SCSI read operations are allowed on members of the default zone which are SCSI targets, on this VSAN. If true, then only SCSI read operations are permitted. So, this default zone becomes a read-only default zone on this VSAN. If false, then both SCSI read and write operations are permitted.
Default Zone QoS	Specifies whether the QoS attribute for the default zone on this VSAN is enabled. If true, then QoS attribute for the default zone on this VSAN is enabled. If false, then the QoS attribute for the default zone on this VSAN is disabled.
Default Zone QoS Priority	Specifies the QoS priority value.
Default Zone Broadcast	Specifies if broadcast zoning is enabled on this default zone on this VSAN. If true, then it is enabled. If false, then it is disabled.
Smart Zoning	Specifies if the smart zoning feature is enabled or disabled at the VSAN level
Propagation	Controls the way zoneset information is propagated during Merge/Change protocols on this VSAN
Read From	Specifies whether the management station wishes to read from the effective database or from the copy database.

Zone Set Active Zones Attributes

Field	Description
Name	Zone name.
Read Only	Indicates if only SCSI read operations are allowed on members of the default zone which are SCSI targets on this VSAN. If true, then only SCSI read operations are permitted. So, this default zone becomes a read-only default zone on this VSAN. If false, then both SCSI read and write operations are permitted.
QoS	Specifies whether the QoS attribute for the default zone on this VSAN is enabled. If true, then QoS attribute for the default zone on this VSAN is enabled. If false, then the QoS attribute for the default zone on this VSAN is disabled.
QoS Priority	Specifies QoS priority value (Low, Medium, or High).

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
Broadcast	Specifies if broadcast zoning is enabled on this default zone on this VSAN. If true, then it is enabled. If false, then it is disabled.
Smart Zoning	Specifies if the smart zoning feature is enabled. on this VSAN. If the check box is unchecked, then it is disabled.

Zone Set Enhanced

Field	Description
Action	When set to basic(1), results in the zone server operating in the basic mode as defined by FC-GS4 standards. When set to enhanced(2), results in the zone server operating in the enhanced mode as defined by FC-GS4 standards.
Result	The outcome of setting the mode of operation of the local zone server on this VSAN.
Config DB Locked By	Specifies the owner for this session.
Config DB Discard Changes	Assists in committing or clearing the contents of the copy database on this session.
Config DB Result	Indicates the outcome of setting the corresponding instance of czseSessionCntl to commitChanges(1).
Enforce Full DB Merge	Controls the zone merge behavior. If this object is set to allow, then the merge takes place according to the merge rules. If set to restrict, then if the merging databases are not exactly identical, the Inter-Switch Link (ISL) between the devices is isolated.
Read From	Specifies whether the management station wishes to read from the effective database or from the copy database.

Smart Zoning

Field	Description
Switch	Specifies the type of which where smart zoning feature exists.
Status	Specifies if the smart zoning feature is enabled or disabled.
Command	Specifies the switch level command for smart zoning. If the command is disabled in one switch then smart zoning will be disabled in the whole fabric.
Last Command	Specified the previous command mode of the switch. Enabled or disabled.
Result	Specifies if the enable or disable action has been successful or unsuccessful.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Zone Set Read Only Violations

Field	Description
Violations	The number of data-protected Check Condition error responses sent by the local zone server.

Zone Set Statistics

Field	Description
Merge Req Tx	The number of merge request frames sent by this zone server to other zone servers in the fabric on this VSAN.
Merge Req Rx	The number of merge request frames received by this zone server from other zone servers in the fabric on this VSAN.
Merge Acc Tx	The number of merge accept frames sent by this zone server to other zone servers in the fabric on this VSAN.
Merge Acc Rx	The number of merge accept frames received by this zone server from other zone servers in the fabric on this VSAN.
Change Req Tx	The number of change requests sent by this zone server to other zone servers in the fabric on this VSAN.
Change Req Rx	The number of change requests received by this zone server from other zone servers in the fabric on this VSAN.
Change Acc Tx	The number of change responses sent by this zone server to other zone servers in the fabric on this VSAN.
Change Acc Rx	The number of change responses received by this zone server from other zone servers in the fabric on this VSAN.
GS3 Rej Tx	The number of GS3 requests rejected by this zone server on this VSAN.
GS3 Req Rx	The number of GS3 requests received by this zone server on this VSAN.

Zone Set LUN Zoning Statistics

Field	Description
INQUIRY	The number of SCSI INQUIRY commands that have been received by the local zone server.
REPORT LUN	The number of SCSI Report LUNs commands that have been received by the local zone server. Typically the Report LUNs command is sent only for LUN 0.
SENSE	The number of SCSI SENSE commands that have been received by the local zone server.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
Other Cmds	The number of SCSI Read, Write, Seek commands received by the local zone server.
BadInquiry Errors	The number of No LU error responses sent by the local zone server.
Illegal Errors	The number of Illegal Request Check Condition responses sent by the local zone server.

Zone Set Members

Field	Description
Zone	Default zone.
Type	FCID.
Device Type	Specifies if the end device type is host, storage, or both.
Switch Interface	Switch interface to which the zone member is connected to.
Name	Zone member name.
WWN	Zone member WWN.
FcId	Zone member FC ID.
Luns	Zone member LUN.
Status	<ul style="list-style-type: none"> Not in Fabric: If zone member is not in the fabric. Not in VSAN: If zone member is not present in the VSAN. n/a: Cannot determine status. <p>Empty: Member is present in fabric and correct VSAN and can communicate with other members of the zone.</p>

Additional References

For additional information related to implementing VSANs, see the following section:

- [Related Document, page 1-45](#)
- [Standards, page 1-45](#)
- [RFCs, page 1-45](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- [MIBs, page 1-45](#)

Related Document

Related Topic	Document Title
Cisco MDS 9000 Family Command Reference	<i>Cisco MDS 9000 Family Command Reference</i>

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified.	—

MIBs

MIBs	MIBs Link
<ul style="list-style-type: none">• CISCO-ZS-EXT-MIB• CISCO-ZS-MIB	For more information, go to the following URL: http://www.cisco.com/en/US/products/ps5989/prod_technical_reference_list.html .

Send documentation comments to dcnm-san-docfeedback@cisco.com



CHAPTER 1

Distributing Device Alias Services

All switches in the Cisco MDS 9000 Family support Distributed Device Alias Services (device alias) on a per-VSAN basis and on a fabric-wide basis. Device alias distribution allows you to move host bus adapters (HBAs) between VSANs without manually reentering alias names.

This chapter includes the following topics:

- [Information About Device Aliases, page 1-1](#)
- [Guidelines and Limitations, page 1-6](#)
- [Default Settings, page 1-6](#)
- [Configuring Device Aliases, page 1-6](#)
- [Field Descriptions for Device Aliases, page 1-9](#)
- [Additional References, page 1-10](#)

Information About Device Aliases

When the port WWN (pWWN) of a device must be specified to configure different features (zoning, QoS, port security) in a Cisco MDS 9000 Family switch, you must assign the correct device name each time you configure these features. An incorrect device name can cause unexpected results. You can avoid this problem if you define a user-friendly name for a port WWN and use this name in all of the configuration commands as required. These user-friendly names are referred to as *device aliases* in this chapter.

- [About Device Alias Modes, page 1-2](#)
- [Changing Mode Settings, page 1-2](#)
- [Device Alias Mode Distribution, page 1-2](#)
- [Merging Device Alias, page 1-3](#)
- [Resolving Merge and Device Alias Mode Mismatch, page 1-3](#)
- [Device Alias Features, page 1-3](#)
- [Device Alias Requirements, page 1-4](#)
- [Zone Aliases Versus Device Aliases, page 1-4](#)
- [Device Alias Databases, page 1-4](#)
- [About Device Alias Distribution, page 1-5](#)
- [About Creating a Device Alias, page 1-5](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- [Fabric Lock Override, page 1-5](#)
- [About Legacy Zone Alias Configuration Conversion, page 1-5](#)

About Device Alias Modes

Device alias supports two modes: basic and enhanced mode.

- When device alias runs in the basic mode, all applications function like the applications on the Cisco SAN-OS Release 3.0 switches. When you configure the basic mode using device aliases, the application immediately expands to pWWNs. This operation continues until the mode is changed to enhanced.
- When device alias runs in the enhanced mode, all applications accept the device-alias configuration in the native format. The applications store the device alias name in the configuration and distribute it in the device alias format instead of expanding to pWWN. The applications track the device alias database changes and take actions to enforce it.

A native device-alias configuration is not accepted in the interop mode VSAN. IVR zoneset activation will fail in interop mode VSANs if the corresponding twilight zones being injected are native device alias members.

Changing Mode Settings

When the device alias mode is changed from basic to enhanced mode, the applications are informed about the change. The applications start accepting the device alias-based configuration in the native format.



Note

Because the device alias was previously running in the basic mode, the applications do not have any prior native device alias configuration.

The applications check for an existing device alias configuration in the native format. If the device alias is in the native format, the applications reject the request and device alias mode cannot be changed to basic.

All native device alias configurations (both on local and remote switches) must be explicitly removed, or all device alias members must be replaced with the corresponding pWWN before changing the mode back to basic.

Device Alias Mode Distribution

If the device alias distribution is turned on, it is distributed to the other switches in the network whenever there is a change in the mode. You cannot change the mode from basic to enhanced unless all the switches are upgraded to Cisco SAN-OS Release 3.1. The device alias enhancements will not apply unless the entire fabric is upgraded to Cisco SAN-OS Release 3.1.



Note

When all the switches are upgraded to Cisco SAN-OS Release 3.1, you cannot automatically convert to enhanced mode. You do not need to change to enhanced mode, you can continue working in the basic mode.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Merging Device Alias

If two fabrics are running different device alias modes and are joined together, the device alias merge will fail. There is no automatic conversion of one mode to the other during the merge process. You will need to resolve the issue.

**Note**

Release 3.0 switches run in basic mode.

At the application level, a merger takes place between the applications and the fabric. For example, zone merge occurs when the E port is up and the IVR/PSM/DPVM merge occurs due to CFS. This merge is completely independent of the device alias merge.

If the application running on an enhanced fabric has a native device alias configuration, the application must fail the merge. The application has to fail the merge even though the other fabric is can support the native device alias-based configuration, but is running in the basic mode. You will need to resolve the issue. Once the device alias merge issue is resolved, each application must be fixed accordingly.

Resolving Merge and Device Alias Mode Mismatch

If two fabrics are running in different modes and the device alias merge fails between the fabrics, the conflict can be resolved by selecting one mode or the other. If you choose the enhanced mode, ensure that all the switches are running at least the Cisco SAN-OS Release 3.1. Otherwise, the enhanced mode cannot be turned on. If you choose the basic mode, the applications running on the enhanced fabric have to comply with the device alias merge.

The device alias merge fails because of mode mismatch, but the application merge succeeds if it does not have any native device alias configurations.

If the native device alias configuration is attempted on an application from a Release 3.1 switch, the commit must be rejected because of device alias mode mismatch on some of the applications.

**Note**

The applications should not accept any native device alias configuration over SNMP if the device alias is running in the basic mode on that particular switch.

**Note**

Confcheck is added when the enhanced mode is turned on and removed when it is turned off. Applications have to add confcheck if they have a device alias configuration in the native format. They have to remove confcheck once the configuration is removed.

Device Alias Features

Device aliases have the following features:

- The device alias information is independent of your VSAN configuration.
- The device alias configuration and distribution is independent of the zone server and the zone server database.
- You can import legacy zone alias configurations without losing data.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- The device alias application uses the Cisco Fabric Services (CFS) infrastructure to enable efficient database management and distribution. Device aliases use the coordinated distribution mode and the fabric-wide distribution scope (refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*).
- When you configure zones, IVR zones, or QoS features using device aliases, and if you display these configurations, the device aliases are automatically displayed along with their respective pWWNs.

Device Alias Requirements

Device aliases have the following requirements:

- You can only assign device aliases to pWWNs.
- The mapping between the pWWN and the device alias to which it is mapped must have a one-to-one relationship. A pWWN can be mapped to only one device alias and vice versa.
- A device alias name is restricted to 64 alphanumeric characters and may include one or more of the following characters:
 - a to z and A to Z
 - 1 to 9
 - - (hyphen) and _ (underscore)
 - \$ (dollar sign) and ^ (up caret)

Zone Aliases Versus Device Aliases

[Table 1-1](#) compares the configuration differences between zone-based alias configuration and device alias configuration.

Table 1-1 Comparison Between Zone Aliases and Device Aliases

Zone-Based Aliases	Device Aliases
Aliases are limited to the specified VSAN.	You can define device aliases without specifying the VSAN number. You can also use the same definition in one or more VSANs without any restrictions.
Zone aliases are part of the zoning configuration. The alias mapping cannot be used to configure other features.	Device aliases can be used with any feature that uses the pWWN.
You can use any zone member type to specify the end devices.	Only pWWNs are supported along with new device aliases such as IP addresses.
Configuration is contained within the Zone Server database and is not available to other features.	Device aliases are not restricted to zoning. Device alias configuration is available to the FCNS, zone, fcping, traceroute, and IVR applications.

Device Alias Databases

The device alias feature uses two databases to accept and implement device alias configurations:

- Effective database—The database currently used by the fabric.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Pending database—Your subsequent device alias configuration changes are stored in the pending database.

If you modify the device alias configuration, you need to commit or discard the changes as the fabric remains locked during this period.

About Device Alias Distribution

By default, device alias distribution is enabled. The device alias feature uses the coordinated distribution mechanism to distribute the modifications to all switches in a fabric.

If you have not committed the changes and you disable distribution, then a commit task will fail.

About Creating a Device Alias

When you perform the first device alias task (regardless of which device alias task), the fabric is automatically locked for the device alias feature. Once you lock the fabric, the following situations apply:

- No other user can make any configuration changes to this feature.
- A copy of the effective database is obtained and used as the pending database. Modifications from this point on are made to the pending database. The pending database remains in effect until you commit the modifications to the pending database or discard (**abort**) the changes to the pending database.

Fabric Lock Override

If you have performed a device alias task and have forgotten to release the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your changes to the pending database are discarded and the fabric lock is released.



Tip

The changes are only available in the volatile directory and are subject to being discarded if the switch is restarted.

About Legacy Zone Alias Configuration Conversion

You can import legacy zone alias configurations to use this feature without losing data if they follow the following restrictions:

- Each zone alias has only one member.
- The member type is pWWN.
- The name and definition of the zone alias should not be the same as any existing device alias name.

If any name conflict exists, the zone aliases are not imported.

Send documentation comments to dcnm-san-docfeedback@cisco.com



Tip

Ensure that you copy any required zone aliases to the device alias database as required by your configuration.

When an import operation is complete, the modified alias database is distributed to all other switches in the physical fabric when you perform the **commit** operation. At this time if you do not want to distribute the configuration to other switches in the fabric, you can perform the **abort** operation and the merge changes are completely discarded.

Guidelines and Limitations

This section explains the database guidelines for this feature.

For information about CFS merge support, refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide* for detailed concepts.

When merging two device alias databases, follow these guidelines:

- Verify that two device aliases with different names are not mapped to the same pWWN.
- Verify that two different pWWNs are not mapped to the same device aliases
- Verify that the combined number of the device aliases in both databases does not exceed 8191 (8K). For example, if database N has 6000 device aliases and database M has 2192 device aliases, this merge operation will fail.

Default Settings

Table 1-2 lists the default settings for device alias parameters.

Table 1-2 Default Device Alias Parameters

Parameters	Default
Database in use	Effective database.
Database to accept changes	Pending database.
Device alias fabric lock state	Locked with the first device alias task.

Configuring Device Aliases

This section includes the following topics:

- [Creating Device Aliases, page 1-7](#)
- [Distributing the Device Alias Database, page 1-7](#)
- [Committing Changes, page 1-7](#)
- [Discarding Changes, page 1-8](#)
- [Using Device Aliases or FC Aliases, page 1-8](#)
- [Populating Device Alias to Interface Description, page 1-9](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

Creating Device Aliases

Detailed Steps

To lock the fabric and create a device alias in the pending database, follow these steps:

-
- | | |
|---------------|---|
| Step 1 | Expand End Devices , and then select Device Alias in the Physical Attributes pane.
You see the device alias configuration in the Information pane. |
| Step 2 | Click the Configuration tab and click the Create Row icon.
You see the Device Alias Creation dialog box. |
| Step 3 | Select a switch from the drop-down menu. |
| Step 4 | Complete the Alias name and pWWN fields. |
| Step 5 | Click Create to create this alias or click Close to discard any unsaved changes. |
-

Distributing the Device Alias Database

Detailed Steps

To enable the device alias distribution, follow these steps:

-
- | | |
|---------------|---|
| Step 1 | Expand End Devices and then select Device Alias in the Physical Attributes pane.
You see the device alias configuration in the Information pane.
The CFS tab is the default tab. |
| Step 2 | Select enable from the Global drop-down menus to enabled switch aliases. |
| Step 3 | Select commit from the Config Action drop-down menu for the newly enabled switches. |
| Step 4 | Click Apply Changes to commit and distribute these changes or click Undo Changes to discard any unsaved changes. |
-

Committing Changes

If you commit the changes made to the pending database, the following events occur:

1. The pending database contents overwrites the effective database contents.
2. The pending database is emptied of its contents.
3. The fabric lock is released for this feature.

Detailed Steps

To commit the changes to the device alias database, follow these steps:

Send documentation comments to dcnm-san-docfeedback@cisco.com

-
- Step 1** Expand **End Devices**, and then select **Device Alias** in the Physical Attributes pane.
You see the device alias configuration in the Information pane. The **CFS** tab is the default tab.
- Step 2** Select **enable** from the Global drop-down menus to enable switch aliases.
- Step 3** Select **commit** from the Config Action drop-down menu for the newly enabled switches.
- Step 4** Click **Apply Changes** to commit and distribute these changes or click **Undo Changes** to discard any unsaved changes.
-

Discarding Changes

If you discard the changes made to the pending database, the following events occur:

1. The effective database contents remain unaffected.
2. The pending database is emptied of its contents.
3. The fabric lock is released for this feature.

Detailed Steps

To discard the device alias session, follow these steps:

-
- Step 1** Expand **End Devices** and then select **Device Alias** in the Physical Attributes pane.
You see the device alias configuration in the Information pane. The **CFS** tab is the default tab.
- Step 2** Select **abort** from the Config Action drop-down menu.
- Step 3** Click **Apply Changes** to discard the session.
-

Using Device Aliases or FC Aliases

You can change whether DCNM-SAN uses FC aliases or global device aliases from DCNM-SAN Client without restarting Cisco DCNM for SAN.

Detailed Steps

To change whether DCNM-SAN uses FC aliases or global device aliases, follow these steps:

-
- Step 1** Click **Server > Admin**.
You see the Admin dialog box.
- Step 2** For each fabric that you are monitoring with Cisco DCNM for SAN, check the **Device Alias** check box to use global device aliases, or uncheck to use FC aliases.
- Step 3** Click **Apply** to save these changes or click **Close** to exit the dialog box without saving any changes.
-

Send documentation comments to dcnm-san-docfeedback@cisco.com

Populating Device Alias to Interface Description

When an end device is not logged into the switch, the Device Alias is blank. To find out what device is supposed to connect to an FC port when the device is logged out, you can populate the interface description with the device alias when the devices are logged in.

Detailed Steps

To populate the interface description with the device alias, follow these steps:

-
- Step 1** From the Physical Attributes pane, expand End Devices.
 - Step 2** From the right pane, click the **General** tab.
 - Step 3** Select the rows of FC interfaces.
 - Step 4** Click the **Alias->Description** button.
 - Step 5** Click the **commit** button.
-

Field Descriptions for Device Aliases

This section displays the field descriptions for this feature.

Device Alias Configuration

Field	Description
Device Alias	The device alias of this entry. A device can have only one alias configured.
WWN	The Fibre Channel device which is given a device alias.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Device Alias Mode

Field	Description
ConfigMode	Specifies the mode in which the device aliases can be configured. When it is set to basic, the device aliases operate in basic mode of operation. When basic mode is turned on, all MIBs which are using device aliases should internally convert them to their equivalent pWWNs and use the pWWNs. The mechanism to be followed for this conversion is implementation specific. When it is set to enhanced, the Device aliases operate in enhanced mode of operation. When enhanced mode is turned on, all MIBs which are using device aliases should use them as is without any conversion. Since the device aliases are used directly without any conversion, this is the native mode of operation of device aliases.

Device Alias Discrepancies

Field	Description
Discrepancy	Represents the checksum computed over the database represented by cfdaConfigTable and the cfdaConfigMode object. This object is used by a network manager to check if the above mentioned objects have changed on the local device. The method used to compute the checksum is implementation specific.

Additional References

For additional information related to implementing VSANs, see the following section:

- [Related Document, page 1-11](#)
- [Standards, page 1-11](#)
- [RFCs, page 1-11](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- [MIBs, page 1-11](#)

Related Document

Related Topic	Document Title
Cisco MDS 9000 Family Command Reference	<i>Cisco MDS 9000 Family Command Reference</i>

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified.	—

MIBs

MIBs	MIBs Link
<ul style="list-style-type: none">CISCO-FC-DEVICE-ALIAS-MIB	To locate and download MIBs, go to the following URL: http://www.cisco.com/en/US/products/ps5989/prod_technical_reference_list.html

Send documentation comments to dcnm-san-docfeedback@cisco.com



CHAPTER 1

Configuring FCoE

This chapter describes how to configure Fibre Channel over Ethernet (FCoE) on a Cisco Nexus 5000 Series Switch, Cisco Nexus 7000 Series Switch, and Cisco 9000 Family MDS switch.

This chapter includes the following sections:

- [About FCoE, page 1-1](#)
- [Guidelines and Limitations, page 1-1](#)
- [Configuring FCoE, page 1-1](#)
- [Field Descriptions for FCoE, page 1-4](#)
- [Additional References, page 1-5](#)

About FCoE

Cisco Nexus 5000 Series Switch, Cisco Nexus 7000 Series Switch, and Cisco MDS 9000 family switches support Fibre Channel over Ethernet (FCoE), which allows Fibre Channel and Ethernet traffic to be carried on the same physical Ethernet connection between the switch and the servers. FCoE requires the underlying Ethernet to be full duplex and to provide lossless behavior for Fibre Channel traffic.

The FCoE Initialization Protocol (FIP) allows the switch to discover and initialize FCoE-capable entities that are connected to an Ethernet LAN.

Guidelines and Limitations

When configuring FCoE, note the following guidelines and limitations:

- FCoE is supported on 10-Gigabit Ethernet interfaces.
- FCoE is not supported on private VLANs.
- DPVM supports MAC-based device mapping for FCoE devices. DPVM does not support pWWN mapping for FCoE devices.

Configuring FCoE

This section describes how to configure FCoE on a switch and includes the following topics:

- [Enabling FCoE, page 1-2](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- [Configuring FCoE Using DCNM for SAN, page 1-3](#)
- [Configuring FCoE Using Device Manager, page 1-4](#)

Enabling FCoE

Fibre Channel over Ethernet (FCoE) provides a method of transporting Fibre Channel traffic over a physical Ethernet connection. By default, each Ethernet interface attempts to enable FCoE by advertising that it has FCoE to the adapter. If the FCoE negotiation fails, you can configure the Cisco Nexus 5000 Series switch to disable FCoE for this interface.



Note

In Cisco Nexus 5000 Series switches, FCoE is supported on all 10-Gigabit Ethernet interfaces.

Detailed Steps

To enable or disable FCoE features on a switch using Device Manager, follow these steps:

Step 1 Launch Device Manager from the Cisco Nexus 5000 Series switch.



Note

Use the Control tab to enable FCoE on a Cisco Nexus 5000 Series switch.

Step 2 Choose **Admin > Feature Control**.

You see the Feature Control dialog box.



Note

You cannot enable FCoE using Device Manager on Cisco Nexus 7000 series and Cisco MDS 900 family switches. Cisco Nexus 7000 series and Cisco MDS 9000 Family switches uses a feature set to display FCoE information.

Step 3 In the dialog box, in the table, click the **fcoe_mgr** row, and then click the **Action** cell in the fcoe_mgr row. From the drop-down list, choose **enable** to enable the FCoE feature in the switch.



Note

You can also disable the FCoE feature in the switch. To do so, from the drop-down list in the Action column, choose **disable**.

Step 4 Click **Apply**.



Note

If the Cisco Nexus 5000 Series switch is running a Cisco NX-OS release prior to Release 4.2(1), you must do the following after you enable or disable FCoE on the switch:

- In the confirmation dialog box that appears, click **Yes** to enable the FCoE feature in the switch.
- Reboot the switch before you use the FCoE feature.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuring FCoE Using DCNM for SAN

From Cisco NX-OS Release 5.2, FCoE is supported on MDS and Cisco Nexus 7000 switches. To enable or disable FCoE, Cisco MDS 9000 Family and Cisco Nexus 7000 switches use feature set MIBs.

Detailed Steps

To configure FCoE on a switch, follow these steps:

Step 1 In the Physical Attributes pane, choose **Switches > FC Services > FCoE**.

You see the FCoE information pane shown in [Figure 1-1](#).

The Config tab displays the FCoE parameters for each Cisco Nexus 5000 Series, Cisco Nexus 7000 Series, and Cisco MDS 9000 Family switches. Use the VLAN-VSAN mapping tab to create mappings. [Table 1-1](#) lists the FCoE parameters for a switch.

For more information on configuring Cisco Nexus 5000 Series and Nexus 7000 Series switches, see the Cisco Nexus 5000 Series and Nexus 7000 Series Configuration Guides.

Figure 1-1 FCoE Information Pane

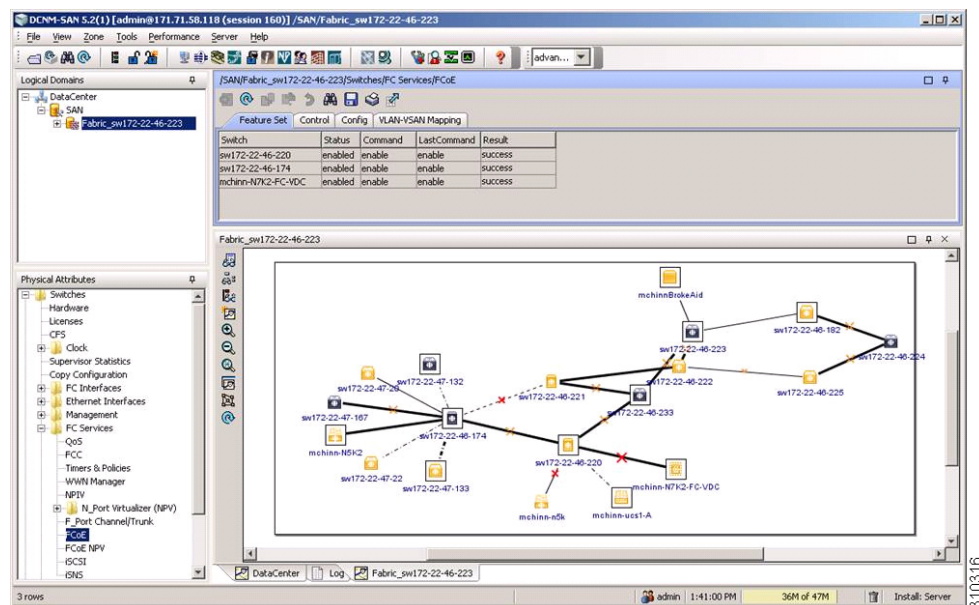


Table 1-1 FCoE Parameters

Parameter	Description
Feature Set	Enables or disables the FCoE feature set on Cisco MDS 9000 Family or Nexus 7000 Series switches.
Control	Enables or disables FCoE on Cisco Nexus 5000 Series switches.
Config	Displays the FCoE configuration information on the switch. For example, FC Map and FCF Priority.
VLAN-VSAN Mapping	Displays the VSAN and VLAN IDs with their operational status.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Step 2** Double-click the relevant FCoE parameter for a switch, and modify the value of the parameter.
- Step 3** In the Information pane toolbar, click the **Apply Changes** icon to save the changes.

Configuring FCoE Using Device Manager

Detailed Steps

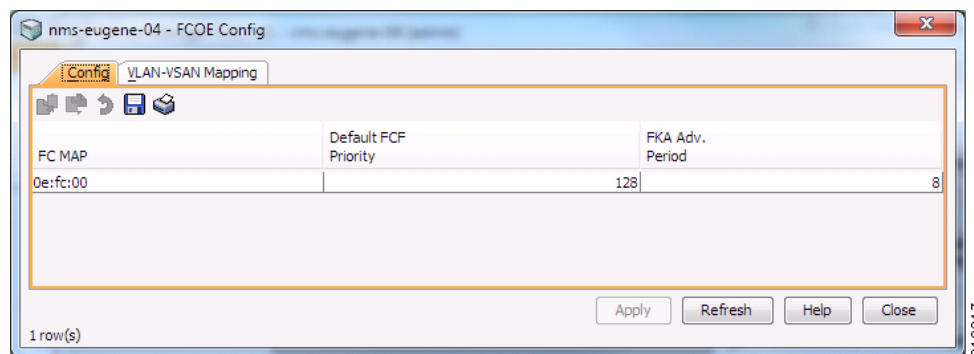
To configure FCoE on a switch using Device Manager, follow these steps:

- Step 1** Launch Device Manager from the Cisco Nexus 5000 Series switch.
- Step 2** Choose **FCoE > Config**.

You see the FCoE Config dialog box shown in [Figure 1-2](#).

The Config tab displays the FCoE parameters, such as FC Map, default FCF priority value, and FKA advertisement period, for each Cisco Nexus 5000 Series switch that runs Cisco NX-OS Release 4.1(3) or later releases. [Table 1-1](#) lists the FCoE parameters for a switch.

Figure 1-2 FCoE Config Dialog Box



- Step 3** Double-click the relevant FCoE parameter for a switch, and modify the value of the parameter.
- Step 4** Click **Apply** to save the changes.

Field Descriptions for FCoE

Feature Set

Field	Description
Status	<i>Display only.</i> Displays the FCoE status on the switch.
Command	<i>Display only.</i> Displays the feature set command.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
Last Command	<i>Display only.</i> Displays the last feature set command executed on the switch.
Result	<i>Display only.</i> Displays the result of the last feature set command executed on the switch.

Control

Field	Description
Status	<i>Display only.</i> Displays the FCoE status on the switch.
Command	<i>Display only.</i> Displays the feature set command.
Last Command	<i>Display only.</i> Displays the last feature set command executed on the switch.
Result	<i>Display only.</i> Displays the result of the last feature set command executed on the switch.

Config

Field	Description
FC Map	The FCoE Mac Address Prefix used to associate the FCoE Node (ENode).
Default FCF Priority	The default FCoE Initialization Protocol (FIP) priority value advertised by the Fibre Channel Forwarder (FCF) to ENodes.
FKA Adv. Period	The time interval at which FIP Keepalive (FKA) messages are transmitted to the MAC address of the ENode.

VSAN-VLAN Mapping

Field	Description
VSAN Id	The ID of the VSAN.
VLAN Id	The ID of the VLAN.
Oper State	Shows the operational state of this VLAN-VSAN association entry.

Additional References

For additional information related to implementing FCoE, see the following section:

- [Related Document, page 1-6](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- [Standards, page 1-6](#)
- [MIBs, page 1-6](#)
- [RFCs, page 1-7](#)

Related Document

Related Topic	Document Title
Cisco MDS 9000 Family Command Reference	<i>Cisco MDS 9000 Family Command Reference</i>

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIB	MIBs Link
No new or modified MIBs are supported by this feature, and support for existing MIBs has not been modified by this feature.	To locate and download MIBs, go to the following URL: http://www.cisco.com/en/US/products/ps5989/prod_technical_reference_list.html

Send documentation comments to dcnm-san-docfeedback@cisco.com

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified.	—

Feature History for FCoE

[Table 1-2](#) lists the release history for this feature. Only features that were introduced or modified in 5.0(1a) or a later release appear in the table.

Table 1-2 ***Feature History for FCoE***

Feature Name	Releases	Feature Information
Configuring FCoE	5.2(1)	<p>Added information about discovering Cisco Nexus 7000 and Cisco MDS 9000 Family switches using the FCoE wizard. FICON tape acceleration over FCIP efficiently utilizes the tape device by decreasing idle time.</p> <p>The following sections provide information about this feature:</p> <ul style="list-style-type: none">• Enabling FCoE, page 1-2• Configuring FCoE Using DCNM for SAN, page 1-3• Configuring FCoE Using Device Manager, page 1-4

Send documentation comments to dcnm-san-docfeedback@cisco.com



CHAPTER 1

Configuring Fibre Channel Routing Services and Protocols

This chapter describes Fibre Channel routing services and protocols.

This chapter includes the following sections:

- [Information About FSPF, page 1-1](#)
- [Licensing Requirements for FSPF, page 1-8](#)
- [Default Settings, page 1-8](#)
- [Configuring FSPF, page 1-9](#)
- [Verifying FSPF Configuration, page 1-14](#)
- [Configuration Examples for FSPF, page 1-15](#)
- [Field Descriptions for FSPF, page 1-17](#)
- [Additional References, page 1-21](#)

Information About FSPF

Fabric Shortest Path First (FSPF) is the standard path selection protocol used by Fibre Channel fabrics. The FSPF feature is enabled by default on all Fibre Channel switches. Except in configurations that require special consideration, you do not need to configure any FSPF services. FSPF automatically calculates the best path between any two switches in a fabric. FSPF provides these features:

- Dynamically compute routes throughout a fabric by establishing the shortest and quickest path between any two switches.
- Select an alternative path in the event of the failure of a given path. FSPF supports multiple paths and automatically computes an alternative path around a failed link. It provides a preferred route when two equal paths are available.

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

- Supports multipath routing.
- Bases path status on a link state protocol.
- Routes hop by hop, based only on the domain ID.
- Runs only on E ports or TE ports and provides a loop free topology.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Runs on a per VSAN basis. Connectivity in a given VSAN in a fabric is guaranteed only for the switches configured in that VSAN.
- Uses a topology database to keep track of the state of the links on all switches in the fabric and associates a cost with each link.
- Guarantees a fast reconvergence time in case of a topology change. Uses the standard Dijkstra algorithm, but there is a static dynamic option for a more robust, efficient, and incremental Dijkstra algorithm. The reconvergence time is fast and efficient as the route computation is done on a per VSAN basis.

This section includes the following topics:

- [FSPF Global Configuration, page 1-2](#)
- [About SPF Computational Hold Times, page 1-3](#)
- [About Link State Record Defaults, page 1-3](#)
- [About FSPF Link Cost, page 1-3](#)
- [About Hello Time Intervals, page 1-3](#)
- [About Dead Time Intervals, page 1-3](#)
- [About Retransmitting Intervals, page 1-4](#)
- [About Disabling FSPF for Specific Interfaces, page 1-4](#)
- [FSPF Routes, page 1-4](#)
- [About Fibre Channel Routes, page 1-4](#)
- [About Broadcast and Multicast Routing, page 1-6](#)
- [About Multicast Root Switch, page 1-6](#)
- [In-Order Delivery, page 1-6](#)
- [About Reordering Network Frames, page 1-6](#)
- [About Reordering PortChannel Frames, page 1-7](#)
- [About Enabling In-Order Delivery, page 1-8](#)
- [About Flow Statistics, page 1-8](#)

FSPF Global Configuration

By default, FSPF is enabled on switches in the Cisco MDS 9000 Family.

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



Note

FSPF is enabled by default. Generally, you do not need to configure these advanced features.



Caution

The default for the backbone region is 0 (zero). You do not need to change this setting unless your region is different from the default. If you are operating with other vendors using the backbone region, you can change this default to be compatible with those settings.

Send documentation comments to dcnm-san-docfeedback@cisco.com

About SPF Computational Hold Times

The SPF computational hold time sets the minimum time between two consecutive SPF computations on the VSAN. Setting this to a small value means that FSPF reacts faster to any fabric changes by recomputing paths on the VSAN. A small SPF computational hold time uses more switch CPU time.

About Link State Record Defaults

Each time a new switch enters the fabric, a link state record (LSR) is sent to the neighboring switches, and then flooded throughout the fabric. [Table 1-1](#) displays the default settings for switch responses.

Table 1-1 LSR Default Settings

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

The LSR minimum arrival time is the period between receiving LSR updates on this VSAN. Any LSR updates that arrive before the LSR minimum arrival time are discarded.

The LSR minimum interval time is the frequency at which this switch sends LSR updates on a VSAN.

About FSPF Link Cost

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

About Hello Time Intervals

You can set the FSPF Hello time interval to specify the interval between the periodic hello messages sent to verify the health of the link. The integer value can range from 1 to 65,535 seconds.



Note

This value must be the same in the ports at both ends of the ISL.

About Dead Time Intervals

You can set the FSPF dead time interval to specify the maximum interval for which a hello message must be received before the neighbor is considered lost and removed from the database. The integer value can range from 1 to 65,535 seconds.

Send documentation comments to dcnm-san-docfeedback@cisco.com

**Note**

This value must be the same in the ports at both ends of the ISL.

**Caution**

An error is reported at the command prompt if the configured dead time interval is less than the hello time interval.

About Retransmitting Intervals

You can specify the time after which an unacknowledged link state update should be transmitted on the interface. The integer value to specify retransmit intervals can range from 1 to 65,535 seconds.

**Note**

This value must be the same on the switches on both ends of the interface.

About Disabling FSPF for Specific Interfaces

You can disable the FSPF protocol for selected interfaces. By default, FSPF is enabled on all E ports and TE ports. This default can be disabled by setting the interface as passive.

**Note**

FSPF must be enabled at both ends of the interface for the protocol to work.

FSPF Routes


FSPF routes traffic across the fabric, based on entries in the FSPF database. These routes can be learned dynamically, or configured statically.

About Fibre Channel Routes

Each port implements forwarding logic, which forwards frames based on its FC ID. Using the FC ID for the specified interface and domain, you can configure the specified route (for example FC ID 111211 and domain ID 3) in the switch with domain ID 1 (see [Figure 1-1](#)).

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 1-1 Fibre Channel Routes



Note

Other than in VSANs, runtime checks are not performed on configured and suspended static routes.

Send documentation comments to dcnm-san-docfeedback@cisco.com

About Broadcast and Multicast Routing

Broadcast and multicast in a Fibre Channel fabric uses the concept of a distribution tree to reach all switches in the fabric.

FSPF provides the topology information to compute the distribution tree. Fibre Channel defines 256 multicast groups and one broadcast address for each VSAN. Switches in the Cisco MDS 9000 Family only use broadcast routing. By default, they use the principal switch as the root node to derive a loop-free distribution tree for multicast and broadcast routing in a VSAN.

**Caution**

All switches in the fabric should run the same multicast and broadcast distribution tree algorithm to ensure the same distribution tree.

To interoperate with other vendor switches (following FC-SW3 guidelines), the Cisco SAN-OS and Cisco NX-OS Release 4.1(1b) and later releases uses the lowest domain switch as the root to compute the multicast tree in interop mode.

About Multicast Root Switch

By default, the native (non-interop) mode uses the principal switch as the root. If you change the default, be sure to configure the same mode in all switches in the fabric. Otherwise, multicast traffic could encounter potential loop and frame-drop problems.

**Note**

The operational mode can be different from the configured interop mode. The interop mode always uses the lowest domain switch as the root.

In-Order Delivery

In-order delivery (IOD) of data frames guarantees frame delivery to a destination in the same order that they were sent by the originator.

Some Fibre Channel protocols or applications cannot handle out-of-order frame delivery. In these cases, switches in the Cisco MDS 9000 Family preserve frame ordering in the frame flow. The source ID (SID), destination ID (DID), and optionally the originator exchange ID (OX ID) identify the flow of the frame.

On any given switch with IOD enabled, all frames received by a specific ingress port and destined to a certain egress port are always delivered in the same order in which they were received.

Use IOD only if your environment cannot support out-of-order frame delivery.

**Tip**


If you enable the in-order delivery feature, the graceful shutdown feature is not implemented.

About Reordering Network Frames

When you experience a route change in the network, the new selected path may be faster or less congested than the old route.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 1-2 Route Change Delivery



In [Figure 1-2](#), the new path from Switch 1 to Switch 4 is faster. In this scenario, Frame 3 and Frame 4 may be delivered before Frame 1 and Frame 2.


If the in-order guarantee feature is enabled, the frames within the network are treated as follows:

- Frames in the network are delivered in the order in which they are transmitted.
- Frames that cannot be delivered in order within the network latency drop period are dropped inside the network.

About Reordering PortChannel Frames

When a link change occurs in a PortChannel, the frames for the same exchange or the same flow can switch from one path to another faster path.

Figure 1-3 Link Congestion Delivery



In [Figure 1-3](#), the port of the old path (red dot) is congested. In this scenario, Frame 3 and Frame 4 can be delivered before Frame 1 and Frame 2.

The in-order delivery feature attempts to minimize the number of frames dropped during PortChannel link changes when the in-order delivery is enabled by sending a request to the remote switch on the PortChannel to flush all frames for this PortChannel.



Note

Both switches on the PortChannel must be running Cisco SAN-OS Release 3.0(1) for this IOD enhancement. For earlier releases, IOD waits for the switch latency period before sending new frames.

When the in-order delivery guarantee feature is enabled and a PortChannel link change occurs, the frames crossing the PortChannel are treated as follows:

- Frames using the old path are delivered before new frames are accepted.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- The new frames are delivered through the new path after the switch latency drop period has elapsed and all old frames are flushed.

Frames that cannot be delivered in order through the old path within the switch latency drop period are dropped. See the [“Configuring the Drop Latency Time” section on page 1-14](#).

About Enabling In-Order Delivery

You can enable the in-order delivery feature for a specific VSAN or for the entire switch. By default, in-order delivery is disabled on switches in the Cisco MDS 9000 Family.



Tip

We recommend that you only enable this feature when devices that cannot handle any out-of-order frames are present in the switch. Load-balancing algorithms within the Cisco MDS 9000 Family ensure that frames are delivered in order during normal fabric operation. The load-balancing algorithms based on source FC ID, destination FC ID, and exchange ID are enforced in hardware without any performance degradation. However, if the fabric encounters a failure and this feature is enabled, the recovery will be delayed because of an intentional pausing of fabric forwarding to purge the fabric of resident frames that could potentially be forwarded out-of-order.

About Flow Statistics

If you enable flow counters, you can enable a maximum of 1 K entries for aggregate flow and flow statistics for Generation 1 modules, and 2 K entries for Generation 2 modules. Be sure to assign an unused flow index to a module for each new flow. Flow indexes can be repeated across modules. The number space for flow index is shared between the aggregate flow statistics and the flow statistics.

Generation 1 modules allow a maximum of 1024 flow statements per module. Generation 2 modules allow a maximum of 2048-128 flow statements per module.



Note

For each session, fcflow counter will increment only on locally connected devices and should be configured on the switch where the initiator is connected.

Licensing Requirements for FSPF

The following table shows the licensing requirements for this feature:

License	License Description
ENTERPRISE_PKG	The enterprise license is required to enable Fibre Channel routing services and protocols. For a complete explanation of the licensing scheme, see the <i>Cisco MDS 9000 Family NX-OS Licensing Guide</i> .

Default Settings

[Table 1-2](#) lists the default settings for FSPF features.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Table 1-2 Default FSPF Settings

Parameters	Default
FSPF	Enabled on all E ports and TE ports.
SPF computation	Dynamic.
SPF hold time	0.
Backbone region	0.
Acknowledgment interval (RxmtInterval)	5 seconds.
Refresh time (LSRefreshTime)	30 minutes.
Maximum age (MaxAge)	60 minutes.
Hello interval	20 seconds.
Dead interval	80 seconds.
Distribution tree information	Derived from the principal switch (root node).
Routing table	FSPF stores up to 16 equal cost paths to a given destination.
Load balancing	Based on destination ID and source ID on different, equal cost paths.
In-order delivery	Disabled.
Drop latency	Disabled.
Static route cost	If the cost (metric) of the route is not specified, the default is 10.
Remote destination switch	If the remote destination switch is not specified, the default is direct.
Multicast routing	Uses the principal switch to compute the multicast tree.

Configuring FSPF

This section includes the following topics:

- [Configuring FSPF on a VSAN, page 1-10](#)
- [Resetting FSPF to the Default Configuration, page 1-10](#)
- [Enabling or Disabling FSPF, page 1-10](#)
- [Configuring FSPF Link Cost, page 1-11](#)
- [Configuring Hello Time Intervals, page 1-11](#)
- [Configuring Dead Time Intervals, page 1-11](#)
- [Configuring Retransmitting Intervals, page 1-12](#)
- [Disabling FSPF for Specific Interfaces, page 1-12](#)
- [Configuring Fibre Channel Routes, page 1-12](#)
- [Setting the Multicast Root Switch, page 1-13](#)
- [Enabling In-Order Delivery Globally, page 1-13](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- [Enabling In-Order Delivery for a VSAN, page 1-14](#)
- [Configuring the Drop Latency Time, page 1-14](#)

Configuring FSPF on a VSAN

Detailed Steps

To configure an FSPF feature for the entire VSAN, follow these steps:

-
- | | |
|---------------|--|
| Step 1 | Expand a Fabric, expand a VSAN and select FSPF for a VSAN that you want to configure for FSPF. You see the FSPF configuration in the Information pane. |
| Step 2 | The RegionID, Spf Comp Holdtime, LSR Min Arrival, and LSR Min Interval field values are applied across all interfaces on the VSAN. You can change them here or, if they do not exist create them here. |
| Step 3 | Click Apply Changes to save these changes, or click Undo Changes to discard any unsaved changes. |
-

Resetting FSPF to the Default Configuration

Detailed Steps

To return the FSPF VSAN global configuration to its factory default, follow these steps:

-
- | | |
|---------------|--|
| Step 1 | Expand a Fabric, expand a VSAN, and select FSPF for a VSAN that you want to configure for FSPF. You see the FSPF configuration in the Information pane. |
| Step 2 | Check the SetToDefault check box for a switch. |
| Step 3 | Click Apply Changes to save these changes, or click Undo Changes to discard any unsaved changes. |
-

Enabling or Disabling FSPF

Detailed Steps

To enable or disable FSPF, follow these steps:

-
- | | |
|---------------|--|
| Step 1 | Expand a Fabric, expand a VSAN, and select FSPF for a VSAN that you want to configure for FSPF. You see the FSPF configuration in the Information pane. |
| Step 2 | Set the Status Admin drop-down menu to up to enable FSPF or to down to disable FSPF. |
| Step 3 | Click Apply Changes to save these changes, or click Undo Changes to discard any unsaved changes. |
-

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuring FSPF Link Cost

Detailed Steps

To configure FSPF link cost, follow these steps:

-
- | | |
|---------------|--|
| Step 1 | Expand Switches , expand FC Interfaces , and then select Physical .
You see the interface configuration in the Information pane. |
| Step 2 | Click the FSPF tab.
You see the FSPF interface configuration in the Information pane. |
| Step 3 | Double-click in the Cost field of a switch and change the value. |
| Step 4 | Click Apply Changes to save these changes, or click Undo Changes to discard any unsaved changes. |
-

Configuring Hello Time Intervals

Detailed Steps

To configure the FSPF Hello time interval, follow these steps:

-
- | | |
|---------------|--|
| Step 1 | Expand Switches , expand FC Interfaces , and then select Physical .
You see the interface configuration in the Information pane. |
| Step 2 | Click the FSPF tab.
You see the FSPF interface configuration in the Information pane. |
| Step 3 | Change the Hello Interval field for a switch. |
| Step 4 | Click Apply Changes to save these changes, or click Undo Changes to discard any unsaved changes. |
-

Configuring Dead Time Intervals

Detailed Steps

To configure the FSPF dead time interval, follow these steps:

-
- | | |
|---------------|--|
| Step 1 | Expand Switches , expand FC Interfaces , and then select Physical .
You see the interface configuration in the Information pane. |
| Step 2 | Click the FSPF tab.
You see the FSPF interface configuration in the Information pane. |
| Step 3 | Double-click the Dead Interval field for a switch and provide a new value. |

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Step 4** Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.
-

Configuring Retransmitting Intervals

Detailed Steps

To configure the FSPF retransmit time interval, follow these steps:

-
- Step 1** Expand **Switches**, expand **FC Interfaces**, and then select **Physical**.
You see the interface configuration in the Information pane.
- Step 2** Click the **FSPF** tab.
You see the FSPF interface configuration in the Information pane.
- Step 3** Double-click the ReTx Interval field and enter a value.
- Step 4** Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.
-

Disabling FSPF for Specific Interfaces

You can disable the FSPF protocol for selected interfaces. By default, FSPF is enabled on all E ports and TE ports. This default can be disabled by setting the interface as passive.

You can disable the FSPF protocol for selected interfaces. By default, FSPF is enabled on all E ports and TE ports. This default can be disabled by setting the interface as passive.

Detailed Steps

To disable FSPF for a specific interface, follow these steps:

-
- Step 1** Expand **Switches**, expand **FC Interfaces**, and then select **Physical**.
You see the interface configuration in the Information pane.
- Step 2** Click the **FSPF** tab.
You see the FSPF interface configuration in the Information pane.
- Step 3** Set a switch Admin Status drop-down menu to **down**.
- Step 4** Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.
-

Configuring Fibre Channel Routes

Detailed Steps

If you disable FSPF, you can manually configure a Fibre Channel route.

Send documentation comments to dcnm-san-docfeedback@cisco.com

To configure a Fibre Channel route using Device Manager, follow these steps:

-
- Step 1** Click **FC > Advanced > Routes**.
You see the FC Static Route Configuration dialog box.
 - Step 2** Click **Create** to create a static route.
You see the Create Route dialog box.
 - Step 3** Select the VSAN ID that you are configuring this route.
 - Step 4** Fill in the destination address and destination mask for the device you are configuring a route.
 - Step 5** Select the interface that you want to use to reach this destination.
 - Step 6** Select the next hop domain ID and route metric.
 - Step 7** Select either the **local** or **remote** radio button.
 - Step 8** Click **Create** to save these changes or click **Close** to discard any unsaved changes.
-

Setting the Multicast Root Switch

Detailed Steps

To use the lowest domain switch for the multicast tree computation, follow these steps:

-
- Step 1** Expand a fabric, expand a VSAN, and then select **Advanced** for the VSAN that you want to configure FSPF on.
You see the advanced Fibre Channel configuration in the Information pane.
 - Step 2** Click the **Multicast Root** tab.
You see the multicast root configuration in the Information pane.
 - Step 3** Set the Config Mode drop-down menu to **lowestDomainSwitch**.
 - Step 4** Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.
-

Enabling In-Order Delivery Globally

To ensure that the in-order delivery parameters are uniform across all VSANs on an Cisco MDS 9000 Family switch, enable in-order delivery globally.

Only enable in-order delivery globally if this is a requirement across your entire fabric. Otherwise, enable IOD only for the VSANs that require this feature.



Note

Enable in-order delivery on the entire switch before performing a downgrade to Cisco MDS SAN-OS Release 1.3(3) or earlier.

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

Enabling In-Order Delivery for a VSAN

When you create a VSAN, that VSAN automatically inherits the global in-order-guarantee value. You can override this global value by enabling or disabling in-order-guarantee for the new VSAN.

Detailed Steps

To use the lowest domain switch for the multicast tree computation, follow these steps:

-
- Step 1** Expand a fabric and select **All VSANS**.
 - Step 2** Select the **Attributes** tab.
You see the general VSAN attributes in the Information pane.
 - Step 3** Check the **InOrder Delivery** check box to enable IOD for the switch.
 - Step 4** Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.
-

Configuring the Drop Latency Time

You can change the default latency time for a network, a specified VSAN in a network, or for the entire switch.

Detailed Steps

To configure the drop latency time for a switch, follow these steps:

-
- Step 1** Expand a fabric and select **All VSANS**.
You see the VSAN configuration in the Information pane.
 - Step 2** Click the **Attributes** tab.
You see the general VSAN attributes in the Information pane.
 - Step 3** Double-click the Network Latency field and change the value.
 - Step 4** Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.
-

Verifying FSPF Configuration

For detailed information about the fields in the output from these commands, refer to the *Cisco MDS 9000 Family Command Reference*.

This section contains the following topics:

- [Displaying the FSPF Database, page 1-15](#)
- [Displaying FSPF Statistics, page 1-15](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

Displaying the FSPF Database

The FSPF database for a specified VSAN includes the following information:

- Link State Record (LSR) type
- Domain ID of the LSR owner
- Domain ID of the advertising router
- LSR age
- LSR incarnation member
- Number of links

To display the FSPF database using Device Manager, follow these steps:

-
- Step 1** Choose **FC > Advanced > FSPF**.
You see the FSPF dialog box.
- Step 2** Click the **LSDB LSRs** tab.
You see the FSPF database information.
- Step 3** Click **Close** to close the dialog box.
-

Displaying FSPF Statistics

To view FSPF statistics using DCNM-SAN, follow these steps:

-
- Step 1** Expand a Fabric, expand a VSAN, and then select **FSPF** in the Logical Domains pane.
You see the FSPF configuration dialog box.
- Step 2** Click the **Statistics** tab.
You see the FSPF VSAN statistics in the Information pane.
- Step 3** Click the **Interface Statistics** tab.
You see the FSPF interface statistics in the Information pane.
-

Configuration Examples for FSPF

This section provides examples of topologies and applications that demonstrate the benefits of FSPF.



Note


The FSPF feature can be used on any topology.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Fault Tolerant Fabric

Figure 1-4 depicts a fault tolerant fabric using a partial mesh topology. If a link goes down anywhere in the fabric, any switch can still communicate with all others in the fabric. In the same way, if any switch goes down, the connectivity of the rest of the fabric is preserved.

Figure 1-4 Fault Tolerant Fabric




For example, if all links are of equal speed, the FSPF calculates two equal paths from A to C: A-D-C (green) and A-E-C (blue).

Redundant Links

To further improve on the topology in Figure 1-4, each connection between any pair of switches can be replicated; two or more links can be present between a pair of switches. Figure 1-5 shows this arrangement. Because switches in the Cisco MDS 9000 Family support PortChanneling, each pair of physical links can appear to the FSPF protocol as one single logical link.

By bundling pairs of physical links, FSPF efficiency is considerably improved by the reduced database size and the frequency of link updates. Once physical links are aggregated, failures are not attached to a single link but to the entire PortChannel. This configuration also improves the resiliency of the network. The failure of a link in a PortChannel does not trigger a route change, thereby reducing the risks of routing loops, traffic loss, or fabric downtime for route reconfiguration.

Figure 1-5 Fault Tolerant Fabric with Redundant Links



For example, if all links are of equal speed and no PortChannels exist, the FSPF calculates four equal paths from A to C: A1-E-C, A2-E-C, A3-D-C, and A4-D-C. If PortChannels exist, these paths are reduced to two.

Failover Scenarios for PortChannels and FSPF Links

The SmartBits traffic generator was used to evaluate the scenarios displayed in Figure 1-6. Two links between switch 1 and switch 2 exist as either equal-cost ISLs or PortChannels. There is one flow from traffic generator 1 to traffic generator 2. The traffic was tested at 100 percent utilization at 1 Gbps in two scenarios:

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Disabling the traffic link by physically removing the cable (see [Table 1-3](#)).
- Shutting down either switch 1 or switch 2 (see [Table 1-4](#)).

Figure 1-6 Failover Scenario Using Traffic Generators




Table 1-3 Physically Removing the Cable for the SmartBits Scenario

PortChannel Scenario		FSPF Scenario (Equal cost ISL)	
Switch 1	Switch 2	Switch 1	Switch 2
110 msec (~2K frame drops)		130+ msec (~4k frame drops)	
100 msec (hold time when a signal loss is reported as mandated by the standard)			

Table 1-4 Shutting Down the Switch for the SmartBits Scenario

PortChannel Scenario		FSPF Scenario (Equal cost ISL)	
Switch 1	Switch 2	Switch 1	Switch 2
~0 msec (~8 frame drops)	110 msec (~2K frame drops)	130+ msec (~4K frame drops)	
No hold time needed	Signal loss on switch 1	No hold time needed	Signal loss on switch 1

Field Descriptions for FSPF

This section displays the field descriptions for this feature.

FSPF General

Field	Description
AdminStatus	The desired state of FSPF on this VSAN.
OperStatus	State of FSPF on this VSAN.
SetToDefault	Enabling this changes each value in this row to its default value. If all the configuration parameters have their default values and if the VSAN is suspended, then the row is deleted automatically.
RegionId	The autonomous region of the local switch on this VSAN.
DomainId	The domain ID of the local switch on this VSAN.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
SpfHoldTime	The minimum time between two consecutive SPF computations on this VSAN. The smaller value means that routing will react to the changes faster but the CPU usage is greater.
SpfDelay	The time between when FSPF receives topology updates and when it starts the Shortest Path First (SPF) computation on this VSAN. The smaller value means that routing will react to the changes faster but the CPU usage is greater.
MinLsArrival	The minimum time after accepting a Link State Record (LSR) on this VSAN before accepting another update of the same LSR on the same VSAN. An LSR update that is not accepted because of this time interval is discarded.
MinLsInterval	The minimum time after this switch sends an LSR on this VSAN before it will send another update of the same LSR on the same VSAN.
LsRefreshTime	The interval between transmission of refresh LSRs on this VSAN.
LSRMaxAge	The maximum age an LSR will be retained in the FSPF database on this VSAN. It is removed from the database after MaxAge is reached.
CreateTime	When this entry was last created.
Checksum	The total checksum of all the LSRs on this VSAN.

FSPF Interfaces

Field	Description
SetToDefault	Enabling this changes each value in this row to its default value. If all the configuration parameters have their default values and if the interface is down, then the row is deleted automatically.
Cost	<p>The administrative cost of sending a frame on this interface on this VSAN. The value 0 means that the cost has not been configured. Once the value has been configured, the value cannot again be 0; so, obviously the value cannot be set to 0. If the value is 0 and the corresponding interface is up, the agent sets a value calculated using the ifSpeed of the interface. Otherwise, the value is used as the cost.</p> <p>The following formula is used to calculate the link cost: $\text{Link Cost} = \begin{cases} \text{fspfIfCost} & \text{if } \text{fspfIfCost} > 0 \\ (1.0625 \times 10^{12} / \text{Baud Rate}) & \text{if } \text{fspfIfCost} == 0 \end{cases}$ where Baud Rate is the ifSpeed of the interface.</p>
AdminStatus	The desired state of FSPF on this interface on this VSAN.
HelloInterval	Interval between the periodic hello messages sent on this interface on this VSAN to verify the link health. Note that this value must be same on both the interfaces on each end of the link on this VSAN.
DeadInterval	<p>Maximum time for which no hello messages can be received on this interface on this VSAN. After this time, the interface is assumed to be broken and removed from the database.</p> <p>This value must be greater than the hello interval specified on this interface on this VSAN.</p>

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
RetransmitInterval	Time after which an unacknowledged link update is retransmitted on this interface on this VSAN.
Neighbour State	The state of FSPF's neighbor state machine, which is the operational state of the interaction with the neighbor's interface which is connected to this interface.
Neighbour DomainId	The domain ID of the neighbor on this VSAN.
Neighbour PortIndex	The index, as known by the neighbor, of the neighbor's interface which is connected to this interface on this VSAN.
CreateTime	When this entry was last created.

FSPF Interface Stats

Field	Description
CreateTime	When this entry was last created.
ErrorRxPkts	Number of invalid FSPF control frames received on this interface on this VSAN since the creation of the entry.
InactivityExpirations	Number of times the inactivity timer has expired on this interface on this VSAN since the creation of the entry.
LsuRxPkts	Number of Link State Update (LSU) frames received on this interface on this VSAN since the creation of the entry.
LsuTxPkts	Number of Link State Update (LSU) frames transmitted on this interface on this VSAN since the creation of the entry.
RetransmittedLsuTxPkts	Number of LSU frames retransmitted on this interface on this VSAN since the creation of the entry.
LsaRxPkts	Number of Link State Acknowledgement (LSA) frames received on this interface on this VSAN since the creation of the entry.
LsaTxPkts	Number of Link State Acknowledgement (LSA) frames transmitted on this interface on this VSAN since the creation of the entry.
HelloTxPkts	Number of HELLO frames transmitted on this interface on this VSAN since the creation of the entry.
HelloRxPkts	Number of HELLO frames received on this interface on this VSAN since the creation of the entry.

FSPF LSDB Links

Field	Description
NbrDomainId	The domain ID of the neighbor on the other end of this link on this VSAN.
PortIndex	The source E_port of this link, as indicated by the index value in the LSR received from the switch identified by the domain ID.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
NbrPortIndex	The destination E_port of this link, as indicated by the index value in the LSR received from the switch identified by NbrDomainId.
Cost	The cost of sending a frame on this link on this VSAN. Link cost is calculated using a formula link cost = S * (1.0625e12/Baud Rate) where S (value of Cost on the interface on the switch corresponding to the domain Id) is the administratively set cost factor for this interface.

FSPF LSDB LSRs

Field	Description
AdvDomainId	Domain ID of the switch that is advertising the LSR on the behalf of the switch owning it.
Age	Time since this LSR was inserted into the database.
IncarnationNumber	The link state incarnation number of this LSR. This is used to identify most recent instance of an LSR while updating the topology database when an LSR is received. The updating of an LSR includes incrementing its incarnation number prior to transmission of the updated LSR. So most recent LSR is the one with larger incarnation number.
Checksum	The checksum of the LSR.
Links	Number of entries associated with this LSR.
External	Indicates of this is an external LSR advertised by local switch.

FSPF Statistics

Field	Description
SpfComputations	The number of times the SPF computation has been done on this VSAN since the creation of the entry.
ErrorRxPkts	Number of invalid FSPF control frames received on all the interface on this VSAN since the creation of the entry.
ChecksumErrors	The number of FSPF checksum errors occurred on this on this VSAN since the creation of the entry.
LsuRxPkts	Total number of Link State Update (LSU) frames received on all the interfaces on this VSAN since the creation of the entry.
LsuTxPkts	Total number of Link State Update (LSU) frames transmitted on all the interfaces on this VSAN since the creation of the entry.
RetransmittedLsuTxPkts	Total number of LSU frames retransmitted on all the interfaces on this VSAN since the creation of the entry.
LsaRxPkts	Total number of Link State Acknowledgement (LSA) frames received on all the interfaces on this VSAN since the creation of the entry.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
LsaTxPkts	Total number of Link State Acknowledgement (LSA) frames transmitted on all the interfaces on this VSAN since the creation of the entry.
HelloTxPkts	Total number of HELLO frames transmitted on all interfaces on this VSAN since the creation of the entry.
HelloRxPkts	Total number of HELLO frames received on all the interfaces on this VSAN since the creation of the entry.
MaxAgeCount	The number of times any LSR reached fspfMaxAge in this VSAN since the creation of the entry.

Additional References

For additional information related to implementing VSANs, see the following section:

- [Related Document, page 1-21](#)
- [Standards, page 1-21](#)
- [RFCs, page 1-22](#)
- [MIBs, page 1-22](#)

Related Document

Related Topic	Document Title
Cisco MDS 9000 Family Command Reference	<i>Cisco MDS 9000 Family Command Reference</i>

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

Send documentation comments to dcnm-san-docfeedback@cisco.com

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified.	—

MIBs

MIBs	MIBs Link
<ul style="list-style-type: none">CISCO-FC-ROUTE-MIB	To locate and download MIBs, go to the following URL: http://www.cisco.com/en/US/products/ps5989/prod_technical_reference_list.html



CHAPTER 1

Configuring Dense Wavelength Division Multiplexing

This chapter includes the following topics:

- [Information About DWDM, page 1-1](#)
- [Configuring X2 DWDM Transceiver Frequency, page 1-1](#)
- [Monitoring DWDM Links, page 1-2](#)
- [Field Descriptions for DPVM, page 1-2](#)
- [Additional References, page 1-4](#)

Information About DWDM

Dense Wavelength-Division Multiplexing (DWDM) multiplexes multiple optical carrier signals on a single optical fiber. DWDM uses different wavelengths to carry various signals.

To establish a DWDM link, both ends of an Inter Switch Link (ISL) need to be connected with DWDM SFPs (small form-factor pluggable) at each end of the link. To identify a DWDM link, DCNM-SAN discovers the connector type on the Fiber Channel (FC) ports. If the ISL link is associated with the FC ports at each end, then the FC port uses DWDM SFP to connect the links.

Cisco DCNM for SAN discovers FC ports with DWDM SFPs and the ISLs associated with the FC ports. The DCNM-SAN Client displays ISL with DWDM attribute on the topology map.



Note

The Fabric Shortest Path First (FSPF) database only displays an ISL link, which is connected with DWDM SFPs at both ends.

Configuring X2 DWDM Transceiver Frequency

Restrictions

This feature is supported only in MDS 9134 modules. With MDS 9134 modules, the 10-Gigabit Ethernet ports must be in a down state when you configure the X2 transceiver frequency.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Detailed Steps

To configure the X2 DWDM transceiver frequency using Device Manager, follow these steps:

-
- Step 1** From the Device Manager menu bar, select **Physical > Modules**.
The module configuration window is displayed.
 - Step 2** Choose an **XcvrFrequencyConfig** option button.
 - Step 3** Click **Apply**.
-

To configure the X2 DWDM transceiver frequency, follow these steps:

-
- Step 1** From the Physical Attributes pane, select **Hardware**.
The module configuration window is displayed.
 - Step 2** Click the **Card Module Config** tab.
 - Step 3** In the X2 XcvrFrequencyConfig column, choose an option.
 - Step 4** Click **Apply**.
-

Monitoring DWDM Links

The DCNM-SAN Client displays DWDM links with a “dash-dash” pattern. The tooltip for the link displays “DWDM” to indicate its link type.

To view the DWDM link, follow these steps:

-
- Step 1** Select the switch in the Logical Domain region.
 - Step 2** Select ISL in the Physical Attributes region.
The Information pane displays the ISL’s information.
 - Step 3** Click the **Physical** tab.
You see the ISL in the Information pane.
The ISL’s Physical table displays the connector type as sfpDwdm.
Move the mouse over the link to see the tooltip as DWDM indicating the link type.
 - Step 4** Perform a Dump Discovery of ISL to list all ISLs. DWDM links are listed with [DWDM].
-

Field Descriptions for DPVM

This section displays the following field descriptions for this feature.

Send documentation comments to dcnm-san-docfeedback@cisco.com

DPVM Actions

Field	Description
Action	Helps in activating the set of bindings.
Result	Indicates the outcome of the activation.
Status	Indicates the state of activation. If true, then activation has been attempted as the most recent operation. If false, then an activation has not been attempted as the most recent operation.
CopyActive to Config	When set to copy(1), results in the active (enforced) binding database to be copied on to the configuration binding database. The learned entries are also copied.
Auto Learn Enable	Helps to learn the configuration of devices logged into the local device on all its ports and the VSANs to which they are associated.
Auto Learn Clear	Assists in clearing the auto-learned entries.
Clear WWN	Represents the port WWN (pWWN) to be used for clearing its corresponding auto-learned entry.

DPVM Config Database

Field	Description
Type	Specifies the type of the corresponding instance of cdpvmLoginDev object.
WWN or Name	Represents the logging in device.
VSAN Id	Represents the VSAN to be associated to the port on the local device on which the device represented by cdpvmLoginDev logs in.
Switch Interface	Represents the device alias.

DPVM Active Database

Field	Description
Type	Specifies the type of the corresponding instance of cdpvmEnfLoginDev.
WWN or Name	Represents the logging in device address.
VSAN Id	Represents the VSAN of the port on the local device through which the device represented by cdpvmEnfLoginDev logs in.
Interface	Represents the device alias.
IsLearnt	Indicates whether this is a learned entry or not. If true, then it is a learned entry. If false, then it is not.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Additional References

For additional information related to implementing VSANs, see the following section:

- [Related Document, page 1-4](#)
- [Standards, page 1-4](#)
- [RFCs, page 1-4](#)
- [MIBs, page 1-4](#)

Related Document

Related Topic	Document Title
Cisco MDS 9000 Family Command Reference	<i>Cisco MDS 9000 Family Command Reference</i>

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified.	—

MIBs

MIBs	MIBs Link
No new or modified MIBs are supported by this feature, and support for existing MIBs has not been modified by this feature.	To locate and download MIBs, go to the following URL: http://www.cisco.com/dc-os/mibs



CHAPTER 1

Managing FLOGI, Name Server, FDMI, and RSCN Databases

This chapter describes the fabric login (FLOGI) database, the name server features, the Fabric-Device Management Interface, and Registered State Change Notification (RSCN) information provided in the Cisco MDS 9000 Family Switches. It includes the following sections:

- [Information About FLOGI, page 1-1](#)
- [Default Settings, page 1-5](#)
- [Registering Name Server Proxies, page 1-5](#)
- [Verifying the Database Configuration, page 1-8](#)
- [Field Descriptions for Databases, page 1-9](#)
- [Additional References, page 1-13](#)

Information About FLOGI

In a Fibre Channel fabric, each host or disk requires an FC ID. If the required device is displayed in the FLOGI table, the fabric login is successful. Examine the FLOGI database on a switch that is directly connected to the host HBA and connected ports. See the [“Default Company ID List” section on page 1-4](#) and the [“Switch Interoperability” section on page 1-5](#).

This section includes the following topics:

- [Name Server Proxy, page 1-2](#)
- [About Registering Name Server Proxies, page 1-2](#)
- [About Rejecting Duplicate pWWN, page 1-2](#)
- [About Name Server Database Entries, page 1-2](#)
- [FDMI, page 1-2](#)
- [RSCN, page 1-3](#)
- [About the multi-pid Option, page 1-3](#)
- [RSCN Timer Configuration Distribution Using CFS, page 1-3](#)
- [RSCN Timer Configuration Distribution, page 1-4](#)
- [Locking the Fabric, page 1-5](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

Name Server Proxy

The name server functionality maintains a database containing the attributes for all hosts and storage devices in each VSAN. Name servers allow a database entry to be modified by a device that originally registered the information.

The proxy feature is useful when you want to modify (update or delete) the contents of a database entry that was previously registered by a different device.

About Registering Name Server Proxies

All name server registration requests come from the same port whose parameter is registered or changed. If it does not, then the request is rejected.

This authorization enables WWNs to register specific parameters for another node.

About Rejecting Duplicate pWWN

You can prevent a malicious or accidental login when using another device's pWWN. These pWWNs are allowed to log in to the fabric and replace the first device in the name server database.

About Name Server Database Entries

The name server stores name entries for all hosts in the FCNS database. The name server permits an Nx port to register attributes during a PLOGI (to the name server) to obtain attributes of other hosts. These attributes are deregistered when the Nx port logs out either explicitly or implicitly.

In a multiswitch fabric configuration, the name server instances running on each switch shares information in a distributed database. One instance of the name server process runs on each switch.

FDMI

Cisco MDS 9000 Family switches provide support for the Fabric-Device Management Interface (FDMI) functionality, as described in the FC-GS-4 standard. FDMI enables management of devices such as Fibre Channel host bus adapters (HBAs) through in-band communications. This addition complements the existing Fibre Channel name server and management server functions.

Using the FDMI functionality, the NX-OS software can extract the following management information about attached HBAs and host operating systems without installing proprietary host agents:

- Manufacturer, model, and serial number
- Node name and node symbolic name
- Hardware, driver, and firmware versions
- Host operating system (OS) name and version number

All FDMI entries are stored in persistent storage and are retrieved when the FDMI process is started.

Send documentation comments to dcnm-san-docfeedback@cisco.com

RSCN

The Registered State Change Notification (RSCN) is a Fibre Channel service that informs hosts about changes in the fabric. Hosts can receive this information by registering with the fabric controller (through SCR). These notifications provide a timely indication of one or more of the following events:

- Disks joining or leaving the fabric.
- A name server registration change.
- A new zone enforcement.
- IP address change.
- Any other similar event that affects the operation of the host.

Apart from sending these events to registered hosts, a switch RSCN (SW-RSCN) is sent to all reachable switches in the fabric.



Note

The switch sends an RSCN to notify registered nodes that a change has occurred. It is up to the nodes to query the name server again to obtain the new information. The details of the changed information are not delivered by the switch in the RSCN sent to the nodes.

About the multi-pid Option

If the RSCN **multi-pid** option is enabled, then RSCNs generated to the registered Nx ports may contain more than one affected port IDs. In this case, zoning rules are applied before putting the multiple affected port IDs together in a single RSCN. By enabling this option, you can reduce the number of RSCNs. For example, suppose you have two disks (D1, D2) and a host (H) connected to switch 1. Host H is registered to receive RSCNs. D1, D2, and H belong to the same zone. If disks D1 and D2 are online at the same time, then one of the following applies:

- The **multi-pid** option is disabled on switch 1— Two RSCNs are generated to host H, one for the disk D1 and another for disk D2.
- The **multi-pid** option is enabled on switch 1— A single RSCN is generated to host H, and the RSCN payload lists the affected port IDs (in this case, both D1 and D2).

Some Nx ports might not support multi-pid RSCN payloads. If this situation occurs, disable the RSCN **multi-pid** option.

RSCN Timer Configuration Distribution Using CFS

Because the timeout value for each switch is configured manually, a misconfiguration occurs when different switches time out at different times. This means different N ports in a network can receive RSCNs at different times. Cisco Fabric Services (CFS) alleviates this situation by automatically distributing configuration information to all switches in a fabric. This also reduces the number of SW-RSCNs.

RSCN supports two modes, distributed and nondistributed. In distributed mode, RSCN uses CFS to distribute configuration to all switches in the fabric. In nondistributed mode, only the configuration commands on the local switch are affected.

Send documentation comments to dcnm-san-docfeedback@cisco.com

**Note**

All configuration commands are not distributed. Only the **rscn event-tov tov vsan vsan** command is distributed.

The RSCN timer is registered with CFS during initialization and switchover. For high availability, if the RSCN timer distribution crashes and restarts or a switchover occurs, it resumes normal functionality from the state prior to the crash or switchover.

**Note**

Before performing a downgrade, make sure that you revert the RSCN timer value in your network to the default value. Failure to do so will disable the links across your VSANs and other devices.

Compatibility across various Cisco MDS NX-OS releases during an upgrade or downgrade is supported by **conf-check** provided by CFS. If you attempt to downgrade from Cisco MDS SAN-OS Release 3.0, you are prompted with a **conf-check** warning. You are required to disable RSCN timer distribution support before you downgrade.

By default, the RSCN timer distribution capability is disabled and is therefore compatible when upgrading from any Cisco MDS SAN-OS release earlier than Release 3.0.

RSCN Timer Configuration Distribution

Because the timeout value for each switch is configured manually, a misconfiguration occurs when different switches time out at different times. This means different Nports in a network can receive RSCNs at different times. Cisco Fabric Services (CFS) infrastructure alleviates this situation by automatically distributing the RSCN timer configuration information to all switches in a fabric. This action also reduces the number of SW-RSCNs. Refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*.

RSCN supports two modes, distributed and nondistributed. In distributed mode, RSCN uses CFS to distribute configuration to all switches in the fabric. In nondistributed mode, only the configuration commands on the local switch are affected.

**Note**

All configuration commands are not distributed. Only the **rscn event-tov tov vsan vsan** command is distributed.

**Note**

Only the RSCN timer configuration is distributed.

The RSCN timer is registered with CFS during initialization and switchover. For high availability, if the RSCN timer distribution crashes and restarts or a switchover occurs, it resumes normal functionality from the state prior to the crash or switchover.

**Note**

You can determine the compatibility when downgrading to an earlier Cisco MDS NX-OS release using **show incompatibility system** command. You must disable RSCN timer distribution support before downgrading to an earlier release.

Send documentation comments to dcnm-san-docfeedback@cisco.com

**Note**

By default, the RSCN timer distribution capability is disabled and is compatible when upgrading from any Cisco MDS SAN-OS release earlier than 3.0.

**Note**

For CFS distribution to operate correctly for the RSCN timer configuration, all switches in the fabric must be running Cisco SAN-OS Release 3.0(1) or later, or Cisco NX-OS 4.1(1b).

Locking the Fabric

The first action that modifies the database creates the pending database and locks the feature in the VSAN. Once you lock the fabric, the following situations apply:

- No other user can make any configuration changes to this feature.
- A copy of the configuration database becomes the pending database along with the first active change.

Default Settings

Table 1-1 lists the default settings for RSCN.

Table 1-1 **Default RSCN Settings**

Parameters	Default
RSCN timer value	2000 milliseconds for Fibre Channel VSANs 1000 milliseconds for FICON VSANs
RSCN timer configuration distribution	Disabled

Registering Name Server Proxies

This section includes the following topics:

- [Registering Name Server Proxies, page 1-6](#)
- [Configuring the multi-pid Option, page 1-6](#)
- [Suppressing Domain Format SW-RSCNs, page 1-6](#)
- [Configuring the RSCN Timer with CFS, page 1-7](#)
- [Configuring the RSCN Timer, page 1-7](#)
- [Committing the RSCN Timer Configuration Changes, page 1-7](#)
- [Discarding the RSCN Timer Configuration Changes, page 1-7](#)
- [Clearing a Locked Session, page 1-8](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

Registering Name Server Proxies

Detailed Steps

To register the name server proxy, follow these steps:

-
- Step 1** Expand a fabric, expand a VSAN, and then select **Advanced**.
You see the VSAN advanced configuration in the Information pane.
 - Step 2** Click the **NS Proxies** tab.
You see the existing name server proxy for the selected VSAN.
 - Step 3** Double-click the PortName field to register a new name server proxy.
 - Step 4** Click **Apply Changes** to save these changes, or click **Undo Changes** to cancel any unsaved changes.
-

Configuring the multi-pid Option

Detailed Steps

To configure the **multi-pid** option, follow these steps:

-
- Step 1** Expand a fabric, expand a VSAN, and then select **Advanced**.
You see the VSAN advanced configuration in the Information pane.
 - Step 2** Click the **RSCN Multi-PID** tab.
 - Step 3** Check the **Enable** check box.
 - Step 4** Click **Apply Changes** to save these changes, or click **Undo Changes** to cancel any unsaved changes.
-

Suppressing Domain Format SW-RSCNs

A domain format SW-RSCN is sent whenever the local switch name or the local switch management IP address changes. This SW-RSCN is sent to all other domains and switches over the ISLs. The remote switches can issue GMAL and GIELN commands to the switch that initiated the domain format SW-RSCN to determine what changed. Domain format SW-RSCNs can cause problems with some non-Cisco MDS switches (refer to the [Cisco MDS 9000 Family Switch-to-Switch Interoperability Configuration Guide](#)).



Note

You cannot suppress transmission of port address or area address format RSCNs.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuring the RSCN Timer with CFS

Detailed Steps

To configure the RSCN timer with CFS, follow these steps:

-
- | | |
|---------------|---|
| Step 1 | Expand a fabric, expand a VSAN, and then select Advanced in the Logical Domains pane. |
| Step 2 | Click the RSCN Event tab.

You see the VSAN advanced configuration in the Information pane. |
| Step 3 | Double-click the TimeOut value to change the value (in milliseconds) for the selected VSAN. |
| Step 4 | Click Apply Changes to save these changes, or click Undo Changes to cancel any unsaved changes. |
-

Configuring the RSCN Timer

RSCN maintains a per-VSAN event list queue, where the RSCN events are queued as they are generated. When the first RSCN event is queued, a per VSAN timer starts. Upon time-out, all the events are dequeued and coalesced RSCNs are sent to registered users. The default timer values minimize the number of coalesced RSCNs sent to registered users. Some deployments require smaller event timer values to track changes in the fabric.



Note

The RSCN timer value must be the same on all switches in the VSAN. See the [“RSCN Timer Configuration Distribution Using CFS” section on page 1-3](#).



Note

Before performing a downgrade, make sure that you revert the RSCN timer value in your network to the default value. Failure to do so will disable the links across your VSANs and other devices.

Committing the RSCN Timer Configuration Changes

If you commit the changes made to the active database, the configuration is committed to all the switches in the fabric. On a successful commit, the configuration change is applied throughout the fabric and the lock is released.

Discarding the RSCN Timer Configuration Changes

If you discard (abort) the changes made to the pending database, the configuration database remains unaffected and the lock is released.

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

Clearing a Locked Session

If you have changed the RSCN timer configuration and have forgotten to release the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your changes to the pending database are discarded and the fabric lock is released.



Tip

The pending database is only available in the volatile directory and are subject to being discarded if the switch is restarted.

Verifying the Database Configuration

For detailed information about the fields in the output from these commands, refer to the *Cisco MDS 9000 Family Command Reference*.

This section includes the following topics:

- [Displaying FLOGI Details, page 1-8](#)
- [Viewing Name Server Database Entries, page 1-8](#)
- [Displaying FDMI, page 1-9](#)
- [Displaying RSCN Information, page 1-9](#)

Displaying FLOGI Details

To verify that a storage device is in the fabric login (FLOGI) table, follow these steps:

-
- Step 1** Expand **Switches**, expand **Interfaces**, and then select **FC Physical**.
You see the interface configuration in the Information pane.
- Step 2** Click the **FLOGI** tab.
You see all end devices that are logged into the fabric.
-

Viewing Name Server Database Entries

To view the name server database using Device Manager, follow these steps:

-
- Step 1** Select **FC > Name Server**.
You see the Name Server dialog box.
The General tab is the default tab; you see the name server database.
- Step 2** Click the **Statistics** tab.
You see the name server statistics.
-

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 3 Click **Close** to close the dialog box.

Displaying FDMI

To display the FDMI database information using Device Manager, choose **FC > Advanced > FDMI**. You see the FDMI dialog box.

Displaying RSCN Information

To display RSCN information, follow these steps:

-
- Step 1** Expand a fabric, expand a VSAN, and then select **Advanced**.
You see the VSAN advanced configuration in the Information pane.
- Step 2** Click the **RSCN Reg** tab or the **RSCN Statistics** tab.
-

Field Descriptions for Databases

This section contains the field descriptions for this feature.

FC Interfaces FLOGI

Field	Description
FcId	The address identifier that has been assigned to the logged-in Nx_Port.
PortName	The world wide name of the logged-in Nx_Port.
NodeName	The world wide name of the Remote Node the logged-in Nx_Port belongs to.
Original PWWN	The original port WWN for this interface.
Version	The version of FC-PH that the Fx_Port has agreed to support from the Fabric Login.
BBCredit Rx	The maximum number of receive buffers available for holding Class 2, Class 3 received from the logged-in Nx_Port. It is for buffer-to-buffer flow control in the incoming direction from the logged-in Nx_Port to FC-port.
BBCredit Tx	The total number of buffers available for holding Class 2, Class 3 frames to be transmitted to the logged-in Nx_Port. It is for buffer-to-buffer flow control in the direction from FC-Port to Nx_Port. The buffer-to-buffer flow control mechanism is indicated in the respective BbCreditModel.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
CoS	The classes of services that the logged-in Nx_Port has requested the FC-Port to support and the FC-Port has granted the request.
Class2 RxDataSize	The Class 2 Receive Data Field Size of the logged-in Nx_Port. Specifies the largest Data Field Size for an FT_1 frame that can be received by the Nx_Port.
Class2 SeqDeliv	Whether the FC-Port has agreed to support Class 2 sequential delivery during the Fabric Login. This is meaningful only if Class 2 service has been agreed. This is applicable only to Fx_Ports.
Class3 RxDataSize	The Class3 Receive Data Field Size of the logged-in Nx_Port. Specifies the largest Data Field Size for an FT_1 frame that can be received by the Nx_Port.
Class3 SeqDeliv	Whether the FxPort has agreed to support Class 3 sequential delivery during the Fabric Login. This is meaningful only if Class 3 service has been agreed. This is applicable only to Fx_Ports.

FDMI HBAs

Field	Description
Sn	The serial number of this HBA.
Model	The model of this HBA.
ModelDescr	The model description.
OSInfo	The type and version of the operating system controlling this HBA.
MaxCTPayload	The maximum size of the Common Transport (CT) payload including all CT headers but no FC frame header(s), that may be send or received by application software resident in the host containing this HBA.

FDMI Ports

Field	Description
SupportedFC4Type	The supported FC-4 types attribute registered for this port on this VSAN.
SupportedSpeed	The supported speed registered for this port on this VSAN.
CurrentSpeed	The current speed registered for this port on this VSAN.
MaxFrameSize	The maximum frame size attribute registered for this port on this VSAN.
OsDevName	The OS Device Name attribute registered for this port on this VSAN.
HostName	The name of the host associated with this port.

Send documentation comments to dcnm-san-docfeedback@cisco.com

FDMI Versions

Field	Description
Hardware	The hardware version of this HBA.
DriverVer	The version level of the driver software controlling this HBA.
OptROMVer	The version of the Option ROM or the BIOS of this HBA.
Firmware	The version of the firmware executed by this HBA.

RSCN Nx Registrations

Field	Description
RegType	Indicates the type of registration desired by the subscriber. <ul style="list-style-type: none">'fromFabricCtrlr' indicates RSCNs generated by the Fabric Controller.'fromNxPort' indicates RSCNs generated by Nx_Ports.'fromBoth' indicates RSCNs generated by Fabric Controller and Nx_Ports.

RSCN Multi-PID Support

Field	Description
Enable	Specifies whether the multi-pid option is enabled on this VSAN.

RSCN Event

Field	Description
TimeOut (msec)	The time (in seconds) before the RSCN event times out.

RSCN Statistics

Field	Description
SCR Rx	The number of SCRs received from Nx_Ports on this VSAN.
SCR RJT	The number of SCR rejected on this VSAN.
RSCN Rx	The number of RSCNs from Nx_Ports received on this VSAN.
RSCN Tx	The total number of RSCNs transmitted on this VSAN.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
RSCN RJT	The number of RSCN requests rejected on this VSAN.
SW-RSCN Rx	The number of Inter-Switch Registered State Change Notifications (SW_RSCN) received on this VSAN from other switches.
SW-RSCN Tx	The number of Inter-Switch Registered State Change Notifications (SW_RSCN) transmitted on this VSAN to other switches.
SW-RSCN RJT	The number of SW_RSCN requests rejected on this VSAN.

Name Server General

Field	Description
Type	The port type of this port.
PortName	The fibre channel Port_Name (WWN) of this Nx_port.
NodeName	The fibre channel Node_Name (WWN) of this Nx_port.
FC4Type/Features	The FC-4 Features associated with this port and the FC-4 Type. Refer to FC-GS3 specification for the format.
SymbolicPortName	The user-defined name of this port.
SymbolicNodeName	The user-defined name of the node of this port.
FabricPortName	The fabric port name (WWN) of the Fx_port to which this Nx_port is attached.

Name Server Advanced

Field	Description
ClassOfSvc	The class of service indicator.
PortIpAddress	Contains the IP address of the associated port.
NodeIpAddress	The IP address of the node of this Nx_port, as indicated by the Nx_Port in a GS3 message that it transmitted.
SymbolicPortName	The user-defined name of this port.
SymbolicNodeName	The user-defined name of the node of this port.
HardAddress	Extended Link Service (FC-PH-2). Hard Address is the 24-bit NL_Port identifier which consists of - the 8-bit Domain Id in the most significant byte - the 8-bit Area Id in the next most significant byte - the 8-bit AL-PA(Arbitrated Loop Physical Address) which an NL_port attempts acquire during FC-AL initialization in the least significant byte. If the port is not an NL_Port, or if it is an NL_Port but does not have a hard address, then all bits are reported as 0s.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
ProcAssoc	The Fibre Channel initial process associator (IPA).
PermanentPortName	The permanent port name of this Nx port. If multiple port names are associated with this Nx port via FDISC (Discover F Port Service Parameters), the permanent port name is the original port name associated with this Nx port at login.

Name Server Proxy

Field	Description
PortName	Name of the proxy port which can register or deregister for other ports on this VSAN. Users can enable third-party registrations by setting this value.

Name Server Statistics

Field	Description
Queries Rx	The total number of Get Requests received by the local switch on this VSAN.
Queries Tx	The total number of Get Requests sent by the local switch on this VSAN.
Requests Rx Reg	The total number of Registration Requests received by the local switch on this VSAN.
Requests Rx DeReg	The total number of De-registration Requests received by the local switch on this VSAN.
RSCN Rx	The total number of RSCN commands received by the local switch on this VSAN.
RSCN Tx	The total number of RSCN commands sent by the local switch on this VSAN.
Rejects Tx	The total number of requests rejected by the local switch on this VSAN.

Additional References

For additional information related to implementing VSANs, see the following section:

- [Related Document, page 1-14](#)
- [Standards, page 1-14](#)
- [RFCs, page 1-14](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- [MIBs, page 1-14](#)

Related Document

Related Topic	Document Title
Cisco MDS 9000 Family Command Reference	<i>Cisco MDS 9000 Family Command Reference</i>

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified.	—

MIBs

MIBs	MIBs Link
<ul style="list-style-type: none"> • CISCO-FDMI-MIB • CISCO-FDMI-CAPABILITY 	To locate and download MIBs, go to the following URL: http://www.cisco.com/en/US/products/ps5989/prod_technical_reference_list.html



CHAPTER 1

Discovering SCSI Targets

This chapter describes the SCSI LUN discovery feature provided in switches in the Cisco MDS 9000 Family. It includes the following sections:

- [Information About SCSI LUN Discovery, page 1-1](#)
- [Licensing Requirements for SCSI, page 1-2](#)
- [Discovering SCSI Targets, page 1-2](#)
- [Verifying SCSI Targets Configuration, page 1-3](#)
- [Field Descriptions for SCSI Targets, page 1-3](#)
- [Additional References, page 1-8](#)

Information About SCSI LUN Discovery

Small Computer System Interface (SCSI) targets include disks, tapes, and other storage devices. These targets do not register logical unit numbers (LUNs) with the name server.

The name server requires LUN information for the following reasons:

- To display LUN storage device information so an NMS can access this information.
- To report device capacity, serial number, and device ID information.
- To register the initiator and target features with the name server.

The SCSI LUN discovery feature uses the local domain controller Fibre Channel address. It uses the local domain controller as the source FC ID, and performs SCSI INQUIRY, REPORT LUNS, and READ CAPACITY commands on SCSI devices.

The SCSI LUN discovery feature is initiated on demand, through CLI or SNMP. This information is also synchronized with neighboring switches, if those switches belong to the Cisco MDS 9000 Family.

This section includes the following topics:

- [About Starting SCSI LUN Discovery, page 1-1](#)
- [About Initiating Customized Discovery, page 1-2](#)

About Starting SCSI LUN Discovery

SCSI LUN discovery is done on demand.

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

Only Nx ports that are present in the name server database and that are registered as FC4 Type = SCSI_FCP are discovered.

About Initiating Customized Discovery

Customized discovery consists of a list of VSAN and domain pairs that are selectively configured to initiate a discovery. The domain ID is a number from 0 to 255 in decimal or a number from 0x0 to 0xFF in hex.

Licensing Requirements for SCSI

The following table shows the licensing requirements for this feature:

Feature	License Requirement
ENTERPRISE_PKG	The enterprise license is required to enable the SCSI flow statistics. For a complete explanation of the licensing scheme, see the <i>Cisco MDS 9000 Family NX-OS Licensing Guide</i> .
FM_SERVER_PKG	The Cisco DCNM for SAN Package is required to enable the traffic analyzer for SCSI flow statistics. For a complete explanation of the licensing scheme, see the <i>Cisco MDS 9000 Family NX-OS Licensing Guide</i> .

Discovering SCSI Targets

This section includes the following topics:

- [Starting SCSI LUN Discovery, page 1-2](#)
- [Initiating Customized Discovery, page 1-3](#)

Starting SCSI LUN Discovery

Detailed Steps

To begin SCSI LUN discovery using Device Manager, follow these steps:

-
- | | |
|---------------|---|
| Step 1 | Choose FC > Advanced > LUNs .
You see the LUN Configuration dialog box. |
| Step 2 | Set StartDiscovery to local, remote or both. |
| Step 3 | Choose the DiscoveryType and OS. |
| Step 4 | Click Apply to begin discovery. |
-

Send documentation comments to dcnm-san-docfeedback@cisco.com

Initiating Customized Discovery

Detailed Steps

To initiate a customized discovery using Device Manager, follow these steps:

-
- | | |
|---------------|---|
| Step 1 | From the VSAN drop-down menu, select the VSAN in which you want to initiate a customized discovery. |
| Step 2 | Click FC > Advanced > LUNs .
You see the LUN Configuration dialog box. |
| Step 3 | Set StartDiscovery to local, remote or both. |
| Step 4 | Fill in the DiscoveryType and OS fields. |
| Step 5 | Click Apply to begin discovery. |
-

Verifying SCSI Targets Configuration

For detailed information about the fields in the output from these commands, refer to the *Cisco MDS 9000 Family Command Reference*.

- [Displaying SCSI LUN Information, page 1-3](#)

Displaying SCSI LUN Information

To display the results of the discovery using Device Manager, follow these steps:

-
- | | |
|---------------|---|
| Step 1 | Choose FC > Advanced > LUNs .
You see the LUN Configuration dialog box. |
| Step 2 | Click the LUN tab or the Targets tab. |
-

Field Descriptions for SCSI Targets

The following are the field descriptions for SCSI targets.

iSCSI Connection

Field	Description
LocalAddr	The local Internet network address used by this connection.
RemoteAddr	The remote Internet network address used by this connection.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
CID	The iSCSI connection ID for this connection.
State	<p>The current state of this connection, from an iSCSI negotiation point of view.</p> <ul style="list-style-type: none"> login— The transport protocol connection has been established, but a valid iSCSI login response with the final bit set has not been sent or received. full—A valid iSCSI login response with the final bit set has been sent or received. logout— A valid iSCSI logout command has been sent or received, but the transport protocol connection has not yet been closed.
MaxRecvDSLen	The maximum data payload size supported for command or data PDUs in use within this connection. The size is reported in bytes even though the negotiation is in 512K blocks.
SendMarker	Indicates whether or not this connection is inserting markers in its outgoing data stream.
HeaderDigest	The iSCSI header digest scheme in use within this connection.
DataDigest	The iSCSI data digest scheme in use within this connection.

iSCSI Initiators

Field	Description
Name or IP Address	A character string that is a globally unique identifier for the node represented by this entry.
VSAN Membership	The list of configured VSANs the node represented by this entry can access.
Dynamic	If true, then the node represented by this entry is automatically discovered.
Initiator Type	Indicates whether the node is a host that participates in iSCSI load-balancing.
Persistent Node WWN	If true, then the same FC address is assigned to the node if it were to be represented again in the FC domain with the same node name. Note that the node FC address is either automatically assigned or manually configured.
SystemAssigned Node WWNN	If true, the FC address is automatically assigned to this node. If false, then the FC address has to be configured manually.
Node WWN	The persistent FC address of the node.
Persistent Port WWN	If true, then the same FC address is assigned to the ports of the node if it were to be represented again in the FC domain with the same node name.
Port WWN	All the FC port addresses associated with this node.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
AuthUser	This is the only CHAP user name that the initiator is allowed to log in with.
Target UserName	(Optional) The user name to be used for login. If you do not supply a username, the global user name is used.
Target Password	(Optional) The password to be used for login. If you do not supply a password, the global password is used.
Load Metric	A configured load metric of this iSCSI initiator for the purpose of iSCSI load balancing.
Auto Zone Name	The zone name that is used when the system creates automatic zone for this initiator's specific list of targets.

iSCSI Targets

Field	Description
Dynamically Import FC Targets	Check this option to dynamically import FC targets into the iSCSI domain. A target is not imported if it already exists in the iSCSI domain.
iSCSI Name	The iSCSI name of the node represented by this entry.
Dynamic	Indicates if the node represented by this entry was either automatically discovered or configured manually.
Primary Port WWN	The FC address for this target.
Secondary Port WWN	The optional secondary FC address for this target. This is the FC address used if the primary cannot be reached.
LUN Map iSCSI	The configured default logical unit number of this LU.
LUN Map FC Primary	The logical unit number of the remote LU for the primary port address.
LUN Map FC Secondary	The logical unit number of the remote LU for the secondary port address.
Initiator Access All	If true, then all the initiators can access this target even those which are not in the initiator permit list of this target. If false, then only initiators which are in the permit list are allowed access to this target.
Initiator Access List	Lists all the iSCSI nodes that are permitted to access the node represented by this entry. If AllAllowed is false and the value of List is empty, then no initiators are allowed to access this target.
Advertised Interfaces	Lists all the interfaces on which the target could be advertised.
Trespass Mode	The trespass mode for this node. Every iSCSI target represents one or more port(s) on the FC target. If true, the node instructs the FC node to present all LUN I/O requests to secondary port if the primary port is down.
RevertToPrimaryPort	Indicates if it is required to revert back to primary port if the FC target comes back online.

Send documentation comments to dcnm-san-docfeedback@cisco.com

iSCSI Session Initiators

Field	Description
Name or IP Address	The name or IP address of the initiator port.
Alias	The initiator alias acquired at login.

iSCSI Global

Field	Description
AuthMethod	The authentication method.
InitiatorIdleTimeout	The time for which the gateway (representing a FC target) waits from the time of last iSCSI session to a iSCSI initiator went down, before purging the information about that iSCSI initiator.
iSLB ZonesetActivate	Checking this option performs automatic zoning associated with the initiator targets
DynamicInitiator	This field determines how dynamic iSCSI initiators are created. Selecting the iSCSI option (default) creates dynamic iSCSI initiators. If you select iSLB then the an iSLB dynamic initiator is created. Selecting the deny option does not allow dynamic creation of the initiators.
Target UserName	The default user name used for login. If an initiator user name is specified, that user name is used instead.
Target Password	The default password used for login. If an initiator password is specified, that password is used instead.

iSCSI Session Statistics

Field	Description
PDU Command	The count of Command PDUs transferred on this session.
PDU Response	The count of Response PDUs transferred on this session.
Data Tx	The count of data bytes that were transmitted by the local iSCSI node on this session.
Data Rx	The count of data bytes that were received by the local iSCSI node on this session.
Errors Digest	Authentication errors.
Errors CxnTimeout	Connection timeouts.

Send documentation comments to dcnm-san-docfeedback@cisco.com

iSCSI iSLB VRRP

Field	Description
VrId, IpVersion	The virtual router number and the IP version (IPv4, IPv6, or DNS).
Load Balance	Indicates whether load balancing is enabled.

iSCSI Initiator Access

Field	Description
Initiator Name	The iSCSI node name.

iSCSI Initiator PWWN

Field	Description
Port WWN	The FC address for this entry.

iSCSI Sessions

Field	Description
Type	Type of iSCSI session: <ul style="list-style-type: none">normal—Session is a normal iSCSI sessiondiscovery—Session is being used only for discovery.
TargetName	If Direction is Outbound, this will contain the name of the remote target.
Vsan ID	The VSAN to which this session belongs to.
ISID	The initiator-defined portion of the iSCSI session ID.
TSIH	The target-defined identification handle for this session.

iSCSI Sessions Detail

Field	Description
ConnectionNumber	The number of transport protocol connections that currently belong to this session.
ImmediateData	Whether the initiator and target have agreed to support immediate data on this session.

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

Field	Description
Initial	If true, the initiator must wait for a Ready-To-Transfer before sending to the target. If false, the initiator may send data immediately, within limits set by FirstBurstSize and the expected data transfer length of the request.
MaxOutstanding	The maximum number of outstanding Ready-To-Transfers per task within this session.
First	The maximum length supported for unsolicited data sent within this session.
Max	The maximum number of bytes which can be sent within a single sequence of Data-In or Data-Out PDUs.
Sequence	If false, indicates that iSCSI data PDU sequences may be transferred in any order. If true indicates that data PDU sequences must be transferred using continuously increasing offsets, except during error recovery.
PDU	If false, iSCSI data PDUs within sequences may be in any order. If true indicates that data PDUs within sequences must be at continuously increasing addresses, with no gaps or overlay between PDUs.

Additional References

For additional information related to implementing VSANs, see the following section:

- [Related Document, page 1-8](#)
- [Standards, page 1-8](#)
- [RFCs, page 1-9](#)
- [MIBs, page 1-9](#)

Related Document

Related Topic	Document Title
Cisco MDS 9000 Family Command Reference	<i>Cisco MDS 9000 Family Command Reference</i>

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

Send documentation comments to dcnm-san-docfeedback@cisco.com

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified.	—

MIBs

MIBs	MIBs Link
<ul style="list-style-type: none">CISCO-SCSI-FLOW-MIBCISCO-SCSI-MIB	To locate and download MIBs, go to the following URL: http://www.cisco.com/en/US/products/ps5989/prod_technical_reference_list.html

Send documentation comments to dcnm-san-docfeedback@cisco.com



CHAPTER 1

Configuring FICON

Fibre Connection (FICON) interface capabilities enhance the Cisco MDS 9000 Family by supporting both open systems and mainframe storage network environments. The control unit port (CUP) also is supported, which allows in-band management of the switch from FICON processors.



Note

Cisco Fabric Manager release 3.x does not support FICON management of Cisco MDS 9000 Family switches running SAN-OS release 2.(x).

This chapter includes the following topics:

- [Information About FICON, page 1-1](#)
- [Licensing Requirements for FICON, page 1-18](#)
- [Guidelines and Limitations, page 1-19](#)
- [Default Settings, page 1-20](#)
- [Configuring FICON, page 1-20](#)
- [Configuring FICON Ports, page 1-27](#)
- [Verifying FICON Configuration, page 1-34](#)
- [Field Descriptions for FICON, page 1-37](#)
- [Additional References, page 1-40](#)
- [Feature History for FICON, page 1-41](#)

Information About FICON

The Cisco MDS 9000 Family supports the Fibre Channel Protocol (FCP), FICON, iSCSI, and FCIP capabilities within a single, high-availability platform (see [Figure 1-1](#)).

The FICON feature is not supported on:

- Cisco MDS 9120 switches
- Cisco MDS 9124 switches
- Cisco MDS 9140 switches
- The 32-port Fibre Channel switching module
- Cisco Fabric Switch for HP c-Class BladeSystem


Send documentation comments to dcnm-san-docfeedback@cisco.com

- Cisco Fabric Switch for IBM BladeSystem

FCP and FICON are different FC4 protocols and their traffic is independent of each other. Devices using these protocols should be isolated using VSANs.

The fabric binding feature helps prevent unauthorized switches from joining the fabric or disrupting current fabric operations (refer to the *Cisco MDS 9000 Family NX-OS Security Configuration Guide*). The Registered Link Incident Report (RLIR) application provides a method for a switch port to send an LIR to a registered Nx port.

Figure 1-1 Shared System Storage Network



This section includes the following topics:

- [FICON Requirements, page 1-3](#)
- [Cisco MDS-Specific FICON Advantages, page 1-3](#)
- [FICON Cascading, page 1-7](#)
- [FICON VSAN Prerequisites, page 1-7](#)
- [FICON Port Numbering, page 1-8](#)
- [Default FICON Port Numbering Scheme, page 1-8](#)
- [Port Addresses, page 1-10](#)
- [Implemented and Unimplemented Port Addresses, page 1-10](#)
- [About the Reserved FICON Port Numbering Scheme, page 1-11](#)
- [Installed and Uninstalled Ports, page 1-11](#)
- [About Port Numbers for FCIP and PortChannel, page 1-11](#)
- [FC ID Allocation, page 1-12](#)
- [About Enabling FICON on a VSAN, page 1-12](#)
- [FICON Information Refresh, page 1-13](#)
- [About FICON Device Allegiance, page 1-13](#)
- [Automatically Saving the Running Configuration, page 1-13](#)
- [Port Prohibiting, page 1-14](#)
- [About RLIR, page 1-14](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- [Port Swapping, page 1-16](#)
- [FICON Tape Acceleration, page 1-16](#)
- [CUP In-Band Management, page 1-18](#)

FICON Requirements

The FICON feature has the following requirements:

- You can implement FICON features in the following switches:
 - Any switch in the Cisco MDS 9500 Series
 - Any switch in the Cisco MDS 9200 Series (including the Cisco MDS 9222i Multiservice Modular Switch)
 - Cisco MDS 9134 Multilayer Fabric Switch
 - MDS 9000 Family 18/4-Port Multiservice Module
- You need the MAINFRAME_PKG license to configure FICON parameters.
- To extend your FICON configuration over a WAN link using FCIP, you need the appropriate SAN_EXTN_OVER_IP license for the module you are using. For more information, refer to the *Cisco NX-OS Family Licensing Guide*.

Cisco MDS-Specific FICON Advantages

This section explains the additional FICON advantages in Cisco MDS switches and includes the following topics:

- [Fabric Optimization with VSANs, page 1-3](#)
- [FCIP Support, page 1-5](#)
- [PortChannel Support, page 1-5](#)
- [VSANs for FICON and FCP Mixing, page 1-5](#)
- [Cisco MDS-Supported FICON Features, page 1-5](#)

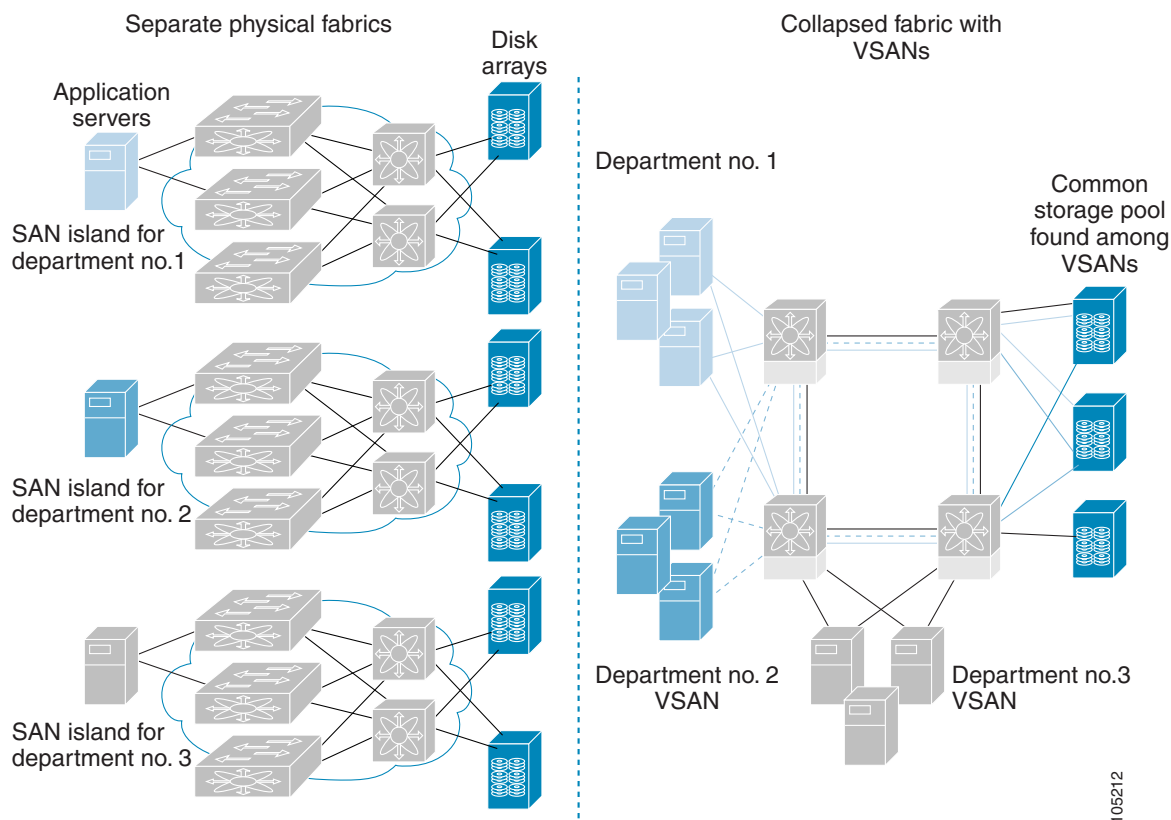
Fabric Optimization with VSANs

Generally, separate physical fabrics have a high level of switch management and have a higher implementation cost. The ports in each island also may be over-provisioned depending on the fabric configuration.

By using the Cisco MDS-specific VSAN technology, you can have greater efficiency between these physical fabrics by lowering the cost of over-provisioning and reducing the number of switches to be managed. VSANs also help you to move unused ports nondisruptively and provide a common redundant physical infrastructure (see [Figure 1-2](#)).

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 1-2 VSAN-Specific Fabric Optimization



VSANs enable global SAN consolidation by allowing you to convert existing SAN islands into virtual SAN islands on a single physical network. It provides hardware-enforced security and separation between applications or departments to allow coexistence on a single network. It also allows virtual rewiring to consolidate your storage infrastructure. You can move assets between departments or applications without the expense and disruption of physical relocation of equipment.



Note

You can configure VSANs in any Cisco MDS switch, but you only can enable FICON in up to eight of these VSANs. The number of VSANs configured depends on the platform.

Mainframe users can think of VSANs as being like FICON LPARs in the MDS SAN fabric. You can partition switch resources into FICON LPARs (VSANs) that are isolated from each other, in much the same way that you can partition resources on a zSeries or DS8000. Each VSAN has its own set of fabric services (such as fabric server and name server), FICON CUP, domain ID, Fabric Shortest Path First (FSPF) routing, operating mode, IP address, and security profile.

FICON LPARs can span line cards and are dynamic in size. For example, one FICON LPAR with 10 ports can span 10 different line cards. FICON LPARs can also include ports on more than one switch in a cascaded configuration. The consistent fairness of the Cisco MDS 9000 switching architecture means that “all ports are created equal,” simplifying provisioning by eliminating the “local switching” issues seen on other vendors’ platforms.

Addition of ports to a FICON LPAR is a nondisruptive process. The maximum number of ports for a FICON LPAR is 255 due to FICON addressing limitations.

Send documentation comments to dcnm-san-docfeedback@cisco.com

FCIP Support

The multilayer architecture of the Cisco MDS 9000 Family enables a consistent feature set over a protocol-agnostic switch fabric. Cisco MDS 9500 Series and 9200 Series switches transparently integrate Fibre Channel, FICON, and Fibre Channel over IP (FCIP) in one system. The FICON over FCIP feature enables cost-effective access to remotely located mainframe resources. With the Cisco MDS 9000 Family platform, storage replication services such as IBM PPRC and XRC can be extended over metro to global distances using ubiquitous IP infrastructure which simplifies business continuance strategies.

Refer to the *Cisco MDS 9000 Family NX-OS IP Services Configuration Guide*.

PortChannel Support

The Cisco MDS implementation of FICON provides support for efficient utilization and increased availability of Inter-Switch Links (ISLs) necessary to build stable large-scale SAN environments. PortChannels ensure an enhanced ISL availability and performance in Cisco MDS switches.

Refer to the *Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide* for more information on PortChannels.

VSANs for FICON and FCP Mixing

Cisco MDS 9000 Family FICON-enabled switches simplify deployment of even the most complex mixed environments. Multiple logical FICON, Z-Series Linux/FCP, and Open-Systems Fibre Channel Protocol (FCP) fabrics can be overlaid onto a single physical fabric by simply creating VSANs as required for each service. VSANs provide both hardware isolation and protocol specific fabric services, eliminating the complexity and potential instability of zone-based mixed schemes.

By default, the FICON feature is disabled in all switches in the Cisco MDS 9000 Family. When the FICON feature is disabled, FC IDs can be allocated seamlessly. Mixed environments are addressed by the Cisco NX-OS software. The challenge of mixing FCP and FICON protocols are addressed by Cisco MDS switches when implementing VSANs.

Switches and directors in the Cisco MDS 9000 Family support FCP and FICON protocol mixing at the port level. If these protocols are mixed in the same switch, you can use VSANs to isolate FCP and FICON ports.



Tip

When creating a mixed environment, place all FICON devices in one VSAN (other than the default VSAN) and segregate the FCP switch ports in a separate VSAN (other than the default VSAN). This isolation ensures proper communication for all connected devices.

Cisco MDS-Supported FICON Features

The Cisco MDS 9000 Family FICON features include:

- Flexibility and investment protection—The Cisco MDS 9000 Family shares common switching and service modules across the Cisco MDS 9500 Series and the 9200 Series.

Refer to the *Cisco MDS 9500 Series Hardware Installation Guide* and the *Cisco MDS 9200 Series Hardware Installation Guide*.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- High-availability FICON-enabled director—The Cisco MDS 9500 Series combines nondisruptive software upgrades, stateful process restart and failover, and full redundancy of all major components for a new standard in director-class availability. It supports up to 528 autosensing, 4/2/1-Gbps, 10-Gbps, FICON or FCP ports in any combination in a single chassis. Refer to the *Cisco MDS 9000 Family NX-OS High Availability and Redundancy Configuration Guide*.
- Infrastructure protection—Common software releases provide infrastructure protection across all Cisco MDS 9000 platforms. Refer to the *Cisco MDS 9000 Family NX-OS Software Upgrade and Downgrade Guide*.
- VSAN technology—The Cisco MDS 9000 Family provides VSAN technology for hardware-enforced, isolated environments within a single physical fabric for secure sharing of physical infrastructure and enhanced FICON mixed support. See [Chapter 1, “Configuring and Managing VSANs.”](#)
- Port-level configurations—There are BB_credits, beacon mode, and port security for each port. Refer to the *Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide* for information about buffer-to-buffer credits, beacon LEDs, and trunking.
- Alias name configuration—Provides user-friendly aliases instead of the WWN for switches and attached node devices. See [Chapter 1, “Configuring and Managing Zones.”](#)
- Comprehensive security framework—The Cisco MDS 9000 Family supports RADIUS and TACACS+ authentication, Simple Network Management Protocol Version 3 (SNMPv3), role-based access control, Secure Shell Protocol (SSH), Secure File Transfer Protocol (SFTP), VSANs, hardware-enforced zoning, ACLs, fabric binding, Fibre Channel Security Protocol (FC-SP), LUN zoning, read-only zones, and VSAN-based access control. Refer to the *Cisco MDS 9000 Family NX-OS Security Configuration Guide* for information about RADIUS, TACACS+, FC-SP, and DHCHAP.
- Traffic encryption—IPsec is supported over FCIP. You can encrypt FICON and Fibre Channel traffic that is carried over FCIP. Refer to the *Cisco MDS 9000 Family NX-OS Security Configuration Guide*.
- Local accounting log—View the local accounting log to locate FICON events. For more information about MSCHAP authentication, and local AAA services, refer to the *Cisco MDS 9000 Family NX-OS Security Configuration Guide*.
- Unified storage management—Cisco MDS 9000 FICON-enabled switches are fully IBM CUP standard compliant for in-band management using the IBM S/A OS/390 I/O operations console. See the [“CUP In-Band Management” section on page 1-18](#).
- Port address-based configurations—Configure port name, blocked or unblocked state, and the prohibit connectivity attributes can be configured on the ports. See the [“Configuring FICON Ports” section on page 1-27](#).
- You can display the following information:
 - Individual Fibre Channel ports, such as the port name, port number, Fibre Channel address, operational state, type of port, and login data.
 - Nodes attached to ports.
 - Port performance and statistics.

See the [“Calculating FICON Flow Load Balance” section on page 1-33](#).

- Configuration files—Store and apply configuration files. See the [“FICON Configuration Files” section on page 1-15](#).
- FICON and Open Systems Management Server features if installed. —See the [“VSANs for FICON and FCP Mixing” section on page 1-5](#).
- Enhanced cascading support—See the [“CUP In-Band Management” section on page 1-18](#).

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Date and time—Set the date and time on the switch. See the “[Allowing the Host to Control the Timestamp](#)” section on page 1-25.
- Configure SNMP trap recipients and community names—See the “[Configuring SNMP Control of FICON Parameters](#)” section on page 1-26.
- Call Home configurations—Configure the director name, location, description, and contact person. Refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*.
- Configure preferred domain ID, FC ID persistence, and principal switch priority—For information about configuring domain parameters, refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*.
- Sophisticated SPAN diagnostics—The Cisco MDS 9000 Family provides industry-first intelligent diagnostics, protocol decoding, and network analysis tools as well as integrated Call Home capability for added reliability, faster problem resolution, and reduced service costs. For information about monitoring network traffic using SPAN, refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*.
- Configure R_A_TOV, E_D_TOV— See the “[Fibre Channel Time-Out Values](#)” section on page 1-2.
- Director-level maintenance tasks—Perform maintenance tasks for the director including maintaining firmware levels, accessing the director logs, and collecting data to support failure analysis. For information about monitoring system processes and logs refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*.

FICON Cascading

The Cisco MDS NX-OS software allows multiple switches in a FICON network. To configure multiple switches, you must enable and configure fabric binding in that switch (see the “[Calculating FICON Flow Load Balance](#)” section on page 1-33 and refer to the *Cisco MDS 9000 Family NX-OS Security Configuration Guide*).

FICON VSAN Prerequisites

To ensure that a FICON VSAN is operationally up, be sure to verify the following requirements:

- Set the default zone to permit, if you are not using the zoning feature. See the “[About the Default Zone](#)” section on page 1-6.
- Enable in-order delivery on the VSAN. See [Chapter 1, “Configuring Fibre Channel Routing Services and Protocols.”](#)
- Enable (and if required, configure) fabric binding on the VSAN. See the “[Calculating FICON Flow Load Balance](#)” section on page 1-33. For more information about Fabric Binding, refer to the *Cisco MDS 9000 Family NX-OS Security Configuration Guide*.
- Verify that conflicting persistent FC IDs do not exist in the switch. For information about configuring domain parameters, refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*.
- Verify that the configured domain ID and requested domain ID match. For information about configuring domain parameters, refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Add the CUP (area FE) to the zone, if you are using zoning. See the “CUP In-Band Management” section on page 1-18.

If any of these requirements are not met, the FICON feature cannot be enabled.

FICON Port Numbering


With reference to the FICON feature, ports in Cisco MDS switches are identified by a statically defined 8-bit value known as the port number. A maximum of 255 port numbers are available. You can use the following port numbering schemes:

- Default port numbers based on the chassis type
- Reserved port numbers

Default FICON Port Numbering Scheme

Default FICON port numbers are assigned by the Cisco MDS NX-OS software based on the module and the slot in the chassis. The first port in a switch always starts with a zero (0) (see Figure 1-3).

Figure 1-3 Default FICON Port Number in Numbering on the Cisco MDS 9000 Family Switch



The default FICON port number is assigned based on the front panel location of the port and is specific to the slot in which the module resides. Thirty-two (32) port numbers are assigned to each slot on all Cisco MDS 9000 Family switches except for the Cisco MDS 9513 Director, which has 16 port numbers assigned for each slot. These default numbers are assigned regardless of the module’s physical presence in the chassis, the port status (up or down), or the number of ports on the module (4, 12, 16, 24, or 48). If a module has fewer ports than the number of port numbers assigned to the slot, then the excess port numbers are unused. If a module has more ports than the number of port numbers assigned to the slot, the excess ports cannot be used for FICON traffic unless you manually assign the port numbers.

Send documentation comments to dcnm-san-docfeedback@cisco.com

**Note**

Follow the steps in “Assigning FICON Port Numbers to Slots” section on page 1-21 to make use of excess ports by manually assigning more port numbers to the slots. Before doing this, however, we recommend that you review the default port number assignments for Cisco MDS 9000 switches shown in Table 1-3 on page 1-20 Table 1-1, and that you read the following sections to gain a complete understanding of FICON port numbering: “About the Reserved FICON Port Numbering Scheme” section on page 1-11, “FICON Port Numbering Guidelines” section on page 1-19, and “Assigning FICON Port Numbers to Slots” section on page 1-21.

**Note**

Only Fibre Channel, PortChannel, and FCIP ports are mapped to FICON port numbers. Other types of interfaces do not have a corresponding port number.

Table 1-1 lists the default port number assignment for the Cisco MDS 9000 Family of switches and directors.

Table 1-1 Default FICON Port Numbering in the Cisco MDS 9000 Family

Product	Slot Number	Implemented Port Allocation		Unimplemented Ports	Notes
		To Ports	To PortChannel/FCIP		
Cisco 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			
Cisco MDS 9222i Series	Slot 1	0 through 31	64 through 89	90 through 253 and port 255	The first 4, 12, 16, or 24 port numbers in a 4-port, 12-port, 16-port, or 24-port module are used and the rest remain unused. Extra 16 ports on 48-port modules are not allocated numbers.
	Slot 2	32 through 63			
Cisco MDS 9506 Director	Slot 1	0 through 31	128 through 153	154 through 253 and port 255	Supervisor modules are not allocated port numbers.
	Slot 2	32 through 63			
	Slot 3	64 through 95			
	Slot 4	96 through 127			
	Slot 5	None			
	Slot 6	None			
Cisco MDS 9134 Director	Slot 1	0 through 33	34 through 59	60 through 253 and port 255	

Send documentation comments to dcnm-san-docfeedback@cisco.com

Table 1-1 *Default FICON Port Numbering in the Cisco MDS 9000 Family (continued)*

Product	Slot Number	Implemented Port Allocation		Unimplemented Ports	Notes
		To Ports	To PortChannel/FCIP		
Cisco MDS 9509 Director	Slot 1	0 through 31	224 through 249	250 through 253 and port 255	The first 4, 12, 16, or 24 port numbers in a 4-port, 12-port, 16-port, or 24-port module are used and the rest remain unused. Extra 16 ports on 48-port modules are not allocated port numbers.
	Slot 2	32 through 63			
	Slot 3	64 through 95			
	Slot 4	96 through 127			Supervisor modules are not allocated port numbers.
	Slot 5	None			
	Slot 6	None			
	Slot 7	128 through 159			The first 4, 12, 16, or 24 port numbers are used for a 4-port, 12-port, 16-port, or 24-port module and the rest remain unused. Extra 16 ports on 48-port modules are not allocated port numbers.
	Slot 8	160 through 191			
	Slot 9	192 through 223			
Cisco MDS 9513 Director	Slot 1	0 through 15	224 through 249	250 through 253 and port 255	The first 4, 12 or 16 port numbers are used for a 4-port, 12-port or 16-port module and the rest remain unused. Extra ports on 24-port, 32-port, and 48-port modules are not allocated port numbers.
	Slot 2	16 through 31			
	Slot 3	32 through 47			
	Slot 4	48 through 63			Supervisor modules are not allocated port numbers.
	Slot 5	64 through 79			
	Slot 6	80 through 95			
	Slot 7	None			The first 4 or 12 port numbers are used for a 4-port or 12-port module and the rest remain unused. Extra ports on 24-port, 32-port, and 48-port modules are not allocated port numbers.
	Slot 8	None			
	Slot 9	96 through 111			
	Slot 10	112 through 127			
	Slot 11	128 through 143			
	Slot 12	144 through 159			
	Slot 13	160 through 175			

Port Addresses

By default, port numbers are the same as port addresses. You can swap the port addresses (see the [“Port Swapping”](#) section on page 1-16).

Implemented and Unimplemented Port Addresses

An implemented port refers to any port address that is assigned by default to a slot in the chassis (see [Table 1-3](#)). An unimplemented port refers to any port address that is not assigned by default to a slot in the chassis (see [Table 1-3](#)).

Send documentation comments to dcnm-san-docfeedback@cisco.com

About the Reserved FICON Port Numbering Scheme

A range of 250 port numbers are available for you to assign to all the ports on a switch. [Table 1-3](#) shows that you can have more than 250 physical ports on a switch and the excess ports do not have port numbers in the default numbering scheme. When you have more than 250 physical ports on your switch, you can have ports without a port number assigned if they are not in a FICON VSAN, or you can assign duplicate port numbers if they are not used in the same FICON VSAN. For example, you can configure port number 1 on interface fc1/1 in FICON VSAN 10 and fc10/1 in FICON VSAN 20.



Note

A VSAN can have a maximum of 250 port numbers.



Note

FICON port numbers are not changed for ports that are active. You must first disable the interfaces using the **shutdown** command.



Note

You can configure port numbers even when no module is installed in the slot.

Installed and Uninstalled Ports

An installed port refers to a port for which all required hardware is present. A specified port number in a VSAN can be implemented, and yet not installed, if any of the following conditions apply:

- The module is not present—For example, if module 1 is not physically present in slot 1 in a Cisco MDS 9509 Director, ports 0 to 31 are considered uninstalled.
- The small form-factor pluggable (SFP) port is not present—For example, if a 16-port module is inserted in slot 2 in a Cisco MDS 9509 Director, ports 48 to 63 are considered uninstalled.
- For slot 1, ports 0 to 31, or 0 to 15 have been assigned. Only the physical port fc1/5 with port number 4 is in VSAN 2. The rest of the physical ports are not in VSAN 2. The port numbers 0 to 249 are considered implemented for any FICON-enabled VSAN. Therefore, VSAN 2 has port numbers 0 to 249 and one physical port, fc1/4. The corresponding physical ports 0 to 3, and 5 to 249 are not in VSAN 2. When the FICON VSAN port address is displayed, those port numbers with the physical ports not in VSAN 2 are not installed (for example, ports 0 to 3, or 5 to 249).

Another scenario is if VSANs 1 through 5 are FICON-enabled, and trunking-enabled interface fc1/1 has VSANs 3 through 10, then port address 0 is uninstalled in VSAN 1 and 2.

- The port is part of a PortChannel—For example, if interface fc 1/1 is part of PortChannel 5, port address 0 is uninstalled in all FICON VSANs. See [Table 1-3](#).

About Port Numbers for FCIP and PortChannel

FCIP and PortChannels cannot be used in a FICON-enabled VSAN unless they are explicitly bound to a port number.

See the “[Configuring FICON Ports](#)” section on page 1-27, “[Configuring FICON Ports](#)” section on page 1-27, “[Reserving FICON Port Numbers for FCIP and PortChannel Interfaces](#)” section on page 1-21, and “[Binding Port Numbers to FCIP Interfaces](#)” section on page 1-28.

Send documentation comments to dcnm-san-docfeedback@cisco.com

You can use the default port numbers if they are available (see [Table 1-1 on page 1-9](#)) or if you reserve port numbers from the pool of port numbers that are not reserved for Fibre Channel interfaces (see the “FICON Port Numbering” section on page 1-8 and the “About the Reserved FICON Port Numbering Scheme” section on page 1-11).

FC ID Allocation

FICON requires a predictable and static FC ID allocation scheme. When FICON is enabled, the FC ID allocated to a device is based on the port address of the port to which it is attached. The port address forms the middle byte of the fabric address. Additionally, the last byte of the fabric address should be the same for all devices in the fabric. By default, the last byte value is 0 and can be configured (see the “Assigning FC ID Last Byte” section on page 1-24).




Note

You cannot configure persistent FC IDs in FICON-enabled VSANs.

Cisco MDS switches have a dynamic FC ID allocation scheme. When FICON is enabled or disabled on a VSAN, all the ports are shut down and restarted to switch from the dynamic to static FC IDs and vice versa (see [Figure 1-4](#)).

Figure 1-4 Static FC ID Allocation for FICON



113134

About Enabling FICON on a VSAN

By default FICON is disabled in all VSANs on the switch.

You can enable FICON on a per VSAN basis in one of the following ways:

- Manually address each prerequisite.
See the “[Information About FICON](#)” section on page 1-1.
- Use Device Manager (refer to the *Fabric Configuration Guide, Cisco DCNM for SAN*).

When you enable the FICON feature in Cisco MDS switches, the following restrictions apply:

- You cannot disable in-order delivery for the FICON-enabled VSAN.
- You cannot disable fabric binding or static domain ID configurations for the FICON-enabled VSAN.
- The load balancing scheme is changed to Source ID (SID)—Destination ID (DID). You cannot change it back to SID—DID—OXID.
- The IPL configuration file is automatically created.

Send documentation comments to dcnm-san-docfeedback@cisco.com

See the “FICON Configuration Files” section on page 1-15.

**Tip**

Using Device Manager, FICON auto-save can be invoked by multiple users logged on to the same FICON-enabled switch. Device Manager performs a periodic auto-save on any FICON-enabled switch causing increments in the FICON key counter. These increments highlight a change that has actually not occurred. To avoid this situation, we recommend that only one instance of Device Manager monitor a FICON-enabled switch.

FICON Information Refresh

When viewing FICON information through the Device Manager dialog boxes, you must manually refresh the display by clicking the **Refresh** button to see the latest updates. You need to take this step whether you configure FICON through the CLI or through the Device Manager.

There is no automatic refresh of FICON information. This information would be refreshed so often that it would affect performance.

About FICON Device Allegiance

FICON requires serialization of access among multiple mainframes, CLI, and SNMP sessions be maintained on Cisco MDS 9000 Family switches by controlling device allegiance for the currently executing session. Any other session is denied permission to perform configuration changes unless the required allegiance is available.

**Caution**

This task discards the currently executing session.

Automatically Saving the Running Configuration

Cisco MDS NX-OS provides an option to automatically save any configuration changes to the startup configuration. This ensures that the new configuration is present after a switch reboot. By default, the Active=Saved option is automatically enabled on any FICON VSAN.

[Table 1-2](#) displays the results of the Active = Saved option and the implicit copy from the running configuration to the startup configuration (**copy running start**) in various scenarios.

When the Active=Saved option is enabled in any FICON-enabled VSAN in the fabric, then the following apply (see Number 1 and 2 in [Table 1-2](#)):

- All configuration changes (FICON-specific or not) are automatically saved to persistent storage (implicit **copy running start**) and stored in the startup configuration.
- FICON-specific configuration changes are immediately saved to the IPL file (see the “FICON Configuration Files” section on page 1-15).

If the Active=Saved option is not enabled in any FICON-enabled VSAN in the fabric, then FICON-specific configuration changes are not saved in the IPL file and an implicit **copy running startup** command is not issued, you must explicitly save the running configuration to the startup configuration (see number 3 in [Table 1-2](#)).

Send documentation comments to dcnm-san-docfeedback@cisco.com

Table 1-2 Saving the Active FICON and Switch Configuration

Number	FICON-enabled VSAN?	active equals saved Enabled?	Implicitcopy running start Issued?	Notes
1	Yes	Yes (in all FICON VSANs)	Implicit	FICON changes written to the IPL file. Non-FICON changes saved to startup configuration and persistent storage.
2	Yes	Yes (even in one FICON VSAN)	Implicit	FICON changes written to IPL file for only the VSAN that has active equals saved option enabled. Non-FICON changes saved to startup configuration and persistent storage.
3	Yes	Not in any FICON VSAN	Not implicit	FICON changes are not written to the IPL file. Non-FICON changes are saved in persistent storage—only if you explicitly issue the copy running start command.
4	No	Not applicable		

Port Prohibiting

To prevent implemented ports from talking to each other, configure prohibits between two or more ports. If you prohibit ports, the specified ports are prevented from communicating with each other.



Tip

You cannot prohibit a PortChannel or FCIP interface.

Unimplemented ports are always prohibited. In addition, prohibit configurations are always symmetrically applied—if you prohibit port 0 from talking to port 15, port 15 is automatically prohibited from talking to port 0.



Note

If an interface is already configured in E or TE mode and you try to prohibit that port, your prohibit configuration is rejected. Similarly, if a port is not up and you prohibit that port, the port is not allowed to come up in E mode or in TE mode.

About RLIR

The Registered Link Incident Report (RLIR) application provides a method for a switch port to send a Link Incident Record (LIR) to a registered Nx port.

When an LIR is detected in FICON-enabled switches in the Cisco MDS 9000 Family from an RLIR Extended Link Service (ELS), the switch sends that record to the members in its Established Registration List (ERL).

In case of multiswitch topology, a Distribute Registered Link Incident Record (DRLIR) Inter-Link Service (ILS) is sent to all reachable remote domains along with the RLIR ELS. On receiving the DRLIR ILS, the switch extracts the RLIR ELS and sends it to the members of the ERL.

The Nx ports interested in receiving the RLIR ELS send the Link Incident Record Registration (LIRR) ELS request to the management server on the switch. The RLIRs are processed on a per-VSAN basis.

Send documentation comments to dcnm-san-docfeedback@cisco.com

The RLIR data is written to persistent storage when you copy the running configuration to the startup configuration.

FICON Configuration Files

You can save up to 16 FICON configuration files on each FICON-enabled VSAN (in persistent storage). The file format is proprietary to IBM. These files can be read and written by IBM hosts using the in-band CUP protocol. Additionally, you can use the Cisco MDS CLI or DCNM-SAN applications to operate on these FICON configuration files.

**Note**

Multiple FICON configuration files with the same name can exist in the same switch, provided they reside in different VSANs. For example, you can create a configuration file named XYZ in both VSAN 1 and VSAN 3.

When you enable the FICON feature in a VSAN, the switches always use the startup FICON configuration file, called IPL. This file is created with a default configuration as soon as FICON is enabled in a VSAN.

**Caution**

When FICON is disabled on a VSAN, all the FICON configuration files are irretrievably lost.

FICON configuration files contain the following configuration for each implemented port address:

- Block
- Prohibit mask
- Port address name

**Note**

Normal configuration files used by Cisco MDS switches include FICON-enabled attributes for a VSAN, port number mapping for PortChannels and FCIP interfaces, port number to port address mapping, port and trunk allowed VSAN configuration for ports, in-order guarantee, static domain ID configuration, and fabric binding configuration.

Only one user can access the configuration file at any given time:

- If this file is being accessed by user 1, user 2 cannot access this file.
- If user 2 does attempt to access this file, an error is issued to user 2.
- If user 1 is inactive for more than 15 seconds, the file is automatically closed and available for use by any other permitted user.

FICON configuration files can be accessed by any host, SNMP, or CLI user who is permitted to access the switch. The locking mechanism in the Cisco NX-OS software restricts access to one user at a time per file. This lock applies to newly created files and previously saved files. Before accessing any file, you must lock the file and obtain the file key. A new file key is used by the locking mechanism for each lock request. The key is discarded when the lock timeout of 15 seconds expires. The lock timeout value cannot be changed.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Port Swapping

The FICON port-swapping feature is only provided for maintenance purposes.

The FICON port-swapping feature causes all configurations associated with *old-port-number* and *new-port-number* to be swapped, including VSAN configurations.

Cisco MDS switches allow port swapping for nonexistent ports as follows:

- Only FICON-specific configurations (prohibit, block, and port address mapping) are swapped.
- No other system configuration is swapped.
- All other system configurations are only maintained for existing ports.
- If you swap a port in a module that has unlimited oversubscription ratios enabled with a port in a module that has limited oversubscription ratios, then you may experience a degradation in bandwidth.



Tip

If you check the Active=Saved check box **active equals saved** is enabled on any FICON VSAN, then the swapped configuration is automatically saved to startup. Otherwise, you must explicitly save the running configuration immediately after swapping the ports.

Once you swap ports, the switch automatically performs the following actions:

- Shuts down both the old and new ports.
- Swaps the port configuration.

If you attempt to bring the port up, you must explicitly shut down the port to resume traffic.



Note

To view the latest FICON information, you must click the **Refresh** button. See the [“Automatically Saving the Running Configuration”](#) section on page 1-26.

FICON Tape Acceleration

The sequential nature of tape devices causes each I/O operation to the tape device over an FCIP link to incur the latency of the FCIP link. Throughput drastically decreases as the round-trip time through the FCIP link increases, leading to longer backup windows. Also, after each I/O operation, the tape device is idle until the next I/O arrives. Starting and stopping of the tape head reduces the lifespan of the tape, except when I/O operations are directed to a virtual tape.

Cisco MDS NX-OS software provides acceleration for the following FICON tape write operations:

- The link between mainframe and native tape drives (both IBM and Sun/STK)
- The back-end link between the VSM (Virtual Storage Management) and tape drive (Sun/STK)

FICON tape acceleration over FCIP provides the following advantages:

- Efficiently utilizes the tape device by decreasing idle time
- More sustained throughput as latency increases
- Similar to FCP tape acceleration, and does not conflict with it

Send documentation comments to dcnm-san-docfeedback@cisco.com

**Note**

FICON tape read acceleration over FCIP is supported from Cisco MDS NX-OS Release 5.0(1). For more information refer to the [“Configuring FICON Tape Read Acceleration”](#) section on page 1-32.

Figure 1-5 through Figure 1-8 show supported configurations.

Figure 1-5 Host Directly Accessing IBM/STK (StorageTek) Library



Figure 1-6 Host Accessing Standalone IBM-VTS (Virtual Tape Server) /STK-VSM (Virtual Shared Memory)






Figure 1-7 Host Accessing Peer-to-Peer VTS (Virtual Tape Server)



Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 1-8 Host Accessing Peer-to-Peer VTS (Virtual Tape Server)



Note

For information about FCIP tape acceleration, refer to the *Cisco MDS 9000 Family NX-OS IP Services Configuration Guide*.

CUP In-Band Management

The CUP protocol configures access control and provides unified storage management capabilities from a mainframe computer. Cisco MDS 9000 FICON-enabled switches are fully IBM CUP standard compliant for in-band management using the IBM S/A OS/390 I/O operations console.



Note

The CUP specification is proprietary to IBM.

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.

Host communication includes control functions such as blocking and unblocking ports, as well as monitoring and error reporting functions.

Licensing Requirements for FICON

The following table shows the licensing requirements for this feature:

License	License Description
MAINFRAME_PKG	The mainframe license is required to enable FICON. For a complete explanation of the licensing scheme, see the <i>Cisco MDS 9000 Family NX-OS Licensing Guide</i> .
XRC_ACCL	The Extended Remote Copy (XRC) acceleration is required to activate FICON XRC acceleration on the Cisco MDS 9222i Switch and on the MSM-18/4 module in the Cisco MDS 9500 Series directors. For a complete explanation of the NX-OS licensing scheme and how to obtain and apply licenses, see the <i>Cisco MDS 9000 Family NX-OS Licensing Guide</i> .

Send documentation comments to dcnm-san-docfeedback@cisco.com

Guidelines and Limitations

This section includes the guidelines and limitations for this feature:

- [FICON Port Numbering Guidelines, page 1-19](#)
- [Port Swapping Guidelines, page 1-19](#)
- [FICON Tape Acceleration Configuration Guidelines, page 1-20](#)

FICON Port Numbering Guidelines

The following guidelines apply to FICON port numbers:

- Supervisor modules do not have port number assignments.
- Port numbers do not change based on TE ports. Since TE ports appear in multiple VSANs, chassis-wide unique port numbers should be reserved for TE ports.
- Each PortChannel must be explicitly associated with a FICON port number.
- When the port number for a physical PortChannel becomes uninstalled, the relevant PortChannel configuration is applied to the physical port.
- Each FCIP tunnel must be explicitly associated with a FICON port number. If the port numbers are not assigned for PortChannels or for FCIP tunnels, then the associated ports will not come up.

See the [“About Port Numbers for FCIP and PortChannel”](#) section on page 1-11.

Port Swapping Guidelines

Be sure to follow these guidelines when using the FICON port swapping feature:

- Port swapping is not supported for logical ports (PortChannels, FCIP links). Neither the *old-port-number* nor the *new-port-number* can be a logical port.
- Port swapping is not supported between physical ports that are part of a PortChannel. Neither the *old-port-number* nor the *new-port-number* can be a physical port that is part of a PortChannel.
- Before performing a port swap, the Cisco NX-OS software performs a compatibility check. If the two ports have incompatible configurations, the port swap is rejected with an appropriate reason code. For example, if a port with BB_credits as 25 is being swapped with an OSM port for which a maximum of 12 BB_credits is allowed (not a configurable parameter), the port swapping operation is rejected.
- Before performing a port swap, the Cisco NX-OS software performs a compatibility check to verify the extended BB_credits configuration.
- If ports have default values (for some incompatible parameters), then a port swap operation is allowed and the ports retain their default values.
- Port tracking information is not included in port swapping. This information must be configured separately (refer to the *Cisco MDS 9000 Family NX-OS Quality of Service Configuration Guide*).



Note

The 32-port module guidelines also apply for port swapping configurations (Refer to the *Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide*).

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

FICON Tape Acceleration Configuration Guidelines

FICON tape acceleration has the following configuration guidelines:

- In addition to the normal FICON configuration, FICON tape acceleration must be enabled on both ends of the FCIP interface. If only one end has FICON tape acceleration enabled, acceleration does not occur.
- FICON tape acceleration is enabled on a per VSAN basis.
- FICON tape acceleration cannot function if multiple ISLs are present in the same VSAN (PortChannels or FSPF load balanced).
- You can enable both Fibre Channel write acceleration and FICON tape acceleration on the same FCIP interface.

Enabling or disabling FICON tape acceleration disrupts traffic on the FCIP interface.

Default Settings

Table 1-3 lists the default settings for FICON features.

Table 1-3 **Default FICON Settings**

Parameters	Default
FICON feature	Disabled.
Port numbers	Same as port addresses.
FC ID last byte value	0 (zero).
EBCDIC format option	US-Canada.
Switch offline state	Hosts are allowed to move the switch to an offline state.
Mainframe users	Allowed to configure FICON parameters on Cisco MDS switches.
Clock in each VSAN	Same as the switch hardware clock.
Host clock control	Allows host to set the clock on this switch.
SNMP users	Configure FICON parameters.
Port address	Not blocked.
Prohibited ports	Ports 90–253 and 255 for the Cisco MDS 9200 Series switches. Ports 250–253 and 255 for the Cisco MDS 9500 Series switches.

Configuring FICON

By default FICON is disabled in all switches in the Cisco MDS 9000 Family. You can enable FICON on a per VSAN basis by using the Device Manager.

This section includes the following topics:

- [Assigning FICON Port Numbers to Slots, page 1-21](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- [Reserving FICON Port Numbers for FCIP and PortChannel Interfaces, page 1-21](#)
- [Enabling FICON on a VSAN, page 1-22](#)
- [Manually Enabling FICON on a VSAN, page 1-23](#)
- [Deleting FICON VSANs, page 1-23](#)
- [Suspending a FICON VSAN, page 1-23](#)
- [Configuring the code-page Option, page 1-24](#)
- [Assigning FC ID Last Byte, page 1-24](#)
- [Allowing the Host to Move the Switch Offline, page 1-25](#)
- [Allowing the Host to Change FICON Port Parameters, page 1-25](#)
- [Allowing the Host to Control the Timestamp, page 1-25](#)
- [Configuring SNMP Control of FICON Parameters, page 1-26](#)
- [Automatically Saving the Running Configuration, page 1-26](#)

Assigning FICON Port Numbers to Slots



Caution

When you assign, change, or release a port number, the port reloads.

Detailed Steps

To assign FICON port numbers to slots using Device Manager, follow these steps:

-
- | | |
|---------------|--|
| Step 1 | Click FICON and then select Port Numbers .
You see the FICON port number. |
| Step 2 | Enter the chassis slot port numbers in the Reserved Port Numbers field. |
| Step 3 | Click Apply . |
-

Reserving FICON Port Numbers for FCIP and PortChannel Interfaces

You must reserve port numbers for logical interfaces, such as FCIP and PortChannels, if you plan to use them.

Detailed Steps

To reserve FICON port numbers for FCIP and PortChannel interfaces using Device Manager, follow these steps:

-
- | | |
|---------------|--|
| Step 1 | Click FICON > Port Numbers .
You see the FICON port numbers dialog box. |
| Step 2 | Click the Logical tab to see the reserved port numbers for the slot. |

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Step 3** Enter the chassis slot port numbers. These are the reserved port numbers for one chassis slot. There can be up to 64 port numbers reserved for each slot in the chassis.
- Step 4** Click **Apply**.


Enabling FICON on a VSAN

Detailed Steps

To create a FICON-enabled VSAN, follow these steps:

- Step 1** Click the **Create VSAN** icon (see [Figure 1-9](#)).

Figure 1-9 Create VSAN Icon



You see the Create VSAN dialog box.

- Step 2** Select the switches you want to be in the VSAN.
- Step 3** Enter a VSAN ID.
- Step 4** Enter the name of the VSAN, if desired.
- Step 5** Select the type of load balancing, the interop value, and the administrative state for this VSAN.
- Step 6** Check the **FICON** check box.



Note You cannot enable interop modes on FICON-enabled VSANs.

- Step 7** Check the option, if appropriate, to enable fabric binding for the selected switches.
- Step 8** Check the All Ports Prohibited option if all ports in this VSAN are prohibited.
- Step 9** Click **Create** to create the VSAN.
- Step 10** Choose **Tools > Device Manager** to open Device Manager for each switch in the FICON VSAN.
- Step 11** Choose **FC > VSANs**.
You see the VSAN dialog box.
- Step 12** Enter the VSAN membership information.
- Step 13** Click the VSAN you want to become a FICON VSAN and select **Add** from the FICON drop-down menu.
- Step 14** Click **Apply** to save these changes.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Manually Enabling FICON on a VSAN

**Note**

This section describes the procedure to manually enable FICON on a VSAN. If you have already enabled FICON on the required VSAN using the automated setup (recommended), skip to the [“Automatically Saving the Running Configuration”](#) section on page 1-26.

Detailed Steps

To manually enable FICON on a VSAN, follow these steps:

-
- Step 1** Choose **VSAN > FICON**.
You see the FICON VSAN configuration information in the Information pane.
- Step 2** Select the switch in the VSAN on which you want to enable FICON.
- Step 3** Click **enable** from the Command drop-down menu.
- Step 4** Click the **Apply Changes** icon to save these changes.
-

Deleting FICON VSANs

Detailed Steps

To delete a FICON VSAN, follow these steps:

-
- Step 1** Select **All VSANS**.
You see the VSAN table in the Information pane.
- Step 2** Click anywhere in the row of the VSAN that you want to delete.
- Step 3** Click **Delete Row** to delete the VSAN.

**Note**

Deleting the VSAN will also delete the associated FICON configuration file, and the file cannot be recovered.

Suspending a FICON VSAN

Detailed Steps

To suspend a FICON VSAN, follow these steps:

-
- Step 1** Click **All VSANS**.
You see all the VSANs listed in the Information pane.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Step 2** Select the VSAN that you want to suspend.
 - Step 3** Set the Admin drop-down menu for a VSAN to **suspended**.
 - Step 4** Click the **Apply Changes** icon to save these changes.
-

**Note**

This command can be issued by the host if the host is allowed to do so (see the [“Allowing the Host to Move the Switch Offline”](#) section on page 1-25).

Configuring the code-page Option

FICON strings are coded in Extended Binary-Coded Decimal Interchange Code (EBCDIC) format. Refer to your mainframe documentation for details on the code-page options.

Cisco MDS switches support **international-5**, **france**, **brazil**, **germany**, **italy**, **japan**, **spain-latinamerica**, **uk**, and **us-canada** (default) EBCDIC format options.

Detailed Steps

To modify the code-page option using Device Manager, follow these steps:

- Step 1** Choose **FICON > VSANs**.
You see the FICON VSAN configuration dialog box. The VSANs tab is the default tab.
 - Step 2** From the CodePage drop-down menu, choose an option for the FICON VSAN you want to configure.
 - Step 3** Click **Apply** to save the changes.
-

Assigning FC ID Last Byte

Restrictions

If the FICON feature is configured in cascaded mode, the Cisco MDS switches use ISLs to connect to other switches.

Detailed Steps

To assign the last byte for the FC ID, follow these steps:

- Step 1** Choose **All VSANs > Domain Manager**.
 - Step 2** Click the **Persistent FCIDs** tab.
 - Step 3** Select **single** in the Mask column and then assign the entire FC ID at once. The single option allows you to enter the FC ID in the ##### format.
 - Step 4** Click the **Apply Changes** icon to save these changes.
-

Send documentation comments to dcnm-san-docfeedback@cisco.com

Allowing the Host to Move the Switch Offline

By default, hosts are allowed to move the switch to an offline state. To do this, the host sends a "Set offline" command (x'FD') to the CUP.

Detailed Steps

To allow the host (mainframe) to move the switch to an offline state, follow these steps:

-
- | | |
|---------------|---|
| Step 1 | Choose VSAN > FICON .
You see a list of switches under the Control tab in the Information pane. |
| Step 2 | Click the VSANs tab.
You see the FICON VSAN configuration information in the Information pane. |
| Step 3 | Check the Host Can Offline Sw check box to allow the mainframe to move a switch to the offline state. |
| Step 4 | Check the Host Can Sync Time check box to allow the mainframe to set the system time on the switch. |
| Step 5 | Click the Apply Changes icon to save the changes. |
-

Allowing the Host to Change FICON Port Parameters

By default, mainframe users are not allowed to configure FICON parameters on Cisco MDS switches—they can only query the switch.

Detailed Steps

To allow the host (mainframe) to configure FICON parameters on the Cisco MDS switch, follow these steps:

-
- | | |
|---------------|---|
| Step 1 | Choose VSAN > FICON .
You see a list of switches under the Control tab in the Information pane. |
| Step 2 | Click the VSANs tab.
You see the FICON VSAN configuration information in the Information pane. |
| Step 3 | Check the Port Control By Host check box to allow the mainframe to control a switch. |
| Step 4 | Click the Apply Changes icon to save the changes. |
-

Allowing the Host to Control the Timestamp

By default, the clock in each VSAN is the same as the switch hardware clock. Each VSAN in a Cisco MDS 9000 Family switch represents a virtual director. The clock and time present in each virtual director can be different. To maintain separate clocks for each VSAN, the Cisco NX-OS software maintains the difference of the VSAN-specific clock and the hardware-based director clock. When a host (mainframe)

Send documentation comments to dcnm-san-docfeedback@cisco.com

sets the time, the Cisco NX-OS software updates this difference between the clocks. When a host reads the clock, it computes the difference between the VSAN-clock and the current director hardware clock and presents a value to the mainframe.

Detailed Steps

To configure host (mainframe) control for the VSAN time stamp, follow these steps:

-
- Step 1** Choose **VSAN > FICON**.
You see a list of switches under the Control tab in the Information pane.
 - Step 2** Click the **VSANs** tab.
You see the FICON VSAN configuration information in the Information pane.
 - Step 3** Check the **Host Can Sync Time** checkbox to allow the mainframe to set the system time on the switch.
 - Step 4** Click the **Apply Changes** icon to save these changes.
-

Configuring SNMP Control of FICON Parameters

By default, SNMP users can configure FICON parameters using Cisco DCNM for SAN.

Restrictions

If you disable SNMP in the Cisco MDS switch, you cannot configure FICON parameters using DCNM-SAN.

Detailed Steps

To configure SNMP control of FICON parameters, follow these steps:

-
- Step 1** Choose **VSAN > FICON**.
You see a list of switches under the Control tab in the Information pane.
 - Step 2** Click the **VSANs** tab.
You see the FICON VSAN configuration information in the Information pane.
 - Step 3** Check the **Port Control By SNMP** checkbox to allow SNMP users to configure FICON on the switch.
 - Step 4** Click the **Apply Changes** icon to save these changes.
-

Automatically Saving the Running Configuration

Detailed Steps

To save the running configuration, follow these steps:

Send documentation comments to dcnm-san-docfeedback@cisco.com

-
- Step 1** Choose **VSAN > FICON**.
- You see a list of switches under the Control tab in the Information pane.
- Step 2** Click the **VSANs** tab.
- You see the FICON VSAN configuration information in the Information pane.
- Step 3** Check the **Active=Saved** check box to automatically save the running configuration to the startup configuration whenever there is a FICON configuration change.
- Step 4** Click the **Apply Changes** icon to save these changes.
-

Configuring FICON Ports

You can perform FICON configurations on a per-port address basis in the Cisco MDS 9000 Family switches.

Even if a port is uninstalled, the port address-based configuration is accepted by the Cisco MDS switch. This configuration is applied to the port when the port becomes installed.

This section includes the following topics:

- [Binding Port Numbers to PortChannels, page 1-27](#)
- [Binding Port Numbers to FCIP Interfaces, page 1-28](#)
- [Configuring Port Blocking, page 1-28](#)
- [Assigning a Port Address Name, page 1-29](#)
- [Specifying an RLIR Preferred Host, page 1-29](#)
- [Applying the Saved Configuration Files to the Running Configuration, page 1-30](#)
- [Editing FICON Configuration Files, page 1-30](#)
- [Copying FICON Configuration Files, page 1-30](#)
- [Swapping Ports, page 1-31](#)
- [Configuring FICON Tape Acceleration, page 1-31](#)
- [Configuring FICON Tape Read Acceleration, page 1-32](#)
- [Configuring XRC Acceleration, page 1-32](#)
- [Placing CUPs in a Zone, page 1-32](#)
- [Calculating FICON Flow Load Balance, page 1-33](#)
- [Receiving FICON Alerts, page 1-33](#)

Binding Port Numbers to PortChannels



Caution

All port number assignments to PortChannels or FCIP interfaces are lost (cannot be retrieved) when FICON is disabled on all VSANs.

You can bind (or associate) a PortChannel with a FICON port number to bring up that interface.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Binding Port Numbers to FCIP Interfaces

You can bind (or associate) an FCIP interface with a FICON port number to bring up that interface.

Configuring Port Blocking

If you block a port, the port is retained in the operationally down state. If you unblock a port, a port initialization is attempted. When a port is blocked, data and control traffic are not allowed on that port.

Physical Fibre Channel port blocks will continue to transmit an Off-line state (OLS) primitive sequence on a blocked port.



Note

The shutdown/no shutdown port state is independent of the block/no block port state.

Restrictions

You cannot block or prohibit the CUP port (0XFE). If a port is shut down, unblocking that port does not initialize the port.

Detailed Steps

To block or unblock port addresses in a VSAN using Device Manager, follow these steps:

-
- Step 1** Choose **FICON > VSANs**.
 - Step 2** Select a VSAN ID and click **Port Configuration**.
You see the FICON Port Configuration dialog box for the selected VSAN.
 - Step 3** Check the **Blocked** check box for the port that you want to block.
 - Step 4** Click **Apply** to save the changes.
-

Configuring the Default State for Port Prohibiting

By default, port prohibiting is disabled on the implemented interfaces on the switch. As of Cisco MDS SAN-OS Release 3.0(2), you can change the default port prohibiting state to enabled in VSANs that you create and then selectively disable port prohibiting on implemented ports, if desired. Also, only the FICON configuration files created after you change the default have the new default setting (see the [“FICON Configuration Files” section on page 1-15](#)).

Configuring Port Prohibiting

Detailed Steps

To prohibit port addresses in a VSAN using Device Manager, follow these steps:

-
- Step 1** Choose **FICON > VSANs**.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Step 2** Select a VASAN ID and click **Port Configuration**.
You see the FICON Port Configuration dialog box.
- Step 3** Set the port prohibit configuration for the selected FICON VSANs.
- Step 4** Click **Apply** to save these changes.
-

Assigning a Port Address Name



Note

To view the latest FICON information, you must click the **Refresh** button. See the [“Automatically Saving the Running Configuration”](#) section on page 1-26.

Detailed Steps

To assign a port address name in Device Manager, follow these steps:

- Step 1** Choose **FICON > VSANs**.
- Step 2** Select a VSAN ID and click **Port Configuration**.
You see the FICON Port Configuration dialog box.
- Step 3** Enter the Port Configuration information.
- Step 4** Click **Apply** to save the configuration information.
-

Specifying an RLIR Preferred Host

As of Cisco MDS SAN-OS Release 3.0(3), you can specify a preferred host to receive RLIR frames. The MDS switch sends RLIR frames to the preferred host only if it meets the following conditions:

- No host in the VSAN is registered for RLIR with the registration function set to “always receive.” If one or more hosts in the VSAN are registered as “always receive,” then RLIR sends only to these hosts and not to the configured preferred host.
- The preferred host is registered with the registration function set to “conditionally receive.”



Note

If all registered hosts have the registration function set to “conditionally receive,” then the preferred host receives the RLIR frames.

You can specify only one RLIR preferred host per VSAN. By default, the switch sends RLIR frames to one of the hosts in the VSAN with the register function set to “conditionally receive” if no hosts have the register function set to “always receive.”

Send documentation comments to dcnm-san-docfeedback@cisco.com

Applying the Saved Configuration Files to the Running Configuration

Detailed Steps

To apply the saved configuration files to the running configuration using Device Manager, follow these steps:

-
- Step 1** Choose **FICON > VSANs**.
 - Step 2** Click the **Files** tab.
You see the FICON Files dialog box.
 - Step 3** Highlight the file you want to apply and click **Apply File** to apply the configuration to the running configuration.
-

Editing FICON Configuration Files

The configuration file submode allows you to create and edit FICON configuration files. If a specified file does not exist, it is created. Up to 16 files can be saved. Each file name is restricted to eight alphanumeric characters.



Note

To view the latest FICON information, you must click the **Refresh** button. See the [“Automatically Saving the Running Configuration”](#) section on page 1-26.

Detailed Steps

To edit the contents of a specified FICON configuration file using Device Manager, follow these steps:

-
- Step 1** Choose **FICON > VSANs**.
 - Step 2** Click the **Files** tab.
You see the FICON VSANs dialog box.
 - Step 3** Select a VSAN ID and then click **Open** to edit the FICON configuration file.
 - Step 4** Select a VSAN ID and then click **Delete** to delete the FICON configuration file.
 - Step 5** Click **Apply** to apply the changed FICON configuration file.
-

Copying FICON Configuration Files

Detailed Steps

To copy an existing FICON configuration file using Device Manager, follow these steps:

-
- Step 1** Choose **FICON > VSANs**.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Step 2** Click the **Files** tab.
You see the FICON VSANs dialog box.
- Step 3** Click **Create** to create a FICON configuration file.
You see the Create FICON VSANs Files dialog box.
- a. Select a VSAN ID for the FICON VSAN you want to configure.
 - b. Enter the file name and the description.
 - c. Click **Create** to create the file.
- Step 4** Click **Copy** to copy the file to a new file.
- Step 5** Click **Apply** to apply the FICON configuration file.
-

Swapping Ports

Detailed Steps

To swap ports using Device Manager, follow these steps:

- Step 1** Select two Fibre Channel ports by holding down the **CTRL** key and clicking them.
- Step 2** Choose **FICON > Swap Selected Ports**.
-

Configuring FICON Tape Acceleration

Detailed Steps

To configure FICON tape acceleration over FCIP, follow these steps:

- Step 1** Expand **ISL** and then select **FCIP** in the Physical Attributes pane.
- Step 2** Click the **Tunnels** tab in the Information pane.
You see a list of available switches.
- Step 3** Click the **Create Row** icon to create an FCIP tunnel.
You see the Create FCIP Tunnel dialog box.
- Step 4** Configure the tunnel with the options.
- Step 5** Check the **TapeAccelerator** check box to enable FICON tape acceleration over this FCIP tunnel.
- Step 6** Click **Create**.
-

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuring FICON Tape Read Acceleration

All the configuration guidelines and restrictions applicable for FICON tape acceleration are also applicable for FICON tape read acceleration. Both FICON tape acceleration and FICON tape read acceleration can coexist.

Configuring XRC Acceleration

IBM z/OS Global Mirror eXtended Remote Copy (XRC) is supported on the MSM-18+4 modules. For XRC to function, XRC acceleration must be enabled on the FCIP tunnel interfaces on both ends. XRC acceleration is disabled by default.

Restrictions

XRC acceleration and FICON tape acceleration cannot be enabled on the same FCIP tunnel interface and cannot exist in the same VSAN.

Detailed Steps

To configure XRC acceleration on a FCIP tunnel interface, follow these steps:

-
- Step 1** Expand **ISL** and then select **FCIP** in the Physical Attributes pane.
 - Step 2** Click the **Tunnels(Advanced)** tab in the Information pane.
You see a list of available FCIP interfaces.
 - Step 3** Check the check box in the XRC Emulator column to enable XRC acceleration over the FCIP tunnel.
 - Step 4** Click **Apply**.
-

To configure XRC acceleration on an FCIP tunnel interface using Device Manager, follow these steps:


-
- Step 1** In the Device Manager window, click **IP** and then select **FCIP** from the menu.
 - Step 2** Click the **Tunnels(Advanced)** tab in the Information pane.
You see a list of FCIP interfaces.
 - Step 3** Check the check box in the XRC Emulator column to enable XRC acceleration over the FCIP tunnel.
 - Step 4** Click **Apply**.
-

Placing CUPs in a Zone

Detailed Steps

To place the CUP in a zone, follow these steps:

Send documentation comments to dcnm-san-docfeedback@cisco.com

-
- Step 1** In DCNM-SAN, choose **Zone > Edit Full Zoneset**, and then choose **Edit > Edit Default Zone Attributes** to set the default zone to permit for the required VSAN.
- Step 2** In Device Manager, choose **FC > Name Server...** for the required VSAN and obtain the FICON:CUP WWN.
-  **Note** If more than one FICON:CUP WWN exists in this fabric, be sure to add all the FICON:CUP pWWNs to the required zone.
-
- Step 3** In DCNM-SAN, choose **Zone > Edit Full Zoneset** and add the FICON:CUP pWWN to the zone database.
-

Calculating FICON Flow Load Balance

The FICON Flow Load Balance Calculator allows you to get the best load balancing configuration for your FICON flows. The calculator does not rely on any switch or flow discovery in the fabric. It is available from the DCNM-SAN Tools menu.

Detailed Steps

To use the FICON Flow Load Balance Calculator, follow these steps:

-
- Step 1** Choose **Tools > Flow Load Balance Calculator**.
- You see the Flow Load Balance Calculator.
- Step 2** Click **Add** to enter the source and destination(s) flows.
- Step 3** Enter source and destination using 2 byte hex (by domain and area IDs). You can copy and paste these IDs, and then edit them if required.
- Step 4** Enter (or select) the number of ISLs between the two switches (for example, between domain ID 0a and 0b).
- Step 5** Select a row to remove it and click **Remove**.
- Step 6** Select the module for which you are calculating the load balance.
- Step 7** Click **Calculate** to show the recommended topology.



Note If you change flows or ISLs, you must click **Calculate** to see the new recommendation.

Receiving FICON Alerts

To receive an alert to indicate any changes in the FICON configuration using Device Manager, follow these steps:

Send documentation comments to dcnm-san-docfeedback@cisco.com

-
- Step 1** Choose **FICON > VSANs**.
You see the FICON VSANs dialog box.
- Step 2** Check the **User Alert Mode** check box to receive an alert when the FICON configuration changes.
- Step 3** Click **Apply** to apply this change.
-

Verifying FICON Configuration

This section includes the following topics:

- [Viewing ESCON Style Ports, page 1-34](#)
- [Displaying RLIR Information, page 1-35](#)
- [Displaying FICON Configuration Files, page 1-35](#)
- [Displaying XRC Acceleration Statistics, page 1-35](#)
- [Displaying FICON Port Address Information, page 1-36](#)
- [Displaying IPL File Information, page 1-36](#)
- [Viewing the History Buffer, page 1-36](#)

Viewing ESCON Style Ports

To view the available and prohibited ESCON style ports using Device Manager, follow these steps:

-
- Step 1** Check the **ESCON Style** check box to see the available and prohibited ESCON style ports.
In [Figure 1-10](#), A stands for available and P stands for prohibited.
When the port address is highlighted red, it represents the E/TE port or multiple interfaces.

Figure 1-10 *ESCON Style*

Port Address	Name	Block	00	01	02	03	04	05	06	07	20	21	2D
00(fc1/1)	host1		A	A	A	A	A	A	A	P	A	A	P
01(fc1/2)		B	A	A	A	A	A	A	A	P	A	A	A
02(fc1/3)	host2		A	A	A	A	A	A	A	P	P	A	A
03(fc1/4)	host1		A	A	A	A	A	A	A	P	A	P	A
04(fc1/5)	host1		A	A	A	A	A	A	A	P	A	A	A
05(fc1/6)	host1		A	A	A	A	A	A	A	P	A	A	A
06(fc1/7)	host1		A	A	A	A	A	A	A	P	A	A	A
07(fc1/8)	storage1		P	P	P	P	P	P	P	P	P	P	P
20(fc2/1)	storage2		A	A	P	A	A	A	A	P	A	A	A
21(fc2/2)	storage3		A	A	A	P	A	A	A	P	A	A	A
2D(fc2/14)	storage4		P	A	A	A	A	A	A	P	A	A	A

11 row(s)

Buttons: Apply, Refresh, Help, Close

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Step 2** Click **Apply** to save the changes.
-

Displaying RLIR Information

To view RLIR information using Device Manager, follow these steps:

- Step 1** Choose **FICON > RLIR ERL**.
You see the Show RLIR ERL dialog box.
- Step 2** Click **Close** to close the dialog box.
-

Displaying FICON Configuration Files

To open and view configuration files in DCNM-SAN, follow these steps:

- Step 1** Choose **FICON > VSAN**.
You see the FICON configuration table in the Information pane.
- Step 2** Click the **Files** tab.
- Step 3** Select the file you want to open.
- Step 4** Click **Open**.
-

Displaying XRC Acceleration Statistics

To display XRC acceleration statistics, follow these steps:

- Step 1** Expand **ISL** and then select **FCIP** in the Physical Attributes pane.
- Step 2** Click the **XRC Statistics** tab in the Information pane.
You see the XRC session statistics.
-

To display XRC acceleration statistics using Device Manager, follow these steps:

- Step 1** In the Device Manager window, click **IP**, and then select **FCIP** from the menu.
- Step 2** Click the **XRC Statistics** tab in the Information pane.
You see the XRC session statistics.
-

Send documentation comments to dcnm-san-docfeedback@cisco.com

Displaying FICON Port Address Information

To display FICON port address information using Device Manager, follow these steps:

-
- Step 1** Choose **FICON > VSANs**.
You see the FICON VSANs dialog box.
 - Step 2** Select a VSAN ID and click **Port Configuration**.
You see the FICON Port Configuration dialog box.
 - Step 3** Click **Close** to close the dialog box.
-

Displaying IPL File Information

To display the IPL file information using Device Manager, follow these steps:

-
- Step 1** Select **VSANs** from the FICON menu.
 - Step 2** Click the **Files** tab.
You see the FICON VSANs dialog box.
 - Step 3** Select the file that you want to view and click **Open**.
-

Viewing the History Buffer

In the directory history buffer, the Key Counter column displays the 32-bit value maintained by Cisco MDS switches. This value is incremented when any port changes state in that VSAN. The key counter (a 32-bit value) is incremented when a FICON-related configuration is changed. Host programs can increment this value at the start of the channel program and then perform operations on multiple ports. The director history buffer keeps a log of which port address configuration was changed for each key-counter value.

The director history buffer provides a mechanism to determine the change in the port state from the previous time when a value was contained in the key counter.

To view the directory history buffer using Device Manager, follow these steps:

-
- Step 1** Choose **FICON > VSANs**.
You see the FICON VSANs dialog box.
 - Step 2** Click the **Director History** button.
You see the history buffer dialog box.
 - Step 3** Click **Close** to close the dialog box.
-

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field Descriptions for FICON

This section displays the field descriptions for this feature.

FICON VSANs

Field	Description
VSAN ID	Uniquely identifies a VSAN within a fabric.
Host Can Offline SW	If true, it allows the host to put the system offline.
Host Can Sync Time	If true, the host can set the system time.
Port Control by Host	If true, the host is allowed to alter FICON Director connectivity parameters.
Port Control by SNMP	If true, SNMP manager is allowed to alter FICON director connectivity parameters.
CUP Name	The name of the control unit device.
CUP Enable	Indicates whether the control unit device is enabled.
Domain ID	Specifies the domain ID of the switch.
CodePage	The Code Page used in this VSAN.
Character Set	Character set for the code page used in this VSAN.
Active=Saved	If true, the active to saved mode is enabled. All changes will be saved to NVRAM.
User Alert Mode	If true, FICON management stations will prompt on changes.
Device Allegiance	If CUP is in allegiance state with a channel, it cannot accept any commands from any logical paths. A CUP goes in an allegiance state when it accepts command from a channel and forms an allegiance with it until the successful completion of the channel program, at which point the CUP goes in an unlocked mode.
VSAN Time	The system time in the VSAN. This could be set either by the host or be the default global time in the FICON Director. The default global time is the local time in the FICON Director.
VSAN State	Controls the state of the ports belonging to a VSAN in the context of the FICON functionality.
VSAN Serial Number	The serial number of the FICON director for this VSAN.

FICON VSANs Files

Field	Description
Description	Configuration file description.
CUP Name	The name of the control unit device.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
Status	Locked indicates no change allowed. Unlocked indicates change allowed.
LastAccessed	The time this file was last accessed.
UserAlertMode	If true, director user alert mode is enabled.

Global

Field	Description
Default Port Prohibited	Check this option to block the default port.

FICON Port Attributes

Field	Description
TypeNumber	The type number for this FICON Director.
SerialNumber	The sequence number assigned to this FICON Director during manufacturing.
Tag	This is the identifier of the peer port. <ul style="list-style-type: none"> If the peer port's unit type is channel, then PortId will be the CHPID (Channel Path Identifier) of the channel path that contains this peer port. If the peer port is controlUnit, then PortId will be 0. If the peer port is fabric, then PortId will be port address of the interface on the peer switch.
FcId	The fabric Id of the other side port (initiator /target). This will be filled only in the case of Fabric ports.
Status	valid—If this information is current. old—If this information is cached. Click Clear Old Attributes to clear the cache.
Name	The FICON port name.
Manufacturer	The name of the company that manufactured this FICON Director.
ModelNumber	The model number for this FICON Director.
PlantOfMfg	The plant code that identifies the plant of manufacture of this FICON Director.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
UnitType	The peer type of the port that this port is communicating. ==Channel - host ==Control Unit - disk == Fabric - ISL
Alert	<p>Displays one of the following:</p> <ul style="list-style-type: none"> bitErrThreshExceeded lossOfSignalOrSync nosReceived primitiveSeqTimeOut invalidPrimitiveSeq <p>Click Clear to acknowledge and clear this alert.</p>

FICON Port Configuration

Field	Description
Show Installed Ports Only	If true, only physically available ports will be listed in the table.
ESCON Style	ESCON Style Port Configuration display is the Port Configuration table in DM displaying the ESCON Style Ports. In the table, A represents the available ports and P represents the prohibited ports.
Port/ Prohibit	Enter the FICON address of the port and the prohibited list. (This is an alternative to the table grid.)
Name	The port name of this port.
Block	If true, this port will be isolated.
Prohibit Grid	Click on the grid to add or remove the ability of ports to communicate with each other.

FICON Port Numbers

Field	Description
Module	The number of the module in the chassis.
Reserved Port Numbers (Physical)	The reserved port numbers for the module.
NumPorts	The number of ports reserved for that module.
Module Name	The name of the module.
Reserved Port Numbers (Logical)	Chassis slot port numbers. Reserved port numbers for one chassis slot. There can be up to 64 port numbers reserved for each slot in the chassis.

Send documentation comments to dcnm-san-docfeedback@cisco.com

FICON VSANs Director History

To view the latest FICON information, you must click the **Refresh** button.

Field	Description
KeyCounter	The key counter.
Ports Address Changed	The list of ports that have configuration change for a value of KeyCounter.

Additional References

For additional information related to implementing VSANs, see the following section:

- [Related Document, page 1-40](#)
- [Standards, page 1-40](#)
- [RFCs, page 1-41](#)
- [MIBs, page 1-41](#)

Related Document

Related Topic	Document Title
Cisco MDS 9000 Family Command Reference	<i>Cisco MDS 9000 Family Command Reference</i>

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

Send documentation comments to dcnm-san-docfeedback@cisco.com

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified.	—

MIBs

MIBs	MIBs Link
<ul style="list-style-type: none"> CISCO-FICON-MIB 	To locate and download MIBs, go to the following URL: http://www.cisco.com/en/US/products/ps5989/prod_technical_reference_list.html

Feature History for FICON

Table 1-4 lists the release history for this feature. Only features that were introduced or modified in 5.0(1a) or a later release appear in the table.

Table 1-4 Feature History for FICON

Feature Name	Releases	Feature Information
FICON Tape Read Acceleration	5.0(1a)	FICON tape acceleration over FCIP efficiently utilizes the tape device by decreasing idle time. The following sections provide information about this feature: <ul style="list-style-type: none"> FICON Tape Acceleration, page 1-16 Configuring FICON Tape Acceleration, page 1-31 Configuring FICON Tape Read Acceleration, page 1-32 The following commands were introduced or modified: ficon-tape-read-accelerator.
XRC Acceleration	4.2(1)	Added information about configuring XRC Acceleration.

Send documentation comments to dcnm-san-docfeedback@cisco.com



CHAPTER 1

Configuring Advanced Fabric Features

This chapter describes the advanced features provided in switches in the Cisco MDS 9000 Family. It includes the following sections:

- [Information About Common Information Model, page 1-1](#)
- [Guidelines and Limitations, page 1-7](#)
- [Default Settings, page 1-7](#)
- [Configuring Timer Across All VSANs, page 1-8](#)
- [Verifying the Advanced Features and Concepts Configuration, page 1-12](#)
- [Additional References, page 1-13](#)

Information About Common Information Model

Common Information Model (CIM) is an object-oriented information model that extends the existing standards for describing management information in a network/enterprise environment.



Note

CIM is not supported in Cisco MDS NX-OS Release 5.2(1), but is supported in Cisco DCNM Release 5.2(1).

CIM messages are independent of platform and implementation because they are encoded in N Extensible Markup Language (XML). CIM consists of a specification and a schema. The specification defines the syntax and rules for describing management data and integrating with other management models. The schema provides the actual model descriptions for systems, applications, networks, and devices.

For more information about CIM, refer to the specification available through the Distributed Management Task Force (DMTF) website at the following URL: <http://www.dmtf.org/>

For further information about Cisco MDS 9000 Family support for CIM servers, refer to the *Cisco MDS 9000 Family CIM Programming Reference Guide*.

A CIM client is required to access the CIM server. The client can be any client that supports CIM.

- [SSL Certificate Requirements and Format, page 1-2](#)
- [Fibre Channel Time-Out Values, page 1-2](#)
- [About fctimer Distribution, page 1-3](#)
- [Fabric Lock Override, page 1-3](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- [World Wide Names, page 1-3](#)
- [Link Initialization WWN Usage, page 1-4](#)
- [FC ID Allocation for HBAs, page 1-4](#)
- [Default Company ID List, page 1-4](#)
- [Switch Interoperability, page 1-5](#)
- [About Interop Mode, page 1-5](#)

SSL Certificate Requirements and Format

To limit access to the CIM server to authorized clients, you can enable the HTTPS transport protocol between the CIM server and client. On the switch side, you must install a Secure Socket Library (SSL) certificate generated on the client and enable the HTTPS server. Certificates may be generated using third-party tools, such as openssl (available for UNIX, Mac, and Windows), and may be certified by a CA or self-signed.

The SSL certificate that you install on the switch must meet the following requirements:

- The certificate file contains the certificate and the private key.
- The private key must be RSA type.
- The certificate file should be in Private Electronic Mail (PEM) style format and have .pem as the extension.

```
-----BEGIN CERTIFICATE-----  
(certificate goes here)  
-----END CERTIFICATE-----  
-----BEGIN RSA PRIVATE KEY-----  
(private key goes here)  
-----END RSA PRIVATE KEY-----
```

Only one certificate file can be installed at a time.

Fibre Channel Time-Out Values

You can modify Fibre Channel protocol related timer values for the switch by configuring the following time-out values (TOVs):

- Distributed services TOV (D_S_TOV)—The valid range is from 5,000 to 10,000 milliseconds. The default is 5,000 milliseconds.
- Error detect TOV (E_D_TOV)—The valid range is from 1,000 to 10,000 milliseconds. The default is 2,000 milliseconds. This value is matched with the other end during port initialization.
- Resource allocation TOV (R_A_TOV)—The valid range is from 5,000 to 10,000 milliseconds. The default is 10,000 milliseconds. This value is matched with the other end during port initialization.



Note

The fabric stability TOV (F_S_TOV) constant cannot be configured.

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

About fctimer Distribution

You can enable per-VSAN fctimer fabric distribution for all Cisco MDS switches in the fabric. When you perform fctimer configurations, and distribution is enabled, that configuration is distributed to all the switches in the fabric.

You automatically acquire a fabric-wide lock when you issue the first configuration command after you enabled distribution in a switch. The fctimer application uses the effective and pending database model to store or commit the commands based on your configuration.

Refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide* for more information on the CFS application.

Fabric Lock Override

If you have performed a fctimer fabric task and have forgotten to release the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your changes to the pending database are discarded and the fabric lock is released.



Tip

The changes are only available in the volatile directory and are subject to being discarded if the switch is restarted.

World Wide Names

The world wide name (WWN) in the switch is equivalent to the Ethernet MAC address. As with the MAC address, you must uniquely associate the WWN to a single device. The principal switch selection and the allocation of domain IDs rely on the WWN. The WWN manager, a process-level manager residing on the switch's supervisor module, assigns WWNs to each switch.

Cisco MDS 9000 Family switches support three network address authority (NAA) address formats (see [Table 1-1](#)).

Table 1-1 Standardized NAA WWN Formats

NAA Address	NAA Type	WWN Format	
IEEE 48-bit address	Type 1 = 0001b	000 0000 0000b	48-bit MAC address
IEEE extended	Type 2 = 0010b	Locally assigned	48-bit MAC address
IEEE registered	Type 5 = 0101b	IEEE company ID: 24 bits	VSID: 36 bits



Caution

Changes to the world-wide names should be made by an administrator or individual who is completely familiar with switch operations.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Link Initialization WWN Usage

Exchange Link Protocol (ELP) and Exchange Fabric Protocol (EFP) use WWNs during link initialization. The usage details differ based on the Cisco NX-OS software release.

Both ELPs and EFPs use the VSAN WWN by default during link initialization. However, the ELP usage changes based on the peer switch's usage:

- If the peer switch ELP uses the switch WWN, then the local switch also uses the switch WWN.
- If the peer switch ELP uses the VSAN WWN, then the local switch also uses the VSAN WWN.

**Note**

As of Cisco SAN-OS Release 2.0(2b), the ELP is enhanced to be compliant with FC-SW-3.

FC ID Allocation for HBAs

Fibre Channel standards require a unique FC ID to be allocated to an N port attached to a Fx port in any switch. To conserve the number of FC IDs used, Cisco MDS 9000 Family switches use a special allocation scheme.

Some HBAs do not discover targets that have FC IDs with the same domain and area. Prior to Cisco SAN-OS Release 2.0(1b), the Cisco SAN-OS software maintained a list of tested company IDs that do not exhibit this behavior. These HBAs were allocated with single FC IDs, and for others a full area was allocated.

The FC ID allocation scheme available in Release 1.3 and earlier, allocates a full area to these HBAs. This allocation isolates them to that area and are listed with their pWWN during a fabric login. The allocated FC IDs are cached persistently and are still available in Cisco SAN-OS Release 2.0(1b) (see the [“FC ID Allocation for HBAs” section on page 1-4](#)).

To allow further scalability for switches with numerous ports, the Cisco NX-OS software maintains a list of HBAs exhibiting this behavior. Each HBA is identified by its company ID (also known as Organizational Unique Identifier, or OUI) used in the pWWN during a fabric login. A full area is allocated to the N ports with company IDs that are listed, and for the others a single FC ID is allocated. Regardless of the kind (whole area or single) of FC ID allocated, the FC ID entries remain persistent.

Default Company ID List

All switches in the Cisco MDS 9000 Family that ship with Cisco SAN-OS Release 2.0(1b) or later, or NX-OS 4.1(1) contain a default list of company IDs that require area allocation. Using the company ID reduces the number of configured persistent FC ID entries. You can configure or modify these entries using the CLI.

Send documentation comments to dcnm-san-docfeedback@cisco.com

**Caution**

Persistent entries take precedence over company ID configuration. If the HBA fails to discover a target, verify that the HBA and the target are connected to the same switch and have the same area in their FC IDs, then perform the following procedure:

1. Shut down the port connected to the HBA.
2. Clear the persistent FC ID entry.
3. Get the company ID from the Port WWN.
4. Add the company ID to the list that requires area allocation.
5. Bring up the port.

The list of company IDs have the following characteristics:

- A persistent FC ID configuration always takes precedence over the list of company IDs. Even if the company ID is configured to receive an area, the persistent FC ID configuration results in the allocation of a single FC ID.
- New company IDs added to subsequent releases are automatically added to existing company IDs.
- The list of company IDs is saved as part of the running and saved configuration.
- The list of company IDs is used only when the fcinterop FC ID allocation scheme is in auto mode. By default, the interop FC ID allocation is set to auto, unless changed.

**Tip**

We recommend that you set the fcinterop FC ID allocation scheme to auto and use the company ID list and persistent FC ID configuration to manipulate the FC ID device allocation.

Switch Interoperability

Interoperability enables the products of multiple vendors to interact with each other. Fibre Channel standards guide vendors towards common external Fibre Channel interfaces.

If all vendors followed the standards in the same manner, then interconnecting different products would become a trivial exercise. However, not all vendors follow the standards in the same way, thus resulting in interoperability modes. This section briefly explains the basic concepts of these modes.

Each vendor has a regular mode and an equivalent interoperability mode, which specifically turns off advanced or proprietary features and provides the product with a more amiable standards-compliant implementation.

**Note**

For more information on configuring interoperability for the Cisco MDS 9000 Family switches, refer to the *Cisco MDS 9000 Family Switch-to-Switch Interoperability Configuration Guide*.

About Interop Mode

Cisco NX-OS software supports the following four interop modes:

- Mode 1— Standards based interop mode that requires all other vendors in the fabric to be in interop mode.
- Mode 2—Brocade native mode (Core PID 0).

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Mode 3—Brocade native mode (Core PID 1).
- Mode 4—McData native mode.

For information about configuring interop modes 2, 3, and 4, refer to the *Cisco MDS 9000 Family Switch-to-Switch Interoperability Configuration Guide*.

Table 1-2 lists the changes in switch behavior when you enable interoperability mode. These changes are specific to switches in the Cisco MDS 9000 Family while in interop mode.

Table 1-2 Changes in Switch Behavior When Interoperability Is Enabled

Switch Feature	Changes if Interoperability Is Enabled
Domain IDs	Some vendors cannot use the full range of 239 domains within a fabric. Domain IDs are restricted to the range 97-127. This is to accommodate McData's nominal restriction to this same range. They can either be set up statically (the Cisco MDS switch accept only one domain ID, if it does not get that domain ID it isolates itself from the fabric) or preferred. (If it does not get its requested domain ID, it accepts any assigned domain ID.)
Timers	All Fibre Channel timers must be the same on all switches as these values are exchanged by E ports when establishing an ISL. The timers are F_S_TOV, D_S_TOV, E_D_TOV, and R_A_TOV.
F_S_TOV	Verify that the Fabric Stability Time Out Value timers match exactly.
D_S_TOV	Verify that the Distributed Services Time Out Value timers match exactly.
E_D_TOV	Verify that the Error Detect Time Out Value timers match exactly.
R_A_TOV	Verify that the Resource Allocation Time Out Value timers match exactly.
Trunking	Trunking is not supported between two different vendor's switches. This feature may be disabled on a per port or per switch basis.
Default zone	The default zone behavior of permit (all nodes can see all other nodes) or deny (all nodes are isolated when not explicitly placed in a zone) may change.
Zoning attributes	Zones may be limited to the pWWN and other proprietary zoning methods (physical port number) may be eliminated. Note Brocade uses the cfsave command to save fabric-wide zoning configuration. This command does not have any effect on Cisco MDS 9000 Family switches if they are part of the same fabric. You must explicitly save the configuration on each switch in the Cisco MDS 9000 Family.
Zone propagation	Some vendors do not pass the full zone configuration to other switches, only the active zone set gets passed. Verify that the active zone set or zone configuration has correctly propagated to the other switches in the fabric.
VSAN	Interop mode only affects the specified VSAN. Note Interop modes cannot be enabled on FICON-enabled VSANs.
TE ports and PortChannels	TE ports and PortChannels cannot be used to connect Cisco MDS to non-Cisco MDS switches. Only E ports can be used to connect to non-Cisco MDS switches. TE ports and PortChannels can still be used to connect an Cisco MDS to other Cisco MDS switches even when in interop mode.

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

Table 1-2 Changes in Switch Behavior When Interoperability Is Enabled (continued)

Switch Feature	Changes if Interoperability Is Enabled
FSPF	The routing of frames within the fabric is not changed by the introduction of interop mode. The switch continues to use src-id, dst-id, and ox-id to load balance across multiple ISL links.
Domain reconfiguration disruptive	This is a switch-wide impacting event. Brocade and McData require the entire switch to be placed in offline mode and/or rebooted when changing domain IDs.
Domain reconfiguration nondisruptive	This event is limited to the affected VSAN. Only Cisco MDS 9000 Family switches have this capability—only the domain manager process for the affected VSAN is restarted and not the entire switch.
Name server	Verify that all vendors have the correct values in their respective name server database.
IVR	IVR-enabled VSANs can be configured in no interop (default) mode or in any of the interop modes.

Guidelines and Limitations

This section explains the database merge guidelines for this feature.

When merging two fabrics, follow these guidelines:

- Be aware of the following merge conditions:
 - The merge protocol is not implemented for distribution of the fctimer values—you must manually merge the fctimer values when a fabric is merged. The per-VSAN fctimer configuration is distributed in the physical fabric.
 - The fctimer configuration is only applied to those switches containing the VSAN with a modified fctimer value.
 - The global fctimer values are not distributed.
- Do not configure global timer values when distribution is enabled.



Note

The number of pending fctimer configuration operations cannot be more than 15. At that point, you must commit or abort the pending configurations before performing any more operations.

For information about CFS merge support, refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*.

Default Settings

Table 1-3 lists the default settings for the features included in this chapter.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Table 1-3 Default Settings for Advanced Features

Parameters	Default
CIM server	Disabled
CIM server security protocol	HTTP
D_S_TOV	5,000 milliseconds.
E_D_TOV	2,000 milliseconds.
R_A_TOV	10,000 milliseconds.
Timeout period to invoke fctrace	5 seconds.
Number of frame sent by the fcping feature	5 frames.
Remote capture connection protocol	TCP.
Remote capture connection mode	Passive.
Local capture frame limit s	10 frames.
FC ID allocation mode	Auto mode.
Loop monitoring	Disabled.
D_S_TOV	5,000 msec
E_D_TOV	2,000 msec
R_A_TOV	10,000 msec
Interop mode	Disabled

Configuring Timer Across All VSANs

You can modify Fibre Channel protocol related timer values for the switch.



Caution

The D_S_TOV, E_D_TOV, and R_A_TOV values cannot be globally changed unless all VSANs in the switch are suspended.



Note

If a VSAN is not specified when you change the timer value, the changed value is applied to all VSANs in the switch.

To configure timers in DCNM-SAN, expand **Switches > FC Services** and then select **Timers & Policies** in the Physical Attributes pane. You see the timers for multiple switches in the Information pane. Click the **Change Timeouts** button to configure the timeout values.

To configure timers in Device Manager, click **FC > Advanced > Timers/Policies**. You see the timers for a single switch in the dialog box.

This section includes the following topics:

- [Task Flow for Configuring Time Across All VSANs, page 1-9](#)
- [Configuring Timer Per-VSAN, page 1-9](#)
- [Enabling fctimer Distribution, page 1-10](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

- [Committing fctimer Changes, page 1-10](#)
- [Discarding fctimer Changes, page 1-10](#)
- [Configuring a Secondary MAC Address, page 1-10](#)
- [Configuring Interop Mode 1, page 1-11](#)

Task Flow for Configuring Time Across All VSANs

Follow these steps to configure time across all VSANs:

-
- | | |
|---------------|---|
| Step 1 | Configure the timer per-VSAN. |
| Step 2 | Enable the fctimer distribution. |
| Step 3 | Make the required configuration changes and commit the fctimer changes. |
| Step 4 | Discard the changes if you choose to discard the configuration changes. |
-

Configuring Timer Per-VSAN

You can also issue the fctimer for a specified VSAN to configure different TOV values for VSANs with special links like FC or IP tunnels. You can configure different E_D_TOV, R_A_TOV, and D_S_TOV values for individual VSANs. Active VSANs are suspended and activated when their timer values are changed.



Caution

You cannot perform a nondisruptive downgrade to any earlier version that does not support per-VSAN FC timers.



Note

This configuration must be propagated to all switches in the fabric—be sure to configure the same value in all switches in the fabric.

If a switch is downgraded to Cisco MDS SAN-OS Release 1.2 or 1.1 after the timer is configured for a VSAN, an error message is issued to warn against strict incompatibilities. Refer to the *Cisco MDS 9000 Family Troubleshooting Guide*.

Detailed Steps

To configure per-VSAN Fiber Channel timers using Device Manager, follow these steps:

-
- | | |
|---------------|---|
| Step 1 | Click FC > Advanced > VSAN Timers .
You see the VSANs Timer dialog box. |
| Step 2 | Fill in the timer values that you want to configure. |
| Step 3 | Click Apply to save these changes. |
-

Send documentation comments to dcnm-san-docfeedback@cisco.com

Enabling fctimer Distribution

Detailed Steps

To enable and distribute fctimer configuration changes using Device Manager, follow these steps:

-
- | | |
|---------------|---|
| Step 1 | Choose FC > Advanced > VSAN Timers .
You see the VSANs Timer dialog box. |
| Step 2 | Fill in the timer values that you want to configure. |
| Step 3 | Click Apply to save these changes. |
| Step 4 | Select commit from the CFS drop-down menu to distribute these changes or select abort from the CFS drop-down menu to discard any unsaved changes. |
-

Committing fctimer Changes

When you commit the fctimer configuration changes, the effective database is overwritten by the configuration changes in the pending database and all the switches in the fabric receive the same configuration. When you commit the fctimer configuration changes without implementing the session feature, the fctimer configurations are distributed to all the switches in the physical fabric.

Discarding fctimer Changes

After making the configuration changes, you can choose to discard the changes by discarding the changes instead of committing them. In either case, the lock is released.

Configuring a Secondary MAC Address

Detailed Steps

To allocate secondary MAC addresses using Device Manager, follow these steps:

-
- | | |
|---------------|---|
| Step 1 | Choose FC > Advanced > WWN Manager .
You see the list of allocated WWNs. |
| Step 2 | Supply the BaseMacAddress and MacAddressRange fields. |
| Step 3 | Click Apply to save these changes, or click Close to discard any unsaved changes. |
-

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuring Interop Mode 1

The interop mode1 in Cisco MDS 9000 Family switches can be enabled disruptively or nondisruptively.



Note

Brocade's **msplmgmtdeactivate** command must explicitly be run prior to connecting from a Brocade switch to either Cisco MDS 9000 Family switches or to McData switches. This command uses Brocade proprietary frames to exchange platform information, which Cisco MDS 9000 Family switches or McData switches do not understand. Rejecting these frames causes the common E ports to become isolated.

Detailed Steps

To configure interop mode 1 for a VSAN, follow these steps:

- Step 1** Choose **VSANxxx > VSAN Attributes** from the Logical Domains pane.
- Step 2** Select **Interop-1** from the Interop drop-down menu.
- Step 3** Click **Apply Changes** to save this interop mode.
- Step 4** Expand **VSANxxx** and then select **Domain Manager** from the Logical Domains pane.
You see the Domain Manager configuration in the Information pane.
- Step 5** Set the Domain ID in the range of 97 (0x61) through 127 (0x7F).
 - a. Click the **Configuration** tab.
 - b. Click in the Configure Domain ID column under the Configuration tab.
 - c. Click the **Running** tab and check that the change has been made.



Note

This is a limitation imposed by the McData switches.



Note

When changing the domain ID, the FC IDs assigned to N ports also change.

- Step 6** Change the Fibre Channel timers (if they have been changed from the system defaults).



Note

The Cisco MDS 9000, Brocade, and McData FC error detect (ED_TOV) and resource allocation (RA_TOV) timers default to the same values. They can be changed if needed. The RA_TOV default is 10 seconds, and the ED_TOV default is 2 seconds. Per the FC-SW2 standard, these values must be the same on each switch within the fabric.

- a. Expand **Switches > FC Services**, and then select **Timers and Policies**. You see the timer settings in the Information pane.
 - b. Click **Change Timeouts** to modify the time-out values.
 - c. Click **Apply** to save the new time-out values.
- Step 7** (Optional) Choose **VSANxxx > Domain Manager > Configuration** and select **disruptive** or **nonDisruptive** in the Restart column to restart the domain.

[Send documentation comments to dcnm-san-docfeedback@cisco.com](mailto:dcnm-san-docfeedback@cisco.com)

Verifying the Advanced Features and Concepts Configuration

For detailed information about the fields in the output from these commands, refer to the *Cisco MDS 9000 Family Command Reference*.

- [Verifying the Company ID Configuration, page 1-12](#)
- [Verifying Interoperating Status, page 1-12](#)
- [Displaying WWN Information, page 1-13](#)

Verifying the Company ID Configuration

To view the configured company IDs using Device Manager, choose **FC > Advanced > FcId Area Allocation**.

You can implicitly derive the default entries shipped with a specific release by combining the list of company IDs displayed without any identification with the list of deleted entries.

Some WWN formats do not support company IDs. In these cases, you may need to configure the FC ID persistent entry.

Verifying Interoperating Status

This section highlights the steps used to verify if the fabric is up and running in interoperability mode.



Note

The Cisco MDS name server shows both local and remote entries, and does not time out the entries.

To verify the interoperability status of any switch in the Cisco MDS 9000 Family using DCNM for SAN, follow these steps:

- Step 1** Choose **Switches** in the Physical Attributes pane and check the release number in the Information pane to verify the Cisco NX-OS release.
- Step 2** Expand **Switches > Interfaces**, and then select **FC Physical** to verify the interface modes for each switch.
- Step 3** Expand **Fabricxx** in the Logical Domains pane and then select **All VSANs** to verify the interop mode for all VSANs.
- Step 4** Expand **Fabricxx > All VSANs** and then select **Domain Manager** to verify the domain IDs, local, and principal sWWNs for all VSANs.
- Step 5** Using Device Manager, choose **FC > Name Server** to verify the name server information.
You see the Name Server dialog box.
- Step 6** Click **Close** to close the dialog box.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Displaying WWN Information

To display WWN information using Device Manager, choose **FC > Advanced > WWN Manager**. You see the list of allocated WWNs.

Additional References

For additional information related to implementing VSANs, see the following section:

- [Related Document, page 1-13](#)
- [Standards, page 1-13](#)
- [RFCs, page 1-13](#)
- [MIBs, page 1-13](#)

Related Document

Related Topic	Document Title
Cisco MDS 9000 Family Command Reference	<i>Cisco MDS 9000 Family Command Reference</i>

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

RFCs

RFC	Title
No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified.	—

MIBs

MIBs	MIBs Link
No new or modified MIBs are supported by this feature, and support for existing MIBs has not been modified.	To locate and download MIBs, go to the following URL: http://www.cisco.com/en/US/products/ps5989/prod_technical_reference_list.html

Send documentation comments to dcnm-san-docfeedback@cisco.com



INDEX

A

active zone sets

- considerations [1-13](#)

- enabling distribution [1-26](#)

aliases

- switching between global device aliases and FC aliases [1-8](#)

B

BB_credits

- FICON port swapping [1-19](#)

broadcast

- routing [1-6](#)

Brocade

- native interop mode [1-5](#)

C

CIM

- description [1-1](#)

Cisco SAN-OS features

- changed (table) [1-2](#)

- new (table) [1-2](#)

code pages

- FICON text string formatting [1-24](#)

company IDs

- FC ID allocations [1-4](#)

configuration

- saving automatically for FICON [1-13](#)

configuration files

- FICON [1-15](#)

Configuring FCoE Using DCNM for SAN [1-3](#)

Control Unit Port. See CUP in-band management

CUP in-band management

- blocking restriction [1-28](#)

- description [1-18](#)

- placing CUPs in zones [1-32](#)

D

dead time intervals

- configuring for FSPF [1-11](#)

- description [1-3](#)

default VSANs

- description [1-7](#)

default zones

- configuring access permissions [1-23](#)

- configuring policies [1-32](#)

- configuring QoS priorities [1-31](#)

- description [1-6](#)

- interoperability [1-6](#)

- policies [1-6](#)

destination IDs

- in-order delivery [1-6](#)

- path selection [1-9](#)

device alias database

- committing changes [1-7](#)

- discarding changes [1-8](#)

- distribution to fabric [1-5](#)

- overriding fabric lock [1-5](#)

device aliases

- comparison with zones (table) [1-4](#)

- creating [1-7](#)

- default settings [1-6](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

features [1-3](#)

modifying the database [1-4](#)

requirements [1-4](#)

domain IDs

configuring fcalias members [1-6](#)

interoperability [1-6](#)

DPVM

enabling [1-5](#)

requirements [1-2](#)

using DPVM Setup Wizard (procedure) [1-5](#)

DPVM databases

autolearned entries [1-3](#)

clearing [1-8](#)

comparing differences [1-11](#)

configuring CFS distribution [?? to 1-10](#)

copying [1-4](#)

description [1-2](#)

enabling autolearning [1-7](#)

drop latency time

configuring [1-14](#)

Dynamic Port VSAN Membership. See DPVM

E

EBCDIC

FICON string format [1-24](#)

enhanced zones

advantages over basic zones [1-11](#)

broadcast frames [1-37](#)

changing from basic zones [1-35](#)

creating attribute groups [1-36](#)

enabling [1-36](#)

merging databases [1-11](#)

modifying database [1-36](#)

E ports

FSPF topologies [1-1](#)

recovering from link isolations [1-8](#)

exchange IDs

in-order delivery [1-6](#)

path selection [1-9](#)

Extended Binary-Coded Decimal Interchange Code. See EBCDIC [1-24](#)

F

Fabric-Device Management Interface. See FDMI

Fabric Manager features

changed (table) [1-2](#)

new (table) [1-2](#)

fabric pWWNs

zone membership [1-2](#)

Fabric Shortest Path First. See FSPF

fabric WWNs. See fWWNs

fault tolerant fabrics

example (figure) [1-16](#)

fcaliases

adding members [1-24](#)

cloning [1-30](#)

configuring for zones [1-6](#)

creating [1-23](#)

renaming [1-30](#)

FC ID allocation

FICON implementation [1-12](#)

FC IDs

allocating [1-4](#)

allocating default company ID lists [1-4](#)

allocating for FICON [1-12](#)

allocation for HBAs [1-4](#)

configuring fcalias members [1-6](#)

FCIP

FICON support [1-5](#)

reserving ports for FICON [1-11](#)

FCIP interfaces

binding to FICON port numbers [1-28](#)

FCoE

configuring [1-1](#)

Device Manager [1-4](#)

enabling [1-2](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

FCP

intermixing protocols [1-5](#)

fctimers

distribution [1-3](#)

FDMI

description [1-2](#)

displaying database information [1-9](#)

Fibre Channel

time out values [1-2 to ??](#)

TOVs [1-8](#)

Fibre Channel interfaces

default settings [1-9, 1-7, 1-5, 1-15, 1-6, 1-8, 1-5, 1-20, 1-7](#)

Fibre Channel Protocol. See FCP

Fibre Connection. See FICON

FICON

advantages on MDS switches [1-3 to ??](#)

automatic configuration save [1-13](#)

calculating flow load balance (procedure) [1-33](#)

cascading [1-7](#)

configuration files [1-15 to 1-31](#)

configuring [?? to 1-27](#)

configuring ports [1-27 to 1-30](#)

CUP in-band management [1-18](#)

default settings [1-20](#)

description [?? to 1-8](#)

FC4 protocols [1-2](#)

FC ID allocations [1-12](#)

FCIP support [1-5](#)

host timestamp control [1-25](#)

implemented ports [1-10](#)

installed ports [1-11](#)

manually enabling [1-23](#)

MDS-supported features [1-5](#)

PortChannel support [1-5](#)

port numbering [?? to 1-12](#)

port swapping [1-16 to ??](#)

RLIRs [1-14 to 1-35](#)

saving configuration changes [1-13](#)

suspending a VSAN [1-23](#)

tape acceleration [1-16 to 1-31](#)

text string formatting codes [1-24](#)

unimplemented port [1-10](#)

VSAN offline state [1-23](#)

FICON configuration files

applying to running configuration [1-30](#)

copying [1-30](#)

displaying [1-35](#)

editing [1-30](#)

view latest information [1-30](#)

FICON port numbers

assigning to slots [1-21](#)

default numbering scheme [1-8](#)

FCIP interfaces [1-11](#)

implemented addresses [1-10](#)

installed ports [1-11](#)

logical interfaces [1-11](#)

numbering guidelines [1-19](#)

PortChannel interfaces [1-11](#)

port swapping [1-10](#)

reserved numbering scheme [1-11](#)

unimplemented addresses [1-10](#)

uninstalled ports [1-11](#)

FICON ports

assigning address names using Device Manager [1-29](#)

binding to FCIP interfaces [1-28](#)

binding to PortChannels [1-27](#)

blocking [1-28](#)

configuring prohibiting default state [1-28](#)

displaying address information [1-36](#)

prohibiting [1-14](#)

swapping configurations [1-31](#)

FICON port swapping

configuring (procedure) [1-31](#)

guidelines [1-19](#)

FICON tape acceleration

configuration considerations [1-20](#)

configuring [1-31](#)

description [1-16](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

FICON tape read acceleration

 configuring [1-32](#)

FLOGI

 displaying details [1-8](#)

flow statistics

 description [1-8](#)

FL ports

 DPVM support [1-3](#)

F ports

 DPVM support [1-3](#)

FSPF

 computing link cost [1-11](#)

 configuring globally [1-2 to ??](#)

 configuring Hello time intervals [1-3](#)

 configuring link cost [1-3](#)

 dead time intervals [1-3](#)

 default settings [1-8](#)

 description [1-1](#)

 disabling [1-10](#)

 disabling on interfaces [1-12](#)

 disabling routing protocols [1-10](#)

 displaying databases [1-15](#)

 enabling [1-10](#)

 fail-over with PortChannels [1-16](#)

 fault tolerant fabrics [1-16](#)

 in-order delivery [1-6 to ??](#)

 interoperability [1-7](#)

 link state record defaults [1-3](#)

 multicast root switches [1-6, 1-13](#)

 reconvergence times [1-2](#)

 redundant links [1-16](#)

 resetting to defaults [1-10](#)

 retransmitting intervals [1-4](#)

 routing services [1-1](#)

 topology examples [?? to 1-17](#)

FSPF multicast roots

 configuring switches [1-13](#)

FSPF routes

 configuring [1-12](#)

 description [1-4](#)

FSPF routing

 multicast [1-6](#)

full zone sets

 considerations [1-13](#)

 enabling distribution [1-26](#)

fWWNs

 configuring fcalias members [1-6](#)

Fx ports

 VSAN membership [1-5](#)

H

hard zoning

 description [1-7](#)

HBAs

 device aliases [1-1](#)

 FC ID allocations [1-4](#)

Hello time intervals

 configuring for FSPF [1-11](#)

 description [1-3](#)

host control

 FICON [1-25](#)

IBM PPRC

 FICON support [1-5](#)

in-band management

 CUP [1-18](#)

in-order delivery

 configuring drop latency time [1-14](#)

 enabling for VSANs [1-14](#)

 enabling globally [1-13](#)

 guidelines [1-8](#)

 reordering network frames [1-6](#)

 reordering PortChannel frames [1-7](#)

interfaces

Send documentation comments to dcnm-san-docfeedback@cisco.com

- assigning to VSANs [1-11](#)
- configuring fcalias members [1-7](#)
- default settings [1-9, 1-7, 1-5, 1-15, 1-6, 1-8, 1-5, 1-20, 1-7](#)
- VSAN membership [1-7](#)

interoperability

- configuring interop mode 1 [1-11](#)
- description [1-5](#)
- verifying status [1-12](#)
- VSANs [1-9](#)

interop modes

- configuring mode 1 [1-11](#)
- default settings [1-7](#)
- description [1-5](#)

IOD. See in-order delivery

IPv4 addresses

- configuring fcalias members [1-7](#)

IPv6 addresses

- configuring fcalias members [1-2, 1-7](#)

isolated VSANs

- description [1-7](#)
- displaying membership [1-14](#)

IVR

- SDV limitations [1-6](#)

L

link costs

- configuring for FSPF [1-11](#)
- description [1-3](#)

Link Incident Records. See LIRs

LIRs

- description [1-14](#)

load balancing

- attributes [1-9](#)
- attributes for VSANs [1-6](#)
- configuring [1-12](#)
- description [1-9](#)

logical unit numbers. See LUNs

LUN zoning

- configuring [1-34](#)

- description [1-10](#)

M

MAC addresses

- configuring secondary [1-10](#)

mainframes

- FICON parameters [1-25](#)
- VSAN clock [1-26](#)

management interfaces

- default settings [1-9, 1-7, 1-5, 1-15, 1-6, 1-8, 1-5, 1-20, 1-7](#)

McData

- native interop mode [1-6](#)

mgmt0 interfaces

- default settings [1-9, 1-7, 1-5, 1-15, 1-6, 1-8, 1-5, 1-20, 1-7](#)

multicast root switches

- configuring [1-13](#)
- description [1-6](#)

N

name servers

- displaying database entries [1-2](#)
- interoperability [1-7](#)
- LUN information [1-1](#)
- proxy feature [1-2](#)
- registering proxies [1-2](#)

NL ports

- hard zoning [1-7](#)
- zone enforcement [1-7](#)

node world wide names. See nWWNs

N ports

- hard zoning [1-7](#)
- zone enforcement [1-7](#)
- zone membership [1-2](#)
- See also Nx ports

nWWNs

Send documentation comments to dcnm-san-docfeedback@cisco.com

DPVM [1-1](#)

P

PLOGI

name server [1-2](#)

port addresses

FICON [1-10](#)

PortChannels

binding to FICON port numbers [1-27](#)

FICON support [1-5](#)

in-order guarantee [1-7](#)

interoperability [1-6](#)

link changes [1-7](#)

link failures [1-16](#)

reserving ports for FICON [1-11](#)

port numbers. See FICON port numbers

ports

VSAN membership [1-7](#)

port swapping. See FICON port swapping

port world wide names. See pWWNs

proxies

registering for name servers [1-2](#)

pWWNs

configuring fcalias members [1-6](#)

DPVM [1-1](#)

zone membership [1-2](#)

R

read-only zones

configuration guidelines [1-14](#)

configuring [1-34](#)

description [1-10](#)

redundancy

VSANs [1-5](#)

redundant physical links

example (figure) [1-16](#)

Registered Link Incident Reports. See RLIRs

Registered State Change Notifications. See RSCNs

retransmitting intervals

configuring for FSPF [1-12](#)

description [1-4](#)

RLIRs

description [1-14](#)

displaying information (procedure) [1-35](#)

specifying preferred host [1-29](#)

route costs

computing [1-3](#)

routing

multicast [1-6](#)

See also broadcast routing

See also IP routing

RSCNs

default settings [1-5](#)

description [1-3](#)

displaying information [1-9](#)

multiple port IDs [1-3](#)

suppressing domain format SW-RSCNs [1-6](#)

RSCN timers

configuration distribution using CFS [1-4 to ??](#)

configuring [1-7](#)

runtime checks

static routes [1-5](#)

S

scalability

VSANs [1-5](#)

SCSI

displaying LUN discovery results [1-3](#)

SCSI LUNs

customized discovery [1-2](#)

discovering targets [1-1](#)

displaying information [1-3](#)

starting discoveries [1-2](#)

SDV

Send documentation comments to dcnm-san-docfeedback@cisco.com

IVR limitations [1-6](#)
 secondary MAC addresses
 configuring [1-10](#)
 small computer system interface. See SCSI
 SNMP
 FICON control [1-26](#)
 soft zoning
 description [1-7](#)
 See also zoning
 source IDs
 in-order delivery [1-6](#)
 path selection [1-9](#)
 SPF
 computational hold times [1-3](#)
 static routes
 runtime checks [1-5](#)

T

tape acceleration
 FICON [1-16 to 1-31](#)
 TE ports
 FSPF topologies [1-1](#)
 interoperability [1-6](#)
 recovering from link isolations [1-8](#)
 time out values. See TOVs
 timestamps
 FICON host control [1-25](#)
 TOVs
 configuring across all VSANs [1-8](#)
 configuring for a VSAN [1-9](#)
 default settings [1-7](#)
 interoperability [1-6](#)
 ranges [1-2](#)
 traffic isolation
 VSANs [1-5](#)
 trunking
 interoperability [1-6](#)
 trunking ports

associated with VSANs [1-7](#)

V

virtual interface
 limitations [1-1](#)
 VSAN IDs
 description [1-6](#)
 range [1-5](#)
 VSAN membership [1-5](#)
 VSANs
 advantages [1-4](#)
 broadcast addresses [1-6](#)
 clocks [1-25](#)
 comparison with zones (table) [1-5](#)
 configuring FSPF [1-2](#)
 default settings [1-9](#)
 default VSANs [1-7](#)
 deleting [1-8](#)
 description [?? to 1-5](#)
 fabric optimization for FICON [1-3](#)
 FC IDs [1-1](#)
 features [1-1](#)
 FICON-enabled [1-9](#)
 FSPF connectivity [1-2](#)
 interop mode [1-6](#)
 isolated [1-7](#)
 load balancing [1-9](#)
 load balancing attributes [1-6](#)
 multiple zones [1-13](#)
 names [1-6](#)
 name server [1-2](#)
 operational states [1-8](#)
 port membership [1-7](#)
 states [1-6](#)
 suspending for FICON [1-23](#)
 timer configuration [1-9](#)
 traffic isolation [1-4](#)
 trunking ports [1-7](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

W

wizards

Quick Config Wizard [1-18](#)

world wide names. See WWNs

WWNs

configuring [1-3](#)

displaying information [1-13](#)

link initialization [1-4](#)

secondary MAC addresses [1-10](#)

X

XRC

FICON support [1-5](#)

Z

zone attribute groups

cloning [1-30](#)

zone databases

migrating a non-MDS database [1-31](#)

zone members

adding to zones [1-17](#)

converting to pWWN members [1-24](#)

displaying information [1-23](#)

zones

access control [1-5](#)

adding to zone sets [1-25](#)

adding zone members [1-17](#)

assigning LUNs to storage subsystems [1-34](#)

backing up (procedure) [1-9](#)

changing from enhanced zones [1-35](#)

cloning [1-30](#)

compacting for downgrading [1-38](#)

comparison with device aliases (table) [1-4](#)

comparison with VSANs (table) [1-5](#)

configuring [?? to 1-25](#)

configuring aliases [1-6](#)

configuring broadcasting [1-9](#)

configuring fcaliases [1-6](#)

CUPs [1-32](#)

default policies [1-2](#)

editing full zone databases [1-4](#)

enforcing restrictions [1-7](#)

exporting databases [1-8](#)

features [1-3](#)

importing databases [1-8](#)

LUN-based [1-10](#)

membership using pWWNs [1-5](#)

renaming [1-30](#)

restoring (procedure) [1-9](#)

viewing information [1-38](#)

See also default zones

See also hard zoning; soft zoning

See also LUN zoning

See also read-only zones

See also zoning; zone sets

zone server databases

clearing [1-31](#)

zone sets

activating [1-21](#)

adding member zones [1-25](#)

cloning [1-30](#)

configuring [1-5 to 1-23](#)

considerations [1-13](#)

copying [1-8](#)

creating [1-5, 1-25](#)

distributing configuration [1-7](#)

enabling distribution [1-26](#)

exporting [1-27](#)

exporting databases [1-8](#)

importing [1-27](#)

importing databases [1-8](#)

one-time distribution [1-27](#)

recovering from link isolations [1-8](#)

renaming [1-30](#)

viewing information [1-38](#)

Send documentation comments to dcnm-san-docfeedback@cisco.com

See also active zone sets

See also active zone sets; full zone sets

See also zones; zoning

zone traffic priorities

description [1-9](#)

zoning

configuring broadcasting [1-9](#)

implementation [1-3](#)

Quick Config Wizard [1-18 to 1-20](#)

See also LUN zoning

Send documentation comments to dcnm-san-docfeedback@cisco.com