

Cisco MDS 9000 Series Fabric Configuration Guide, Release 9.x

First Published: 2022-09-02

Last Modified: 2026-03-04

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED “AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: https://www.cisco.com/c/en/us/about/legal/trademarks.html. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2022–2026 Cisco Systems, Inc. All rights reserved.



CONTENTS

Full Cisco Trademarks with Software License ?

PREFACE

Preface xvii

Audience xvii

Document Conventions xvii

Documentation Feedback xix

Related Documentation xix

Communications, Services, and Additional Information xix

CHAPTER 1

New and Changed Information 1

Change Summary 1

CHAPTER 2

Fabric Overview 5

Virtual SANs 5

Dynamic Port VSAN Membership 6

SAN Device Virtualization 6

Zoning 6

Distributed Device Alias Services 7

Fibre Channel Routing Services and Protocols 7

Multiprotocol Support 7

CHAPTER 3

Configuring and Managing VSANs 9

About VSANs 9

VSANs Topologies 10

VSAN Advantages 11

VSANs Versus Zones 12

VSAN Configuration 13
Reserved VSAN Range and Isolated VSAN Range Guidelines 14
VSAN Creation 14
Creating VSANs Statically 14
Creating VSANs 14
Port VSAN Membership 15
Assigning Static Port VSAN Membership 15
Displaying VSAN Static Membership 16
Default VSAN 17
Isolated VSAN 17
Displaying Isolated VSAN Membership 18
Operational State of a VSAN 18
Static VSAN Deletion 18
Deleting Static VSANs 19
Load Balancing 19
Configuring Load Balancing 19
Interop Mode 20
FICON VSANs 20
Displaying Static VSAN Configuration 21
Default Settings 21
Displaying Fabric Switch Information 22

CHAPTER 4 **Creating Dynamic VSANs 23**
About DPVM 23
About DPVM Configuration 24
Enabling DPVM 24
DPVM Device Configuration (Static) 25
Configuring DPVM 25
Activating DPVM 26
DPVM Autolearn 27
Enabling Autolearn 27
Clearing Learned Entries 28
Disabling Autolearn 28
DPVM Distribution 29

About DPVM Distribution 29
Disabling DPVM Distribution 29
About Locking the Fabric 30
Locking the Fabric 30
Committing Changes 31
Discarding Changes 31
Clearing a Locked Session 31
DPVM Configuration Merge Guidelines 32
About Copying DPVM DPVM Configurations 32
Copying DPVM Active Configuration 32
Comparing Database Differences 33
Displaying DPVM Merge Status and Statistics 33
Displaying DPVM Configurations 34
Sample DPVM Configuration 35
Default Settings 38

CHAPTER 5 **Configuring and Managing Zones** 39
Finding Feature Information 39
Feature History for Configuring and Managing Zones 40
About Zoning 40
Zoning Example 42
Zone Implementation 42
Zone Member Configuration Guidelines 43
Active and Full Zoneset Considerations 43
Autozone 45
Guidelines and Limitations for Autozone 46
Configuring Autozone in Automatic Mode 47
Enabling Autozone in Automatic Mode 47
Enabling Autosave 48
Executing Autozone in Manual mode 48
Enabling Autozone in Automatic Mode by a Remotely Authenticated (AAA) User 48
Disabling Autosave 49
Disabling Autozone Automatic Mode 49
Displaying All the Zone Configurations 49

Displaying Pending Zone Configurations 49
Applying Pending Zone Configurations (Manual Mode) 49
Deleting Zones and Zoneset Created by Autozone 50
Example: Configuring Autozone 50
Verifying Autozone Configuration 51
Autozone Example Scenario 53
Zone Configuration 56
Configuring a Zone 56
Filtering End Devices Based on Name, WWN or FC ID 59
Adding Multiple End Devices to Multiple Zones 59
Zone Sets and FC Aliases 60
ZoneSet Creation 60
Activating a Zoneset 61
Overwrite Control for an Active Zoneset 62
Default Zone 63
Configuring the Default Zone Access Permission 64
About FC Alias Creation 64
Creating FC Aliases 65
Creating Zone Sets and Adding Member Zones 67
Adding Multiple Zones to Multiple Zone Sets 68
Zone Enforcement 69
ZoneSet Distribution 69
Enabling Full Zoneset Distribution 70
Enabling a One-Time Distribution 70
About Recovering from Link Isolation 71
Importing and Exporting Zone Sets 71
Zoneset Duplication 72
Copying Zone Sets 72
Renaming Zones, Zone Sets, and Aliases 73
Cloning Zones, Zone Sets, FC Aliases, and Zone Attribute Groups 74
Clearing the Zone Server Database 74
Advanced Zone Attributes 75
About Zone-Based Traffic Priority 75
Configuring Zone-Based Traffic Priority 75

Disabling zone-based traffic priority 77
Configuring Default Zone QoS Priority Attributes 77
About Smart Zoning 78
Smart Zoning Member Configuration 79
Enabling Smart Zoning on a VSAN 79
Setting Default Value for Smart Zoning 80
Converting Zones Automatically to Smart Zoning 80
Configuring Device Types for Zone Members 81
Removing Smart Zoning Configuration 82
Disabling Smart Zoning at Zone Level in the Basic Zoning Mode 82
Disabling Smart Zoning at Zone Level for a VSAN in the Enhanced Zoning Mode 83
Displaying Zone Information 84
Enhanced Zoning 92
About Enhanced Zoning 92
Changing from Basic Zoning to Enhanced Zoning 93
Changing from Enhanced Zoning to Basic Zoning 93
Enabling Enhanced Zoning 94
Modifying the Zone Database 94
Enabling Automatic Zone Pending Diff Display 95
Releasing Zone Database Locks 95
Creating Attribute Groups 96
Merging the Database 97
Merge Process 97
Configuring Zone Merge Control Policies 106
Preventing Zones From Flooding FC2 Buffers 107
Permitting or Denying Traffic in the Default Zone 107
Broadcasting a Zone 107
Configuring System Default Zoning Settings 109
Configuring Zone Generic Service Permission Settings 109
Displaying Enhanced Zone Information 110
Controlling Zoning Configuration Sessions 112
Configuring Zoning Session Limit 113
Compacting the Zone Database for Downgrading 113
Zone and ZoneSet Analysis 114

Zoning Best Practice 117
TCAM Regions 117
Zoning Types 118
Best Practises for Forwarding Engines 121
F, TF, NP, and TNP Port Channels 126
Best Practises for E and TE Port Channels and IVR 127
Enhancing Zone Server Performance 130
Zone Server-Fibre Channel Name Server Shared Database 130
Enabling the Zone Server-FCNS Shared Database 130
Disabling Zone Server-FCNS shared database 130
Zone Server SNMP Optimization 131
Enabling Zone Server SNMP Optimization 131
Disabling Zone Server SNMP Optimization 132
Zone Server Delta Distribution 132
Enabling Zone Server Delta Distribution 133
Disabling Zone Server Delta Distribution 133
Default Settings 134

CHAPTER 6 **Distributing Device Alias Services 135**
Understanding Device Aliases 135
Device Alias Modes 135
Guidelines and Limitations 136
Changing Mode Settings 137
Device Alias Mode Distribution 137
Device Alias Diffs-Only Distribution 137
Configuring Device Alias Diffs-Only Distribution 137
Merging Device Alias with the Diffs-Only Distribution Feature Enabled 139
Merging Device Alias in Different Modes 139
Resolving Merge Failure and Device Alias Mode Mismatch 140
Device Alias Features 140
Device Alias Requirements 141
Zone Aliases Versus Device Aliases 141
Device Alias Databases 142
Creating Device Aliases 142

About Device Alias Distribution 143
About Creating a Device Alias 143
About Device Alias Configuration Best Practices 143
Committing Changes 145
Enabling the Device Alias Pending Diff Display 145
Discarding Changes 146
Fabric Lock Override 147
Clearing Database Content 147
Clearing Statistics 147
Disabling and Enabling Device Alias Distribution 147
About Legacy Zone Alias Configuration Conversion 148
Importing a Zone Alias 149
Device Alias Statistics Cleanup 150
Database Merge Guidelines 150
Device Alias Configuration Verification 150
Default Settings 152
Resolving Device Alias Merge Failures 153
Device Alias Best Practices 153
Resolving Device Alias Mismatches 155
Resolving Merge Failures 156
Resolving Duplicate Device Alias Names (Same Device Alias Name, Different pWWNs) 156
Resolving Duplicate pWWNs (Different Device Alias Names, Same pWWN) 158
Resolving Mode Mismatch 160
Resolving a Validation Failure 162
Resolving Database Conflicts 164
Verifying the Device-Alias Database Status 165

CHAPTER 7 **Configuring Fibre Channel Routing Services and Protocols 167**
About FSPF 167
FSPF Examples 168
Fault Tolerant Fabric 168
Redundant Links 168
Failover Scenarios for PortChannels and FSPF Links 169
FSPF Global Configuration 170

About SPF Computational Hold Times 170
About Link State Record Defaults 170
Configuring FSPF on a VSAN 171
Resetting FSPF to the Default Configuration 171
Enabling or Disabling FSPF 172
Clearing FSPF Counters for the VSAN 172
FSPF Interface Configuration 172
About FSPF Link Cost 172
Configuring FSPF Link Cost 173
About FSPF Cost Multiplier 173
Setting up FSPF Cost Multiplier 173
Displaying FSPF Cost Multiplier 174
About Hello Time Intervals 175
Configuring Hello Time Intervals 175
About Dead Time Intervals 176
Configuring Dead Time Intervals 176
About Retransmitting Intervals 177
Configuring Retransmitting Intervals 177
About Disabling FSPF for Specific Interfaces 177
Disabling FSPF for Specific Interfaces 178
Clearing FSPF Counters for an Interface 178
FSPF Routes 179
About Fibre Channel Routes 179
About Broadcast and Multicast Routing 179
About Multicast Root Switch 180
Setting the Multicast Root Switch 180
Load Balancing 180
Load Balancing Schemes 181
Hashing Methods 182
Configure VSAN hash-method 185
Display VSAN load balancing hash-type 185
In-Order Delivery 186
About Reordering Network Frames 186
About Reordering PortChannel Frames 187

About Enabling In-Order Delivery 187
Enabling In-Order Delivery Globally 188
Enabling In-Order Delivery for a VSAN 188
Displaying the In-Order Delivery Status 189
Configuring the Drop Latency Time 189
Displaying Latency Information 190
Flow Statistics Configuration 190
About Flow Statistics 190
Counting Aggregated Flow Statistics 191
Counting Individual Flow Statistics 191
Clearing FIB Statistics 192
Displaying Flow Statistics 192
Displaying Global FSPF Information 193
Displaying the FSPF Database 193
Displaying FSPF Interfaces 195
Default Settings 196

CHAPTER 8 **Managing FLOGI, Name Server, FDMI, and RSCN Databases** 199
About FLOGI 199
FLOGI Scale Optimization 199
FLOGI Quiesce Timeout 200
Restrictions 200
Enabling FLOGI Scale Optimization and Quiesce Timeout 200
Disabling FLOGI Scale Optimization and Quiesce Timeout 201
Displaying FLOGI Details 202
Name Server 203
Bulk Notification Sent from the Name Server 204
Enabling Name Server Bulk Notification 204
Disabling Name Server Bulk Notification 205
Disabling Name Server Bulk Notification for NX-OS Release 6.2(9) 205
Re-enabling Name Server Bulk Notification 205
Name Server Proxy Registration 206
Registering Name Server Proxies 206
About Rejecting Duplicate pWWN 206

Rejecting Duplicate pWWNs 206
Name Server Database Entries 207
Optimizing Name Server Database Sync 207
Verifying the Number of Name Server Database Entries 208
Displaying Name Server Database Entries 208
FDMI 209
Displaying FDMI 210
VMID 213
Guidelines and Limitations for VMID 215
Configuring the VMID Server 216
Enabling the VMID Server 216
Disabling the VMID Server 216
Configuring a VMID Range 217
Examples: Configuring the VMID Server 217
Verifying a VMID Configuration 217
RSCN 221
About RSCN Information 221
Displaying RSCN Information 222
multi-pid Option 223
Configuring the multi-pid Option 223
Suppressing Domain Format SW-RSCNs 223
Coalesced SW-RSCN 224
Enabling Coalesced SW-RSCNs 224
Disabling Coalesced SW-RSCNs 225
Clearing RSCN Statistics 225
RSCN Timer Configuration Distribution Using CFS 226
Configuring the RSCN Timer 226
Verifying the RSCN Timer Configuration 227
RSCN Timer Configuration Distribution 227
Enabling RSCN Timer Configuration Distribution 228
Locking the Fabric 229
Committing the RSCN Timer Configuration Changes 229
Discarding the RSCN Timer Configuration Changes 229
Clearing a Locked Session 230

Displaying RSCN Configuration Distribution Information 230
Default Settings 231
Enabling Port Pacing 231

CHAPTER 9 **Discovering SCSI Targets** 233
About SCSI LUN Discovery 233
About Starting SCSI LUN Discovery 233
Starting SCSI LUN Discovery 233
About Initiating Customized Discovery 234
Initiating Customized Discovery 234
Displaying SCSI LUN Information 235

CHAPTER 10 **Configuring FICON** 239
About FICON 239
FICON Requirements 240
MDS-Specific FICON Advantages 240
Fabric Optimization with VSANs 240
FCIP Support 242
PortChannel Support 242
VSANs for FICON and FCP Mixing 242
Cisco MDS 9000-Supported FICON Features 242
FICON Cascading 244
FICON VSAN Prerequisites 245
FICON Port Numbering 245
Default FICON Port Numbering Scheme 245
Port Addresses 247
Implemented and Unimplemented Port Addresses 248
About the Reserved FICON Port Numbering Scheme 248
Installed and Uninstalled Ports 248
FICON Port Numbering Guidelines 249
Assigning FICON Port Numbers to Slots 249
Displaying the FICON Port Number Assignments 250
About Port Numbers for FCIP and PortChannel 250
Reserving FICON Port Numbers for FCIP and PortChannel Interfaces 251

FC ID Allocation 252
Configuring FICON 252
About Enabling FICON on a VSAN 252
Enabling FICON on the Switch 253
Setting Up a Basic FICON Configuration 254
Manually Enabling FICON on a VSAN 257
Configuring the code-page Option 258
Allowing the Host to Move the Switch Offline 259
Allowing the Host to Change FICON Port Parameters 259
Allowing the Host to Control the Timestamp 260
Clearing the Time Stamp 261
Configuring SNMP Control of FICON Parameters 261
About FICON Device Allegiance 261
Clearing FICON Device Allegiance 262
Automatically Saving the Running Configuration 262
Configuring FICON Ports 263
Binding Port Numbers to PortChannels 264
Binding Port Numbers to FCIP Interfaces 264
Configuring Port Blocking 265
Port Prohibiting 265
Assigning a Port Address Name 265
About RLIR 266
Specifying an RLIR Preferred Host 266
Displaying RLIR Information 267
Clearing RLIR Information 270
FICON Configuration Files 270
About FICON Configuration Files 271
Applying the Saved Configuration Files to the Running Configuration 271
Editing FICON Configuration Files 272
Displaying FICON Configuration Files 273
Copying FICON Configuration Files 273
Port Swapping 274
FICON Tape Acceleration 274
Configuring FICON Tape Acceleration 275

Configuring FICON Tape Read Acceleration 276
Configuring Zoning in a FICON VSAN 277
Moving a FICON VSAN to an Offline State 278
CUP In-Band Management 278
Displaying Control Unit Information 278
Displaying FICON Information 279
Receiving FICON Alerts 279
Displaying FICON Port Address Information 279
Displaying the Configured FICON State 281
Displaying Buffer Information 281
Viewing the History Buffer 282
Displaying FICON Information in the Running Configuration 282
Displaying FICON Information in the Startup Configuration 283
Displaying FICON-Related Log Information 284
Default Settings 284

CHAPTER 11 **Advanced Features and Concepts** 287
Common Information Model 287
Fibre Channel Time-Out Values 287
Timer Configuration Across All VSANs 288
Timer Configuration Per-VSAN 288
About fctimer Distribution 289
Enabling fctimer Distribution 290
Committing fctimer Changes 290
Discarding fctimer Changes 290
Fabric Lock Override 291
Database Merge Guidelines 291
Displaying Configured fctimer Values 292
Organizationally Unique Identifiers 292
Guidelines and Limitations 292
Adding and Deleting OUIs 293
Configuration Examples for Adding and Deleting OUIs 293
Example: Adding and Deleting OUIs 293
Example: Displaying OUIs 293

World Wide Names 293
Displaying WWN Information 294
Link Initialization WWN Usage 294
Configuring a Secondary MAC Address 295
FC ID Allocation for HBAs 295
Default Company ID List 296
Verifying the Company ID Configuration 297
Default Settings 298

CHAPTER 12 **Configuring Fibre Channel Common Transport Management Security** 299
About Fibre Channel Common Transport 299
Configuration Guidelines 299
Configuring the Fibre Channel Common Transport Query 300
Verifying Fibre Channel Common Transport Management Security 300
Default Settings 301



Preface

This preface describes the audience, organization of, and conventions used in the Cisco MDS 9000 Series Configuration Guides. It also provides information on how to obtain related documentation, and contains the following sections:

- Audience, on page xvii
- Document Conventions, on page xvii
- Documentation Feedback, on page xix
- Related Documentation, on page xix
- Communications, Services, and Additional Information, on page xix

Audience

This publication is for network administrators who install, configure, and maintain Cisco MDS 9000 Series Switches.

Document Conventions

Command descriptions use these conventions:

Convention	Description
bold	Bold text indicates the commands and keywords that you enter literally as shown.
Italic	Italic text indicates arguments for which the user supplies the values.
[x]	Square brackets enclose an optional element (keyword or argument).
[x | y]	Square brackets enclosing keywords or arguments separated by a vertical bar indicate an optional choice.
{x | y}	Braces enclosing keywords or arguments separated by a vertical bar indicate a required choice.

Convention	Description
[x {y | z}]	Nested set of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element.
`variable`	Indicates a variable for which you supply values, in context where italics cannot be used.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.

Examples use these conventions:

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



Warning IMPORTANT SAFETY INSTRUCTIONS

This warning symbol means danger. You are in a situation that could cause bodily injury. Before you work on any equipment, be aware of the hazards involved with electrical circuitry and be familiar with standard practices for preventing accidents. Use the statement number provided at the end of each warning to locate its translation in the translated safety warnings that accompanied this device.

SAVE THESE INSTRUCTIONS

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to mds-docfeedback@cisco.com. We appreciate your feedback.

Related Documentation

The entire Cisco MDS 9000 Series switches documentation set is available at the following URL:

https://www.cisco.com/c/en/us/support/storage-networking/mds-9000-nx-os-san-os-software/series.html

Documentation Roadmap

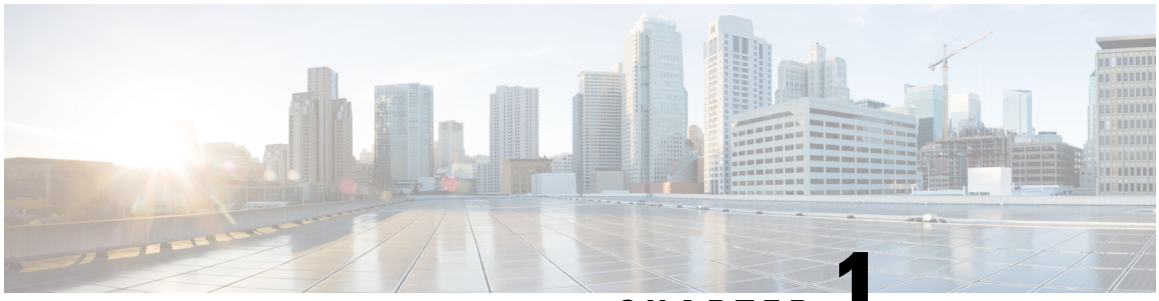
https://www.cisco.com/c/en/us/td/docs/storage/san_switches/mds9000/roadmaps/rel90.html

Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at Cisco Profile Manager.
- To get the business results you're looking for with the technologies that matter, visit Cisco Services.
- To submit a service request, visit Cisco Support.
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit Cisco Marketplace.
- To obtain general networking, training, and certification titles, visit Cisco Press.
- To find warranty information for a specific product or product family, access Cisco Warranty Finder.

Cisco Bug Search Tool

Cisco Bug Search Tool (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.



CHAPTER 1

New and Changed Information

- Change Summary, on page 1

Change Summary

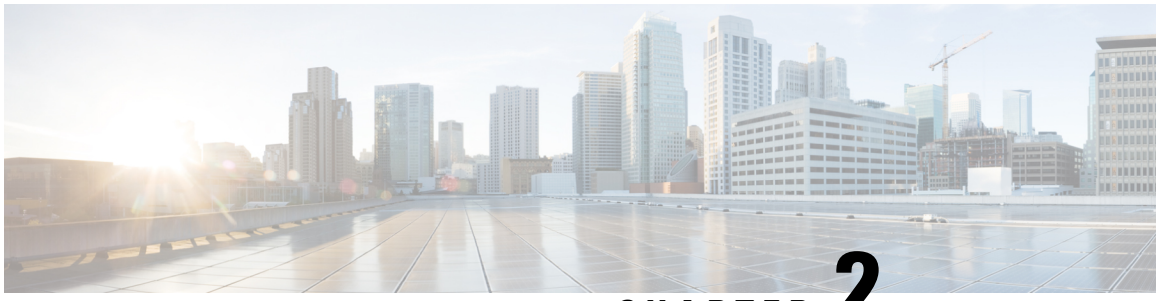
This chapter lists the New and Changed features for this guide.

Table 1: New and Changed Features

Feature	New and Change Topics	Changed in Release	Where Documented
There are no new features in Cisco MDS NX-OS Release 9.4(5).	There are no new features in Cisco MDS NX-OS Release 9.4(5).	9.4(5)	There are no new features in Cisco MDS NX-OS Release 9.4(5).
FSPF Database	The Link Type for the command **show fspf database vsan** is changed from numerical to alphanumerical.	9.4(1)	Displaying the FSPF Database , on page 193
Device Alias	The default device alias mode is changed to enhanced mode.	8.5(1)	Distributing Device Alias Services
Zoning Single Session	A new option to allow only a single configuration session at a time on a switch when in enhanced zoning mode.	8.4(2)	Configuring and Managing Zones
Autozone	Autozone provides options to enable or disable automatically saving of the running configuration to the startup-configuration after making a zoning change, using the **enableautosave** and **disableautosave** options respectively.	8.4(1)	Configuring and Managing Zones
Autozone	The Autozone feature is a mechanism to automate zoning via a single command or to allow minimal effort for deploying of fabric switches that support 32 Gbps or higher speed Fibre Channel interfaces.	8.3(1)	Configuring and Managing Zones
FLOGI Scale Optimization	Changed the default FLOGI quiesce timeout value from 2000 to 0 millisecond.	8.3(1)	Managing FLOGI, Name Server, FDMI, and RSCN Databases

Feature	New and Change Topics	Changed in Release	Where Documented
Virtual Machine Identifier (VMID)	VMID support allows unique identification of virtual machines (VMs) by SAN fabric infrastructure.	8.2(1)	Managing FLOGI, Name Server, FDMI, and RSCN Databases
FLOGI Scale Optimization	The FLOGI scale optimization feature enables users to increase the chassis-wide FLOGI scale limits. This feature is supported only on Cisco MDS 9718 Directors.	8.1(1)	FLOGI Scale Optimization, on page 199
Zone Server Enhancements	The Zone Server performance is enhanced using the following features: • Zone Server FCNS Shared Database • Zone Server SNMP Optimizations • Zone Server Delta Distribution	7.3(0)D1(1)	Configuring and Managing Zones
Device Alias Diffs-only Distribution	When this feature is enabled on all the switches in a fabric, it enables better scalability.	7.3(0)D1(1)	Distributing Device Alias Services
Organizationally Unique Identifiers	This feature introduces a new command which enables dynamic addition of Organizationally Unique Identifiers (OUIs) to the system OUI database.	7.3(0)D1(1)	Organizationally Unique Identifiers, on page 292
Confirm commit device-alias Confirm commit zone	Added pending-diff display on commit for zone and device-alias.	6.2(9)	Distributing Device Alias Services Configuring and Managing Zones
FC and FCOE Scale – Device Alias	Added “Device Alias Configuration Best Practices” section.	6.2(9)	Distributing Device Alias Services
Fibre Channel Common Transport Management Server Query	Configuring Fibre Channel Common Transport Management Server Query	6.2(9)	Configuring Fibre Channel Common Transport Management Security, on page 299
FCNS, RSCN	Added bulk notification feature to improve the performance of all the components listening to FCNS database changes. Added coalesced SWRSCN to improve RSCN performance.	6.2(7)	Managing FLOGI, Name Server, FDMI, and RSCN Databases
	Added “Displaying Fabric Switch Information” section.	6.2(7)	Configuring and Managing Zones
Smart Zoning	Added the command output.	6.2(7)	Configuring and Managing Zones
Smart Zoning	Added the Smart Zoning section.	5.2.6	Configuring and Managing Zones

Feature	New and Change Topics	Changed in Release	Where Documented
FICON Tape Read Acceleration	Added “FICON Tape Acceleration” section.	5.0(1a)	Configuring FICON



CHAPTER 2

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, on page 5
- Dynamic Port VSAN Membership, on page 6
- SAN Device Virtualization, on page 6
- Zoning, on page 6
- Distributed Device Alias Services, on page 7
- Fibre Channel Routing Services and Protocols, on page 7
- Multiprotocol Support, on page 7

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.

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 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.



Note SDV is not supported from Cisco MDS NX-OS Release 4.x and later.

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.



Note LUN zoning and read-only zones are not supported from Cisco MDS NX-OS Release 5.x and later.

- 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 polices 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.

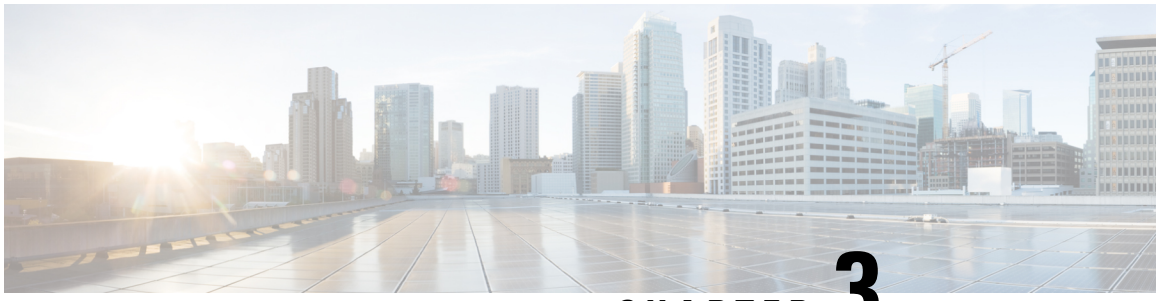
Fibre Channel Routing Services and Protocols

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.



CHAPTER 3

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:

- About VSANs, on page 9
- VSAN Configuration, on page 13
- Displaying Static VSAN Configuration , on page 21
- Default Settings, on page 21
- Displaying Fabric Switch Information, on page 22

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.

This section describes VSANs and includes the following topics:

VSANs Topologies

The switch icons shown in both Figure 1: Logical VSAN Segmentation, on page 10 and Figure 2: Example of Two VSANs, on page 11 indicate that these features apply to any switch in the Cisco MDS 9000 Family.

Figure 1: Logical VSAN Segmentation, on page 10 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.

Figure 1: Logical VSAN Segmentation

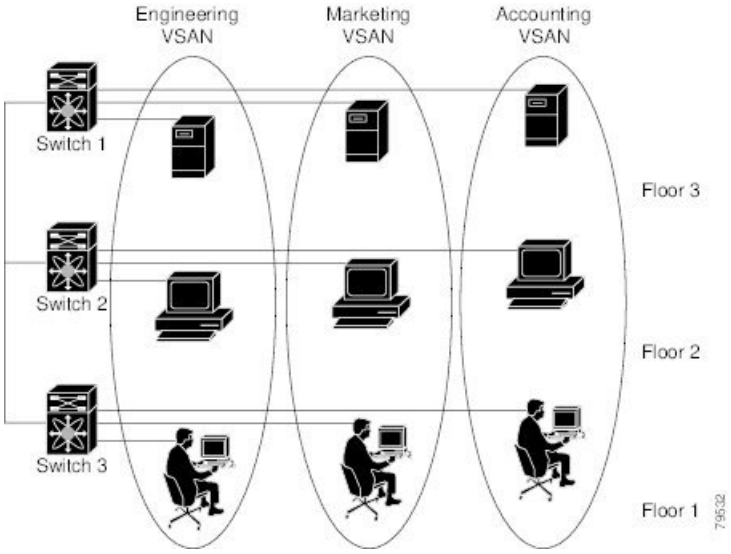
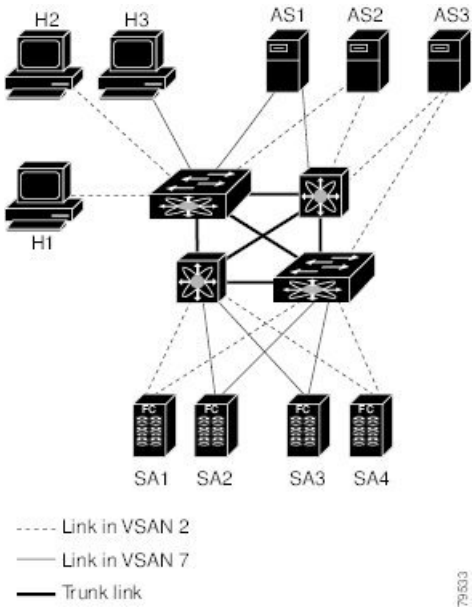


Figure 2: Example of Two VSANs, on page 11 shows a physical Fibre Channel switching infrastructure with two defined VSANs: VSAN 2 (dashed) and VSAN 7 (solid). VSAN 2 includes hosts H1 and H2, application servers AS2 and AS3, and storage arrays SA1 and SA4. VSAN 7 connects H3, AS1, SA2, and SA3.

Figure 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 2: Example of Two VSANs, on page 11 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:

- 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

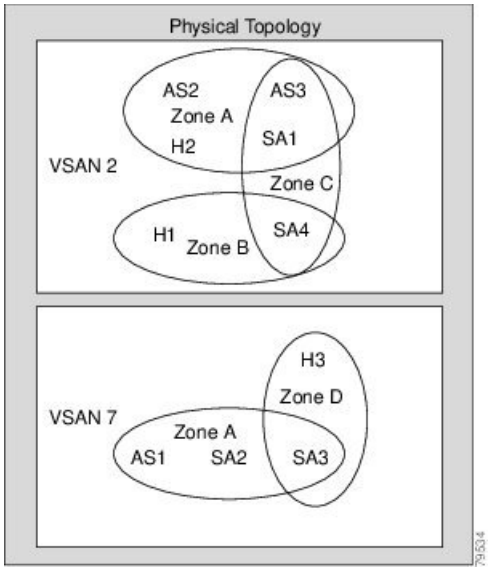
You can define multiple zones in a VSAN. Because two VSANs are equivalent to two unconnected SANs, zone A on VSAN 1 is different and separate from zone A in VSAN 2. Table 2: VSAN and Zone Comparison , on page 12 lists the differences between VSANs and zones.

Table 2: 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 3: VSANS with Zoning, on page 13 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.

Figure 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.

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 switch should work. Generation 2 switching modules support OX ID based load balancing of IVR traffic from IVR-enabled switches.

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

Reserved VSAN Range and Isolated VSAN Range Guidelines

On an NPV switch with a trunking configuration on any interface, or on a regular switch where the f port-channel-trunk command is issued to enable the Trunking F Port Channels feature, follow these configuration guidelines for reserved VSANs and the isolated VSAN:

- If trunk mode is on for any of the interfaces or NP Port Channel is up, the reserved VSANs are 3040 to 4078, and they are not available for user configuration.
- The Exchange Virtual Fabric Protocol (EVFP) isolated VSAN is 4079, and it is not available for user configuration.

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.

Creating VSANs Statically

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

Creating VSANs

To create VSANs, follow these steps:

Procedure

Step 1 switch# **config terminal**

Enters configuration mode.

Step 2 switch(config)# **vsan database**

switch(config-vsan-db)#

Configures the database for a VSAN. Application specific VSAN parameters cannot be configured from this prompt.

Step 3 switch(config-vsan-db)# **vsan 2**

Creates a VSAN with the specified ID (2) if that VSAN does not exist already.

Step 4 switch(config-vsan-db)# **vsan 2 name TechDoc**

updated vsan 2

Updates the VSAN with the assigned name (TechDoc).

Step 5 switch(config-vsan-db)# **vsan 2 suspend**

Suspends the selected VSAN.

Step 6 switch(config-vsan-db)# **no vsan 2 suspend**

Negates the **suspend** command issued in the previous step.

Step 7 switch(config-vsan-db)# **end**

switch#

Returns you to EXEC mode.

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, on page 15.

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

  See create_dynamic_vsan.ditamap#map_2861B3F48B334468BB9FBC52B85CC84A

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*).

Assigning Static Port VSAN Membership

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

Procedure

Step 1 switch# **config terminal**

Enters configuration mode.

Step 2 switch(config)# **vsan database**

switch(config-vsan-db)#

Configures the database for a VSAN.

Step 3 switch(config-vsan-db)# **vsan 2**

Creates a VSAN with the specified ID (2) if that VSAN does not exist already.

Step 4 switch(config-vsan-db)# **vsan 2 interface fc1/8**

Assigns the membership of the fc1/8 interface to the specified VSAN (VSAN 2).

Step 5 switch(config-vsan-db)# **vsan 7**

Creates another VSAN with the specified ID (7) if that VSAN does not exist already.

Step 6 switch(config-vsan-db)# **vsan 7 interface fc1/8**

Updates the membership information of the interface to reflect the changed VSAN.

Step 7 switch(config-vsan-db)# **vsan 1 interface fc1/8**

Removes the interface fc1/8 from VSAN 7 to VSAN 1(the default VSAN).

To remove the VSAN membership of interface fc1/8 from VSAN 7, you should define the VSAN membership of fc1/8 to another VSAN.

The best practice is to assign it back to VSAN 1.

Displaying VSAN Static Membership

To display the VSAN static membership information, use the **show vsan membership** command (see Displays Membership Information for the Specified VSAN, on page 16 through Displays Static Membership Information for a Specified Interface, on page 17).

Displays Membership Information for the Specified VSAN

```
switch # show vsan 1 membership
vsan 1 interfaces:
        fc1/1   fc1/2   fc1/3   fc1/4   fc1/5   fc1/6   fc1/7   fc1/9
        fc1/10  fc1/11  fc1/12  fc1/13  fc1/14  fc1/15  fc1/16  port-channel 99
```

Note Interface information is not displayed if interfaces are not configured on this VSAN.

Displays Static Membership Information for All VSANs

```
switch # show vsan membership

vsan 1 interfaces:
        fc2/16  fc2/15  fc2/14  fc2/13  fc2/12  fc2/11  fc2/10  fc2/9
        fc2/8   fc2/7   fc2/6   fc2/5   fc2/4   fc2/3   fc2/2   fc2/1
        fc1/16  fc1/15  fc1/14  fc1/13  fc1/12  fc1/11  fc1/10  fc1/9
        fc1/7   fc1/6   fc1/5   fc1/4   fc1/3   fc1/2   fc1/1
vsan 2 interfaces:
        fc1/8
vsan 7 interfaces:
```

```
vsan 100 interfaces:
vsan 4094(isolated vsan) interfaces:
```

Displays Static Membership Information for a Specified Interface

```
switch # show vsan membership interface fc1/1
fc1/1
        vsan:1
        allowed list:1-4093
```

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.

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).

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.

Displaying Isolated VSAN Membership

The **show vsan 4094 membership** command displays all ports associated with the isolated VSAN.

Operational State of a VSAN

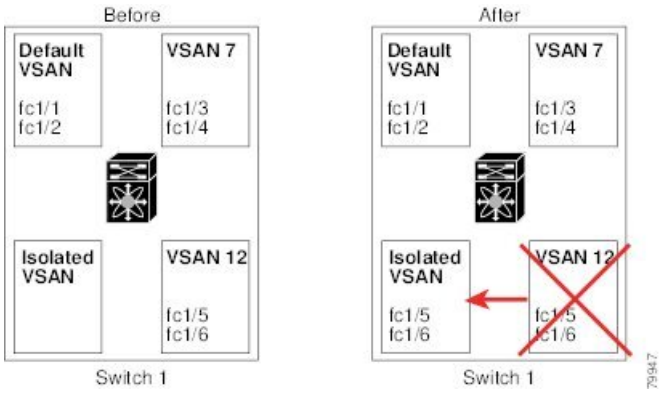
A VSAN is in the operational state if the VSAN is active and at least one port is up. This state indicates that traffic can pass through this VSAN. This state cannot be configured.

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 4: VSAN Port Membership Details, on page 18)

Figure 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.

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.

Deleting Static VSANs

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

Procedure

Step 1 switch# **config terminal**

Enters configuration mode.

Step 2 switch(config)# **vsan database**

Configures the VSAN database.

Step 3 switch-config-db# **vsan 2**

switch(config-vsan-db)#

Places you in VSAN configuration mode.

Step 4 switch(config-vsan-db)# **no vsan 5**

switch(config-vsan-db)#

Deletes VSAN 5 from the database and switch.

Step 5 switch(config-vsan-db)# **end**

switch#

Places you in EXEC mode.

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.

Note For FICON supported switches, load balancing can either be done as a very small disruption producing errors for open exchanges at the time of the change or if the change is required to be 100% disruptive, then either the ISLs can be taken down or the traffic crossing the ISLs can be quiesced.

Configuring Load Balancing

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

Note FICON supported switches do not support the following procedure.

Procedure

Step 1 switch# **config terminal**

Enters configuration mode.

Step 2 switch(config)# **vsan database**

switch(config-vsan-db)#

Enters VSAN database configuration submode

Step 3 switch(config-vsan-db)# **vsan 2**

Specifies an existing VSAN.

Step 4 switch(config-vsan-db)# **vsan 2 loadbalancing src-dst-id**

Enables the load balancing guarantee for the selected VSAN and directs the switch to use the source and destination ID for its path selection process.

Step 5 switch(config-vsan-db)# **no vsan 2 loadbalancing src-dst-id**

Negates the command issued in the previous step and reverts to the default values of the load balancing parameters.

Step 6 switch(config-vsan-db)# **vsan 2 loadbalancing src-dst-ox-id**

Changes the path selection setting to use the source ID, the destination ID, and the OX ID (default).

Step 7 switch(config-vsan-db)# **vsan 2 suspend**

Suspends the selected VSAN.

Step 8 switch(config-vsan-db)# **no vsan 2 suspend**

Negates the **suspend** command issued in the previous step.

Step 9 switch(config-vsan-db)# **end**

switch#

Returns you to EXEC mode.

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.

FICON VSANs

You can enable FICON in up to eight VSANs. See the FICON VSAN Prerequisites, on page 245.

Displaying Static VSAN Configuration

Use the **show vsan** command to display information about configured VSANs (see Examples Displays the Configuration for a Specific VSAN, on page 21 to Displays All VSANs, on page 21).

Displays the Configuration for a Specific VSAN

```
switch# show vsan 100
vsan 100 information
         name:VSAN0100 state:active
         in-order guarantee:no interoperability mode:no
         loadbalancing:src-id/dst-id/oxid
```

Displays the VSAN Usage

```
switch# show vsan usage
4 vsan configured
configured vsans:1-4
vsans available for configuration:5-4093
```

Displays All VSANs

```
switch# show vsan
vsan 1 information
         name:VSAN0001 state:active
         in-order guarantee:no interoperability mode:no
         loadbalancing:src-id/dst-id/oxid
vsan 2 information
         name:VSAN0002 state:active
         in-order guarantee:no interoperability mode:no
         loadbalancing:src-id/dst-id/oxid
vsan 7 information
         name:VSAN0007 state:active
         in-order guarantee:no interoperability mode:no
         loadbalancing:src-id/dst-id/oxid
vsan 100 information
         name:VSAN0100 state:active
         in-order guarantee:no interoperability mode:no
         loadbalancing:src-id/dst-id/oxid
vsan 4094:isolated vsan
```

Default Settings

Table 3: Default VSAN Parameters , on page 22 lists the default settings for all configured VSANs.

Table 3: 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).

Displaying Fabric Switch Information

Use the **show fabric switch information** vsan command to display information about each switch in the fabric in the given VSAN.

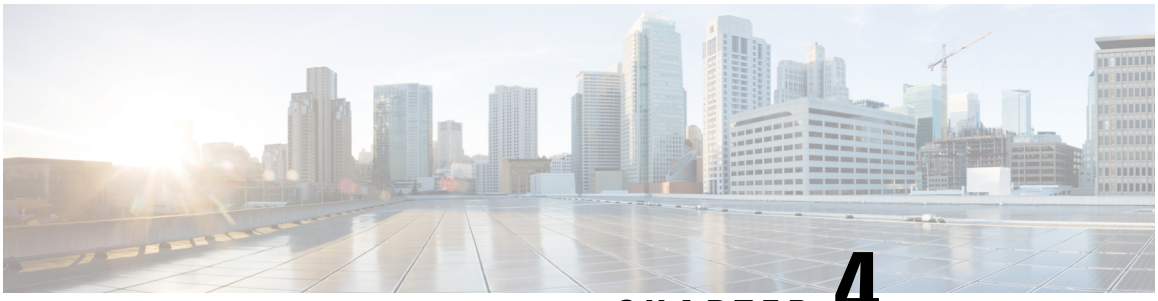
Displays Information about All the Switches in the Fabric

```
switch# show fabric switch information vsan 100
VSAN 1:
--------------------------------------------------------------------------------
SwitchName                      Model               Version     SupMemory
--------------------------------------------------------------------------------
huashan12                       DS-C9148-48P-K9     5.2(2d)           n/a
alishan-bgl-25                  DS-C9250I-K9        6.2(5a)           n/a
Hac18                           DS-C9506            6.2(7)          2 GB
Hac17                           DS-C9506            6.2(5)            n/a
Coco1                           DS-C9222I-K9        6.2(7)                    1 GB
switch#
```

Note This command is not supported prior to Cisco NX-OS Release 6.2(7).

Note SUP memory is not displayed for the switches that are running Cisco NX-OS Release prior to 6.2(7).

Note Without the VSAN option, this command displays information about switches in all the VSANs.



CHAPTER 4

Creating Dynamic VSANs

This chapter includes the following sections:

- About DPVM, on page 23
- DPVM Distribution, on page 29
- DPVM Configuration Merge Guidelines, on page 32
- Displaying DPVM Configurations, on page 34
- Sample DPVM Configuration, on page 35
- Default Settings, on page 38

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 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 Creating Dynamic VSANs, on page 23 .

DPVM configurations are based on port world wide name (pWWN) and node world wide name (nWWN) assignments. DPVM contains mapping information for each device pWWN/nWWN assignment and the corresponding VSAN. The Cisco NX-OS software checks DPVM active configuration 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 Series NX-OS System Management Configuration Guide .

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.

This section describes DPVM and includes the following topics:

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 Series 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).
- Device-alias must be in enhanced mode.

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.

Enabling DPVM

To begin configuring DPVM, you must explicitly enable DPVM on the required switches in the fabric. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family.

The configuration and verification commands for DPVM are only available when DPVM is enabled on a switch. When you disable this feature, all related configurations are automatically discarded.

To enable DPVM on any participating switch, follow these steps:

Procedure

Step 1 switch# **config t**
switch(config)#

Enters configuration mode.

Step 2 switch(config)# **feature dpvm**

Enables DPVM on that switch.

Step 3 switch(config)# **no feature dpvm**

Disables (default) DPVM on that switch.



Note To overwrite the login information with the duplicate pWWN login, enter the **dpvm overwrite-duplicate-pwwn** command.

DPVM Device Configuration (Static)

The DPVM device configuration 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.

Configuring DPVM

To configure DPVM, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **device-alias mode enhanced**

switch(config)# **device-alias commit**

Enables enhanced device alias mode.

Step 3 switch(config)# **dpvm database**

Creates the DPVM config database.

Step 4 switch(config)# **no dpvm database**

(Optional) Deletes the DPVM config database.

Step 5 switch(config-dpvm-db)# **pwwn 12:33:56:78:90:12:34:56 vsan 100**

Maps the specified device pWWN to VSAN 100.

Step 6 switch(config-dpvm-db)# **no pwwn 12:33:56:78:90:12:34:56 vsan 101**

(Optional) Removes the specified device pWWN mapping from the DPVM config database.

Step 7 switch(config-dpvm-db)# **nwwn 14:21:30:12:63:39:72:81 vsan 101**

Maps the specified device nWWN to VSAN 101.

Step 8 switch(config-dpvm-db)# **no nwwn 14:21:30:12:63:39:72:80 vsan 101**

(Optional) Removes the specified device nWWN mapping from the DPVM config database.

Step 9 switch(config-dpvm-db)# **device-alias device1 vsan 102**

Maps the specified device-alias to VSAN 102.

Step 10 switch(config-dpvm-db)# **no device-alias device1 vsan 102**

(Optional) Removes the specified device-alias mapping from the DPVM config database.

Step 11 switch(config-dpvm-db)# **show dpvm pending**

(Optional) When DPVM distribute is enabled (enabled by default when the feature is enabled) all configuration changes are held until they are committed. The list of pending changes can be seen at any time using this command.

Step 12 switch(config-dpvm-db)# **dpvm commit**

(Optional) When DPVM distribute is enabled (enabled by default when the feature is enabled) this command is required to commit the configuration changes.

Step 13 switch(config-dpvm-db)# **show dpvm database**

(Optional) Displays DPVM static device configuration.

Activating DPVM

Activating DPVM enforces the DPVM configuration. Activation may fail if there are conflicts between the already active configuration and the configuration to be activated. Activation can be forced to override the conflicting entries.

DPVM configuration can also be deactivated by issuing the **no dpvm activate** command.

To activate DPVM, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **dpvm activate**

Activates the DPVM configuration.

Step 3 switch(config)# **no dpvm activate**

Deactivates the currently active DPVM configuration.

Step 4 switch(config)# **dpvm activate force**

Forcefully activates the DPVM configuration and overrides the conflicting entries.

Step 5 switch(config)# **dpvm commit**

When DPVM distribute is enabled (enabled by default when the feature is enabled) this command is required to commit the configuration changes.

Step 6 switch(config)# **show dpvm database active**

(Optional) Displays the enforced DPVM device configuration.

DPVM Autolearn

DPVM can be configured to automatically learn (autolearn) new devices within each VSAN. DPVM autolearn can be enabled or disabled at any time. Learned entries are created by populating device pWWNs and VSANs and can be using the **show dpvm database active**. DPVM should be activated before autolearn can be enabled.

Auto learned entries can also be manually deleted. The auto learned entries become permanent when DPVM auto learn is disabled.

Note Autolearn 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, the corresponding autolearn entry is automatically deleted.
- 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.

Enabling Autolearn

To enable autolearn, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **dpvm auto-learn**

Enables autolearn on the switch.

Step 3 switch(config)# **no dpvm auto-learn**

Disables (default) autolearn on the switch.

Step 4 switch(config)# **clear dpvm auto-learn**

Clears the list of autolearned entries.

Step 5 switch(config)# **clear dpvm auto-learn pwwn** *pwwn*

Clears the list of autolearned pWWN entries in the distributed DPVM database.

Step 6 switch(config)# **dpvm commit**

When DPVM distribute is enabled (enabled by default when feature is enabled) any change to DPVM autolearn has to be committed before it can take effect locally and in the fabric.

Clearing Learned Entries

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

- To clear a single autolearn entry, use the **clear dpvm auto-learn pwwn** command.

```
switch# clear dpvm auto-learn pwwn 55:22:33:44:55:66:77:88
```

- To clear all autolearn entries, use the **clear dpvm auto-learn** command.

```
switch# clear dpvm auto-learn
```

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

Disabling Autolearn

To disable autolearn, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **no dpvm auto-learn**

Disables autolearn on the switch.

Note Running the **no dpvm auto-learn** command on other switches in the fabric before running the **dpvm commit** command helps to overcome the learnt conflict.

DPVM Distribution

If the DPVM configuration 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 NX-OS System Management Configuration Guide).

This section describes how to distribute DPVM and includes the following topics:

About DPVM Distribution

Using the CFS infrastructure, each DPVM server learns the DPVM configuration from each of its neighboring switches during the ISL bring-up process. Any configuration changes done locally are distributed in the fabric and updated by all switches in the fabric.

With DPVM distribution enabled, all DPVM configuration changes are stored temporarily and committed only when the **dpvm commit** command is run. Changes include the following tasks:

- Adding, deleting, or modifying DPVM device configuration.
- Activating or deactivating DPVM.
- Enabling or disabling autolearn.
- DPVM copy active configuration.

These changes are distributed to all switches in a fabric with the **dpvm commit** command. Changes can also discard via the **dvpm abort** command.

Tip Temporary changes made can be viewed by the **show dpvm pending** or **show dovm pending-diff** commands.

Disabling DPVM Distribution

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

Procedure

Step 1 switch# **configure terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **no dpvm distribute**

Disables DPVM distribution to the neighboring switches.

Step 3 switch(config)# **dpvm distribute**

Enables (default) DPVM distribution to the neighboring switches.

About Locking the Fabric

The first action that modifies the existing configuration creates the DPVM temporary storage and locks the feature in the fabric. Once the fabric is locked, no other user can make any further configuration changes to this feature.

Locking the Fabric

To lock the fabric and apply changes to the DPVM temporary storage, follow these steps:

Procedure

Step 1 switch# **config terminal**

Enters configuration mode.

Step 2 switch(config)# **dpvm database**

switch(config-dpvm-db)#

Accesses the DPVM configuration.

Step 3 switch(config-dpvm-db)# **pwwn 11:22:33:44:55:66:77:88 vsan 11**

Adds one entry to the DPVM configuration.

Step 4 switch(config-dpvm-db)# **exit**

Exits to configuration mode.

Step 5 switch(config)# **dpvm activate**

Run this command for the recent configuration changes to take effect.

Committing Changes

The **dpvm commit** command commits all the configuration changes made thus far on the local switch and also distributes the configurations to other switches in the fabric. On a successful commit, the configuration change is applied throughout the fabric and the lock is released.

To commit the DPVM configuration changes, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **dpvm commit**

Commits the pending changes. The changes can be viewed using the **show dpvm pending** or the **show dpvm pending-diff** commands.

Discarding Changes

The **dpvm abort** discards all the temporary DPVM changes made thus far. The configurations remain unaffected and the lock is released.

To discard the DPVM configuration changes, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **dpvm abort**

Discards the database entries that are currently in the DPVM pending database. Discards all the pending DPVM changes.

Clearing a Locked Session

If DPVM lock is held and not released either by committing or by discarding the changes, an administrator can still clear the DPVM session from any switch in the fabric. When a DPVM session is cleared, all pending DPVM changes are discarded and the fabric lock is released.

Tip Changes made to DPVM when distribution is enabled and held temporarily until the configuration changes are either committed or discarded. The configuration changes are discarded when the switch is restarted.

To use administrative privileges and release a locked DPVM session, use the **clear dpvm session** command in EXEC mode.

```
switch# clear dpvm session
```

DPVM Configuration Merge Guidelines

DPVM merge refers to a union of DPVM configuration across the fabric. For information about CFS merge support, refer to the Cisco MDS 9000 Family NX-OS System Management Configuration Guide for detailed concepts.

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

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

Caution If these conditions are not met, the merge will fail. The next distribution will forcefully synchronize the configurations and the activation states in the fabric.

This section describes how to merge DPVM configurations and includes the following topics:

About Copying DPVM DPVM Configurations

Note Fabric distribution is enabled and changes must be committed.

Copying DPVM Active Configuration

To copy the currently active DPVM configurations to the DPVM static configuration, use the **dpvm database copy** command.

```
switch# dpvm database copy active
Legend: "+" New Entry, "-" Missing Entry, "*" Possible Conflict Entry
------------------------------------------------------------------------
-  pwwn  12:33:56:78:90:12:34:56  vsan 100
-  nwwn  14:21:30:12:63:39:72:81  vsan 101
```

Comparing Database Differences

Compare the DPVM configurations as follows:

- Use the **dpvm database diff active** command to compare the active DPVM configuration with the static DPVM configuration.

```
switch# dpvm database diff active
Legend: "+" New Entry, "-" Missing Entry, "*" Possible Conflict Entry
-----------------------------------------------------------------------
-  pwwn  44:22:33:44:55:66:77:88  vsan 44
*  pwwn  11:22:33:44:55:66:77:88  vsan 11
```

- Use the **dpvm database diff config** command to compare the static DPVM configuration with the active DPVM configuration.

```
switch# dpvm database diff config
Legend: "+" New Entry, "-" Missing Entry, "*" Possible Conflict Entry
-----------------------------------------------------------------------
+  pwwn  44:22:33:44:55:66:77:88  vsan 44
*  pwwn  11:22:33:44:55:66:77:88  vsan 22
```

- Use the **show dpvm pending-diff** command (when CFS distribution is enabled) to compare the pending DPVM configuration changes.

Displaying DPVM Merge Status and Statistics

To display the DPVM configuration merge statistics, follow these steps:

Command	Purpose
`switch# show dpvm merge statistics`	Displays the DPVM configuration merge statistics.

`switch(config)# clear dpvm merge statistics`
 	
`switch(config)#`
	Clears the DPVM configuration merge statistics.

This example shows the conflicts in DPVM configuration merge:

```
switch# show dpvm merge status
Last Merge Time Stamp     : Fri Aug  8 15:46:36 2008
Last Merge State          : Fail
Last Merge Result         : Fail
Last Merge Failure Reason : DPVM DB conflict found during merge [cfs_status: 76] Last Merge
 Failure Details: DPVM merge failed due to database conflict
Local Switch WWN          : 20:00:00:0d:ec:24:e5:00
Remote Switch WWN         : 20:00:00:0d:ec:09:d5:c0

------------------------------------------------------------------------
        Conflicting DPVM member(s)                    Loc VSAN   Rem VSAN
------------------------------------------------------------------------
dev-alias dpvm_dev_alias_1 [21:00:00:04:cf:cf:45:ba]    1313       1414
```

```
dev-alias dpvm_dev_alias_2 [21:00:00:04:cf:cf:45:bb]    1313       1414
dev-alias dpvm_dev_alias_3 [21:00:00:04:cf:cf:45:bc]    1313       1414
[Total 3 conflict(s)]
rbadri-excal13#
```

This example shows the conflicts in DDAS mode:

```
switch# show dpvm merge status
Last Merge Time Stamp     : Fri Aug  8 15:46:36 2008
Last Merge State          : Fail
Last Merge Result         : Fail
Last Merge Failure Reason : DPVM DB conflict found during merge [cfs_status: 76] Last Merge
 Failure Details: DPVM merge failed due to DDAS mode conflict
Local Switch WWN          : 20:00:00:0d:ec:24:e5:00
Remote Switch WWN         : 20:00:00:0d:ec:09:d5:c0
Local DDAS mode           : Basic
Remote DDAS mode          : Enhanced
```

Displaying DPVM Configurations

Use the **show dpvm** command to display information about WWNs configured on a per VSAN basis (see the following examples).

Displays the DPVM Configuration Status

```
switch# show dpvm status
DB is activated successfully, auto-learn is on
```

Displays the DPVM Current Dynamic Ports for the Specified VSAN

```
switch# show dpvm ports vsan 10
----------------------------------------------------------------
Interface Vsan Device pWWN             Device nWWN
----------------------------------------------------------------
fc1/2     10   29:a0:00:05:30:00:6b:a0 fe:65:00:05:30:00:2b:a0
```

Displays the DPVM Configuration

```
switch# show dpvm database
pwwn  11:22:33:44:55:66:77:88  vsan 11
pwwn  22:22:33:44:55:66:77:88  vsan 22
pwwn  33:22:33:44:55:66:77:88  vsan 33
pwwn  44:22:33:44:55:66:77:88  vsan 44
[Total 4 entries]
```

Displays the DPVM Active Configurations

```
switch# show dpvm database active
pwwn  11:22:33:44:55:66:77:88  vsan 22
pwwn  22:22:33:44:55:66:77:88  vsan 22
pwwn  33:22:33:44:55:66:77:88  vsan 33
[Total 3 entries]
* is auto-learnt entry
```

Displays DPVM Configurations

```
switch# show dpvm database
pwwn  11:22:33:44:55:66:77:88  vsan 11
pwwn  22:22:33:44:55:66:77:88  vsan 22
pwwn  33:22:33:44:55:66:77:88  vsan 33
pwwn  44:22:33:44:55:66:77:88  vsan 44
[Total 4 entries]
```

Shows Pending Changes with Respect to the DPVM Configurations

```
switch# show dpvm pending-diff
Legend: "+" New Entry, "-" Missing Entry, "*" Possible Conflict Entry
--------------------------------------------------------------------
+  pwwn  55:22:33:44:55:66:77:88  vsan 55
-  pwwn  11:22:33:44:55:66:77:88  vsan 11
*  pwwn  44:22:33:44:55:66:77:88  vsan 44
```

Sample DPVM Configuration

To configure a basic DPVM scenario, follow these steps:

Procedure

Step 1 Enable DPVM and enable DPVM distribution.

Example:

```
switch1# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
switch1(config)# feature dpvm
switch1(config)# end

switch1# show dpvm database
switch1# show dpvm database active
switch1# show dpvm status
```

At this stage, the configuration does not have an active DPVM configuration and the **auto-learn** option is disabled.

Step 2 Activate a null (empty) configuration so that it can be populated with autolearned entries.

Example:

```
switch1# configure terminal

Enter configuration commands, one per line.  End with CNTL/Z.
switch1(config)# dpvm activate
switch1(config)# dpvm commit
switch1(config)# end

switch1# show dpvm database

switch1# show dpvm database active

switch1# show dpvm status
```

At this stage, the database is successfully activated and the **auto-learn** option continues to be disabled.

Step 3 Enable the **auto-learn** option and commit the configuration changes.

Example:

```
switch1# configure terminal

Enter configuration commands, one per line.  End with CNTL/Z.
switch1(config)# dpvm auto-learn
switch1(config)# dpvm commit
switch1(config)# end

switch1# show dpvm database active
pwwn 21:00:00:e0:8b:0e:74:8a vsan 4(*)
pwwn 21:01:00:e0:8b:2e:87:8a vsan 5(*)
[Total 2 entries]
* is auto-learnt entry
switch1# show dpvm ports
------------------------------------------------------------
Interface   Vsan        Device pWWN       Device nWWN
------------------------------------------------------------
fc1/24       4   21:00:00:e0:8b:0e:74:8a  20:00:00:e0:8b:0e:74:8a
fc1/27       5   21:01:00:e0:8b:2e:87:8a  20:01:00:e0:8b:2e:87:8a
switch1# show flogi database
--------------------------------------------------------------------------------
INTERFACE  VSAN    FCID            PORT NAME               NODE NAME
--------------------------------------------------------------------------------
fc1/24     4     0xe70100  21:00:00:e0:8b:0e:74:8a  20:00:00:e0:8b:0e:74:8a
fc1/27     5     0xe80100  21:01:00:e0:8b:2e:87:8a  20:01:00:e0:8b:2e:87:8a
Total number of flogi = 2.
switch195# show dpvm status
DB is activated successfully, auto-learn is on
```

At this stage, the currently logged in devices (and their current VSAN assignment) populate the active DPVM configuration. However the entries are not yet permanent in the active DPVM configuration.

The output of the **show dpvm ports** and the **show flogi database** commands displays two other devices that have logged in (referred to as switch9 and switch3 in this sample configuration).

Step 4 Access switch9 and issue the following commands:

Example:

```
switch9# show dpvm database active
pwwn 21:00:00:e0:8b:0e:87:8a vsan 1(*)
pwwn 21:01:00:e0:8b:2e:74:8a vsan 1(*)
[Total 2 entries]
* is auto-learnt entry
switch9# show dpvm status
DB is activated successfully, auto-learn is on
```

Step 5 Access switch3 and issue the following commands:

Example:

```
switch3# show dpvm database active
pwwn 21:00:00:e0:8b:0e:76:8a vsan 1(*)
pwwn 21:01:00:e0:8b:2e:76:8a vsan 1(*)
[Total 2 entries]
* is auto-learnt entry
switch3# show dpvm status
DB is activated successfully, auto-learn is on
```

Step 6 Disable autolearning in switch1 and commit the configuration changes.

Example:

```
switch1# configure terminal

Enter configuration commands, one per line.  End with CNTL/Z.
switch1(config)# no dpvm auto-learn
switch1(config)# dpvm commit
switch1(config)# end

switch1# show dpvm status
DB is activated successfully, auto-learn is off
switch1# show dpvm database active
pwwn 21:00:00:e0:8b:0e:74:8a vsan 4
pwwn 21:01:00:e0:8b:2e:87:8a vsan 5
pwwn 21:00:00:e0:8b:0e:87:8a vsan 1
pwwn 21:01:00:e0:8b:2e:74:8a vsan 1
pwwn 21:00:00:e0:8b:0e:76:8a vsan 1
pwwn 21:01:00:e0:8b:2e:76:8a vsan 1
[Total 6 entries]
* is auto-learnt entry
switch1# show dpvm status
DB is activated successfully, auto-learn is off
```

At this stage, the autolearned entries are made permanent in the active DPVM configuration.

Step 7 Access switch9 and issue the following commands:

Example:

```
switch9# show dpvm database active
pwwn 21:00:00:e0:8b:0e:87:8a vsan 1
pwwn 21:01:00:e0:8b:2e:74:8a vsan 1
pwwn 21:00:00:e0:8b:0e:76:8a vsan 1
pwwn 21:01:00:e0:8b:2e:76:8a vsan 1
pwwn 21:00:00:e0:8b:0e:74:8a vsan 4
pwwn 21:01:00:e0:8b:2e:87:8a vsan 5
[Total 6 entries]
* is auto-learnt entry
switch9# show dpvm status
DB is activated successfully, auto-learn is off
```

Step 8 Access switch3 and issue the following commands:

Example:

```
switch3# show dpvm database active
pwwn 21:00:00:e0:8b:0e:76:8a vsan 1
pwwn 21:01:00:e0:8b:2e:76:8a vsan 1
pwwn 21:00:00:e0:8b:0e:87:8a vsan 1
pwwn 21:01:00:e0:8b:2e:74:8a vsan 1
pwwn 21:00:00:e0:8b:0e:74:8a vsan 4
pwwn 21:01:00:e0:8b:2e:87:8a vsan 5
[Total 6 entries]
* is auto-learnt entry
switch3# show dpvm status
DB is activated successfully, auto-learn is off
```

Note These basic steps help you determine that the information is identical in all the switches in the fabric.

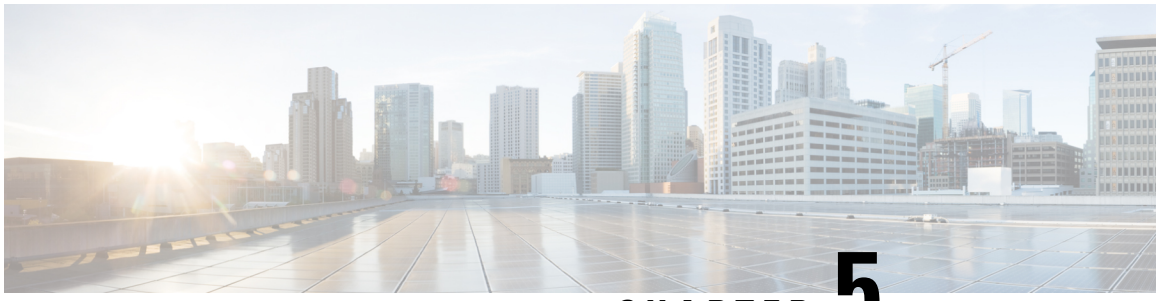
You have now configured a basic DPVM scenario in a Cisco MDS 9000 Series switch.

Default Settings

Table 4: Default DPVM Parameters , on page 38 lists the default settings for DPVM parameters.

Table 4: Default DPVM Parameters

Parameters	Default
DPVM	Disabled.
DPVM distribution	Enabled.
Autolearn	Disabled.



CHAPTER 5

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.

- Finding Feature Information, on page 39
- Feature History for Configuring and Managing Zones, on page 40
- About Zoning, on page 40
- Autozone, on page 45
- Zone Configuration, on page 56
- Zone Sets and FC Aliases, on page 60
- ZoneSet Distribution, on page 69
- Zoneset Duplication, on page 72
- Advanced Zone Attributes, on page 75
- Displaying Zone Information, on page 84
- Enhanced Zoning, on page 92
- Controlling Zoning Configuration Sessions, on page 112
- Compacting the Zone Database for Downgrading, on page 113
- Zone and ZoneSet Analysis, on page 114
- Zoning Best Practice, on page 117
- Enhancing Zone Server Performance, on page 130
- Zone Server SNMP Optimization, on page 131
- Zone Server Delta Distribution, on page 132
- Default Settings , on page 134

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at https://bst.cloudapps.cisco.com/bugsearch and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information"chapter or the Feature History table in this chapter.

Feature History for Configuring and Managing Zones

lists the New and Changed features.

Table 5: New and Changed Features

Feature Name	Releases	Feature Information
Autozone	8.5(1)	• The maximum number of devices supported by zones in Autozone is increased to 250. • Autozone can now be enabled on other VSANs apart from VSAN 1. The **autozone --enable --vsan** *id* command was modified.
Zoning Single Session	8.4(2)	Introduced the single session option for enhanced zoning mode. The following commands were modified: • [**no**] **zone mode enhanced vsan** *id* [**single-session**] • **show zone status vsan** *id*
Autozone	8.4(1)	The **enableautosave** and **disableautosave** options are added to the **autozone** command to enable or disable automatically saving of the running configuration to the startup-configuration after making a zoning change. The following commands were introduced: • **autozone --enable** • **autozone --enableautosave** • **autozone --disableautosave**
Autozone	8.3(1)	The Autozone feature was introduced. The following commands were introduced: • **autozone --delete** • **autozone --disable** • **autozone --help** • **autozone --show** • **autozone --showpending** • **autozone --update**

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 zoneset) is a member of the default zone.
  - Zones can vary in size.
  - Devices can belong to more than one zone.
- A zoneset consists of one or more zones.
  - A zoneset can be activated or deactivated as a single entity across all switches in the fabric.
  - Only one zoneset can be activated at any time.
  - A zone can be a member of more than one zoneset.
  - An MDS switch can have a maximum of 1000 zonesets.
- 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 zoneset. 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 nondisruptively. 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.
  - Symbolic-nodename—Specifies the member symbolic node name. The maximum length is 240 characters.

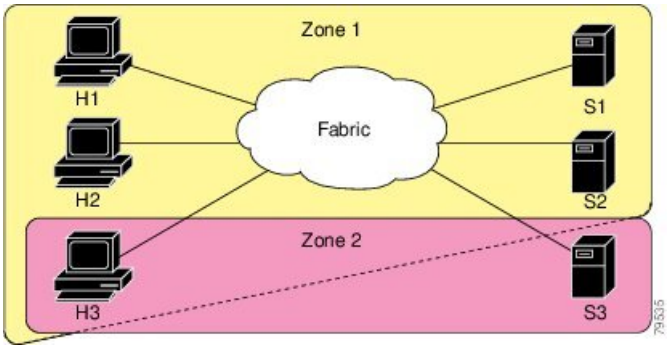
- Default zone membership includes all ports or WWNs that do not have a specific membership association. Access between default zone members is controlled by the default zone policy.

Note For configuration limits on configuring the number of zones, zone members and zone sets, refer to the Cisco MDS NX-OS Configuration Limits.

Zoning Example

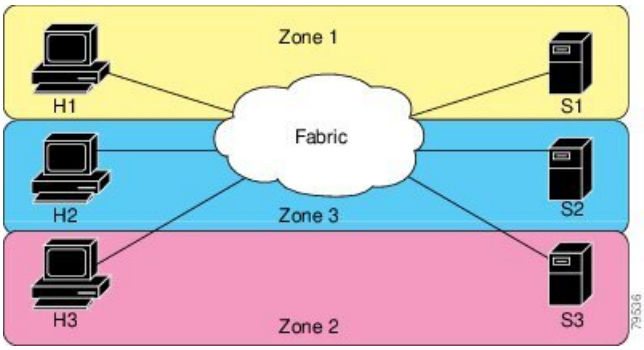
Figure 5: Fabric with Two Zones , on page 42 illustrates a zoneset with two zones, zone 1 and zone 2, in a fabric. Zone 1 provides access from all three hosts (H1, H2, H3) to the data residing on storage systems S1 and S2. Zone 2 restricts the data on S3 to access only by H3. Note that H3 resides in both zones.

Figure 5: Fabric with Two Zones



There are other ways to partition this fabric into zones. Figure 6: Fabric with Three Zones , on page 42 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 6: Fabric with Three Zones



Zone Implementation

All switches in the Cisco MDS 9000 Series 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 zoneset 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 zoneset is active and you activate another zoneset) 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.

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. The best practice is to 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. For this reason, 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.
- While configuring a zone member based on interface type always select a fabric switch which potentially has the highest interface count in the fabric.

Active and Full Zoneset Considerations

Before configuring a zoneset, consider the following guidelines:

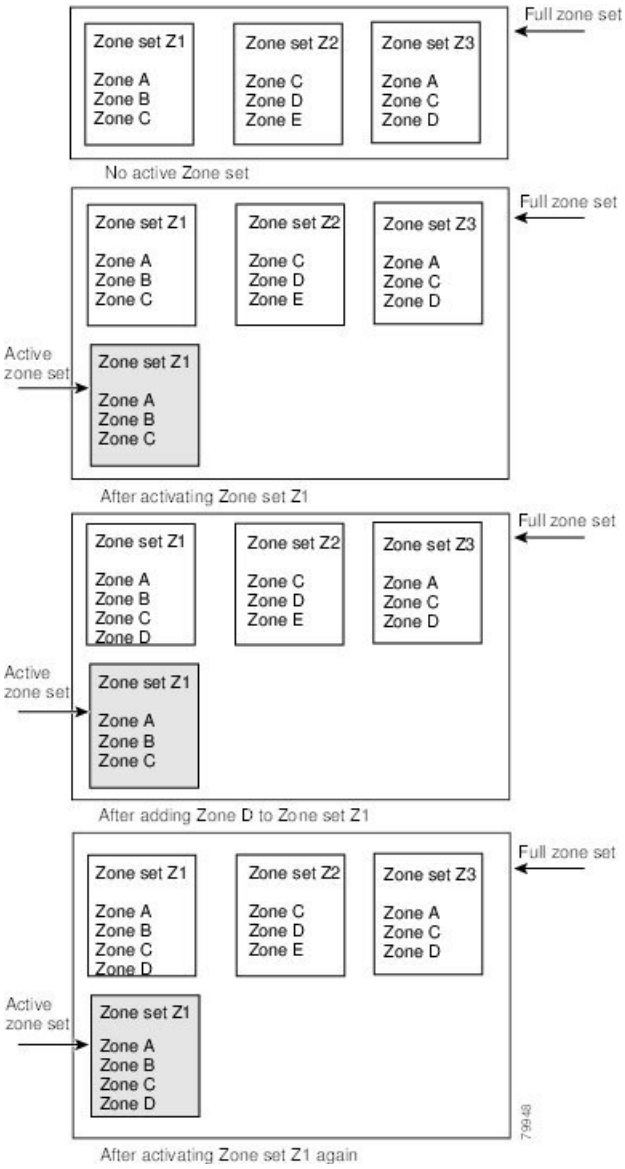
- Each VSAN can have multiple zone sets but only one zoneset can be active at any given time.
- When you create a zoneset, that zoneset becomes a part of the full zoneset.

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

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

Figure shows a zone being added to an activated zoneset.

Figure 7: Active and Full Zone Sets



Autozone

The Autozone feature is a mechanism to automate zoning of devices. This feature can be used to reduce the administrative overhead of manually creating and updating the switch zone configuration each time a device is added to the SAN to a one-time command. An administrator has to configure the Autozone feature after the initial deployment and does not have to manually change or modify the zone configuration each time a new device is added to a fabric. The Autozone feature is intended for fabrics comprise of a single fabric switch that has no more than 100 devices connected.

Initially, Autozone configures zoning that enables connectivity from every initiator to every target, based on the FC4 type registered by each device. The zones that are created are placed in a single zoneset in VSAN 1

and activated. When in automatic mode, a scheduler job is created to scan for newly logged in devices every 5 minutes. New initiators are zoned with all the targets and new targets are zoned with all the initiators. The new zones are then added to the active zoneset. This process allows the switch to be administered with minimal effort by simply plugging in new devices and having automatic connectivity for the devices within minutes. Autozone can be run manually by the administrator if connectivity to the newly logged-in devices is required before the next scheduled scan. Autozone does not change the existing zones created either by Autozone or manually by an administrator. This prevents duplication of any existing zones by Autozone and allows specialized zones to be added manually by an administrator.

Autozone has two modes of operation:

- Automatic mode—The Autozone scheduler job runs every 5 minutes, checking for changes in device logins, and updates the zoneset accordingly.
- Manual mode—No scheduler job is created. The administrator has to run the **autozone --update** command each time a new device is connected to a switch for the device to be added to the zoning configuration.

Guidelines and Limitations for Autozone

- Works only on Cisco MDS 9132T, MDS 9148T, MDS 9396T, MDS 9124V, and MDS 9148V fabric switches.
- Works only on single-switch fabrics.
- From Cisco MDS NX-OS Release 8.5(1), Autozone can be enabled on VSANs apart from VSAN 1, but can be enabled on only one VSAN per switch.
- In Cisco MDS NX-OS Release 8.4(2b) and earlier release, Autozone works only for ports that are logged on VSAN 1. If the administrator moves ports to other VSANs, Autozone does not move them back into VSAN 1 or zone them.
- If Autozone detects an active zoneset with a name different from AUTOZONESET, Autozone exits with a message without making any changes to the existing zone configuration.
- If Autozone detects an Inter-Switch Link (ISL), Autozone exits with a message, and no zones are created.
- Autozone will not work when the default zone is enabled.
- The Autozone feature considers only devices that are registered as FC4 type *init* or *target*. Since devices that are registered as *both* are considered as both *init* and *target*, starting from Cisco MDS NX-OS Release 8.4(2), the Autozone feature will zone these devices with devices that register as *init*, *target*, and *both*. Any other types are ignored and must be zoned manually by the administrator.
- Starting with Cisco MDS NX-OS Release 8.5(1), the Autozone feature zones a maximum of 250 devices. In releases prior to Cisco MDS NX-OS Release 8.5(1), Autozone feature zones a maximum of 100 devices.
- The Autozone feature does not support Smart Zoning.
- Avoid enabling the Autozone feature if you want to use the Inter VSAN Routing (IVR) feature.
- Avoid creating a manual zone of the name format AUTOZONE_<*SwitchSerialNumber*>_<*number*> because this is the format that is used by Autozone to create zone names. Autozone will delete zones with this name format when you use the **autozone --delete** command.

- When Autozone is first run in automatic mode, it creates a scheduler job called AUTOZONE_SCHEDULER_JOB and a schedule called AUTOZONE_SCHEDULER_SCHEDULE to run the **autozone --update** command every 5 minutes. If either the scheduler job or schedule is deleted by the administrator, then periodic zone updates by Autozone will cease.
- When Autozone is enabled and if a zone lock or zone single-session lock is acquired, you will have to clear the zone lock using the **clear zone lock vsan** command and then retry Autozone configuration.
- When Autozone is configured in automatic mode and you execute the **show accounting log** command, there will be entries with an empty command field each time the Autozone scheduler job is run. This is expected.
- Cisco NX-OS releases that support the Autozone feature creates a CLI alias named *autozone* when the switch boots. Even if the **autozone --enable** command is not run, this configuration change triggers the "Unsaved configuration" warning during an upgrade. Ensure that you save the configuration to prevent this message during future upgrades. As a best practice, we recommend that you copy the running configuration to the startup configuration on the switch before upgrading.
- If you have the Autozone feature enabled and you are downgrading from Cisco MDS NX-OS Release 8.3(1) to earlier releases that do not support the Autozone feature, the Autozone scheduler job continues to check for new device logins every 5 minutes. However, the scheduler job fails if any new device logins are detected, and generates a syslog. Therefore, we recommend that you disable Autozone before downgrading.
- If you have the Autozone feature enabled and you are downgrading from Cisco MDS NX-OS Release 8.4(1) to earlier releases that do not support the Autozone feature, the **autozone --enable** CLI alias command is available. However, when you execute the command, it fails. You may delete the autozone CLI alias name using the **cli alias name autozone** command.
- Do not delete the *autozone* CLI alias name to ensure that the Autozone feature functions as documented.
- If the Autozone feature is enabled, during upgrade or downgrade, the Autozone scheduler job might temporarily fail. The scheduler job resumes executing successfully after the upgrade or downgrade is complete.

Configuring Autozone in Automatic Mode

The Autozone feature creates zones and a zoneset on VSAN 1 for unzoned devices and creates a scheduler job to periodically add new device logins on VSAN 1.

Enabling Autozone in Automatic Mode

Before you begin

Review the Guidelines and Limitations for Autozone, on page 46.

Procedure

Enable autozone to automatically create zones, add them to a zoneset, and activate the zoneset as needed every 5 minutes:

switch# **autozone --enable --vsan** *id*

Note **--vsan** *id* is optional and defaults to VSAN 1.

Enabling Autosave

To enable autozone to automatically save the running-configuration to the startup-configuration after it makes a zoning change, perform the step below:

Before you begin

Enable Autozone in Automatic Mode.

Procedure

Enable automatic saving of the autozone configuration:

switch# **autozone --enableautosave**

Executing Autozone in Manual mode

You can execute Autozone manually for updating the zoning information each time a new device is logged into the switch.

To execute Autozone in manual mode, perform this step:

switch# **autozone --update**

Enabling Autozone in Automatic Mode by a Remotely Authenticated (AAA) User

An Autozone scheduler job runs using the identity of a user who has enabled the Autozone feature on the switch. If the user has remote authentication (AAA), the user's credentials must be added manually to the scheduler configuration for the periodic Autozone scheduler job to succeed.

To enable the Autozone feature for a remotely authenticated user, perform these steps:

Procedure

Step 1 Enter global configuration mode:

switch# **configure**

Step 2 Enable the command scheduler:

switch(config)# **feature scheduler**

Step 3 Configure a cleartext password for a remotely authenticated user:

switch(config)# **scheduler aaa-authentication user** *name* **password** *password*

Step 4 Create zones and a zoneset on a VSAN automatically and schedule a timer to check for new device logins:

switch(config)# **autozone --enable --vsan** *id*

Note **--vsan** *id* is optional and defaults to VSAN 1.

Disabling Autosave

To disable autozone from automatically saving the running-configuration to the startup-configuration after it makes a zoning change, perform the step below:

Procedure

Disable automatic saving of the autozone configuration:

switch# **autozone --disableautosave**

Disabling Autozone Automatic Mode

To prevent new devices from being zoned automatically and retaining the existing zone configuration, run this command:

switch# **autozone --disable**

Displaying All the Zone Configurations

To display the Autozone status, the existing zone and zoneset configuration created by Autozone, and the zoning configuration that Autozone would create for unzoned devices that are currently logged in to a switch, run this command:

switch# **autozone --show**

Displaying Pending Zone Configurations

To display only the zone configuration changes which Autozone configures for unzoned devices before the Autozone scheduler job is run, run this command:

switch# **autozone --showpending**

Applying Pending Zone Configurations (Manual Mode)

By default, if the Autozone feature is enabled, the Autozone scheduler job automatically runs every 5 minutes. However, to optionally force Autozone to run before the end of the 5-minute cycle or to run Autozone without creating the Autozone scheduler job, run this command:

switch# **autozone --update**

Deleting Zones and Zoneset Created by Autozone

To delete all the zones and the zoneset created by Autozone on VSAN 1, run this command:

switch# **autozone --delete**

Note Deleting the zones and zoneset created by Autozone does not disable the Autozone feature. To disable the Autozone feature, use the **autozone --disable** command. We recommend that you use the **autozone --disable** command before using the **autozone --delete** command because Autozone reconfigures all the zones again if it is enabled and the devices are still connected. Optionally, you can use both the options together using the **autozone --disable --delete** command.

Example: Configuring Autozone

The following example shows how to enable Autozone in automatic mode. In this mode, all currently logged in devices are zoned and new logins are automatically added periodically. In this example, a device without a suitable FC4 type is detected and not included in the zone configuration.

```
switch# autozone --enable --vsan 1
This command will automatically create and activate single-initiator and single-target zones
 for all end-devices currently logged-in to VSAN 1; all initiators will be zoned to all
targets. This may lead to a large TCAM and RSCN load on the switch. Please use AutoZone
judiciously.

AutoZone feature is enabled

Device with pwwn 10:00:00:de:fb:74:e8:31 is not registered with FC4-type Init or Target.
Hence, it will be ignored for AutoZone configuration.
Configuring zones for vsan 1
                AUTOZONE_JPG21190082_1

Configuring zoneset for vsan 1
Activating the zoneset. Please wait...
Configured zoneset AUTOZONESET for vsan 1 successfully.
```

The following example shows how to enable Autozone in automatic mode on VSAN 2:

```
switch# autozone --enable --vsan 2
This command will automatically create and activate single-initiator and single-target zones
 for all end-devices currently logged-in to VSAN 2; all initiators will be zoned to all
targets. This may lead to a large TCAM and RSCN load on the switch. Please use AutoZone
judiciously.

AutoZone feature is enabled

Device with pwwn 10:00:00:de:fb:74:e8:31 is not registered with FC4-type Init or Target.
Hence, it will be ignored for AutoZone configuration.
Configuring zones for vsan 2
                AUTOZONE_JPG21190082_1

Configuring zoneset for vsan 2
```

```
Activating the zoneset. Please wait...
Configured zoneset AUTOZONESET for vsan 2 successfully.
```

The following example shows how to execute the Autozone feature once to zone all the unzoned devices logged in on VSAN 1, and add them to the active zoneset of VSAN 1 without creating the Autozone scheduler job. A device without a suitable FC4 type is detected and not included in the zone configuration.

```
switch# autozone --update
Device with pwwn 10:00:00:de:fb:74:e8:31 is not registered with FC4-type Init or Target.
Hence, it will be ignored for AutoZone configuration.
Configuring zones for vsan 1
                AUTOZONE_JPG21190082_1
                AUTOZONE_JPG21190082_2
                AUTOZONE_JPG21190082_3
                AUTOZONE_JPG21190082_4
Configuring zoneset for vsan 1
Activating the zoneset. Please wait...
Configured zoneset AUTOZONESET for vsan 1 successfully.
```

The following example shows how to disable the Autozone feature so that newly logged-in devices are not zoned while retaining the existing zone configuration:

```
switch# autozone --disable
This will disable the AutoZone feature. Do you wish to continue? [y/n]|y: y

AutoZone feature disabled successfully.
```

The following example shows how to delete the Autozone and zoneset created for VSAN 1:

```
switch# autozone --delete
Checking if zoneset name AUTOZONESET present on switch...[Found]
Checking if AutoZone is enabled on switch...[Disabled]

This option will only delete the zone/zoneset configurations done by AutoZone feature.
Do you wish to continue? [n]|y: y
Deleting zoneset name AUTOZONESET and all zones for vsan 1 configured by AutoZone
Deleting following zones -
   AUTOZONE_JPG21190082_1
   AUTOZONE_JPG21190082_2
   AUTOZONE_JPG21190082_3
   AUTOZONE_JPG21190082_4
Deactivating zoneset for vsan 1.
Deactivated zoneset for vsan 1.
```

Verifying Autozone Configuration

The following example displays the Autozone status, the zones already created, as well as uncreated (pending) zones, by Autozone:

```
switch# autozone --show
Feature AutoZone : Enabled
AutoSave Configuration : Enabled
The possible zone/zoneset configuration with AutoZone feature for currently logged-in devices
```

```
 is :
zoneset name AUTOZONESET vsan 1
        zone name AUTOZONE_JPG21190082_1 vsan 1
                 member pwwn 20:00:00:11:0d:97:00:01
                 member pwwn 20:01:00:11:0d:97:01:01
        zone name AUTOZONE_JPG21190082_2 vsan 1
                 member pwwn 20:00:00:11:0d:97:00:01
                 member pwwn 20:01:00:11:0d:97:01:00
        zone name AUTOZONE_JPG21190082_3 vsan 1
                 member pwwn 20:00:00:11:0d:97:00:00
                 member pwwn 20:01:00:11:0d:97:01:01
        zone name AUTOZONE_JPG21190082_4 vsan 1
                 member pwwn 20:00:00:11:0d:97:00:00
                 member pwwn 20:01:00:11:0d:97:01:00
```

The following example shows how to check for the zoning configuration that Autozone creates for any unzoned devices and then apply those changes. In this example, Autozone is disabled so that zoning is updated only once and there is no periodic zoning by Autozone.

```
switch# autozone --showpending
Feature AutoZone : Disabled
zoneset name AUTOZONESET vsan 1
        zone name AUTOZONE_JPG21190082_1 vsan 1
                 member pwwn 20:00:00:11:0d:97:00:00
                 member pwwn 20:01:00:11:0d:97:01:00
switch# autozone --update
Configuring zones for vsan 1
                AUTOZONE_JPG21190082_1
Configuring zoneset for vsan 1
Activating the zoneset. Please wait...
Configured zoneset AUTOZONESET for vsan 1 successfully.
```

The following example shows how to verify if the Autozone feature is already enabled and if there are currently unzoned devices:

```
switch# autozone --showpending
Feature AutoZone : Enabled
zoneset name AUTOZONESET vsan 1
        zone name AUTOZONE_JPG21190082_2 vsan 1
                 member pwwn 20:00:00:11:0d:97:00:01
                 member pwwn 20:01:00:11:0d:97:01:00
        zone name AUTOZONE_JPG21190082_3 vsan 1
                 member pwwn 20:00:00:11:0d:97:00:00
                 member pwwn 20:01:00:11:0d:97:01:01
        zone name AUTOZONE_JPG21190082_4 vsan 1
                 member pwwn 20:00:00:11:0d:97:00:01
                 member pwwn 20:01:00:11:0d:97:01:01
```

The following example shows how to obtain information about the **autozone** command:

```
switch# autozone --help
usage: autozone.py [-h] [--enable] [--disable] [--update] [--delete] [--show]
                   [--showpending] [--enableautosave] [--disableautosave]
                   [--vsan VSAN]

Enables AutoZone feature for vsan 1

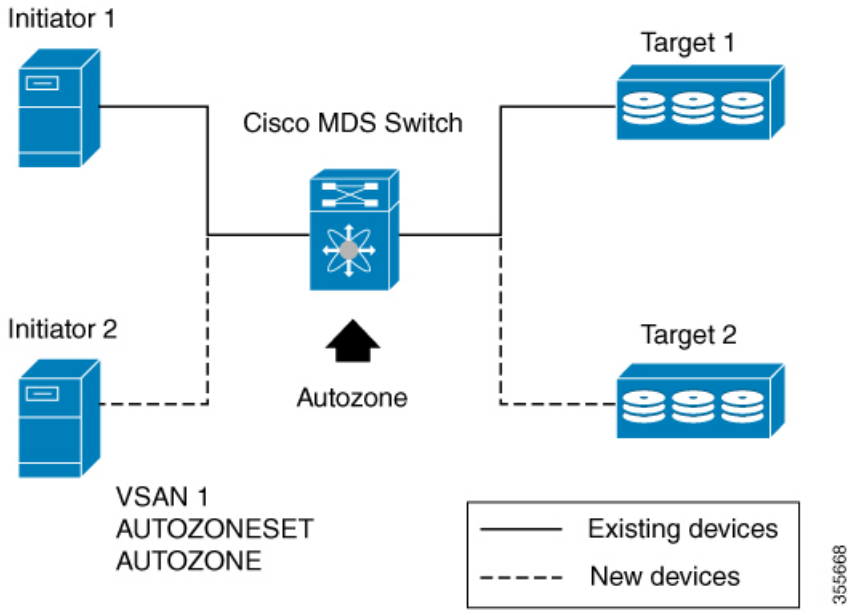
optional arguments:
  -h, --help         show this help message and exit
```

```
--enable            Enables AutoZone automatic mode for VSAN 1. New devices
                    logging in will be zoned automatically. No changes will
                    be done for existing configuration. To have autzone
                    automatically save the running configuration to startup
                    configuration include the --enable argument followed by
                    --enableautosave argument.
--disable           Disables AutoZone feature for VSAN 1. New devices logging
                    in will not be zoned automatically. No changes will be
                    done for existing configuration.
--update            Computes and applies any pending AutoZone configuration
                    to switch for vsan 1
--delete            Deletes zone/zoneset configuration done by AutoZone for VSAN
                    1
--show              Displays the current active zone/zonset configuration done by
                    Autozone for VSAN 1.
--showpending       Displays only zoning configuration that is pending and
                    not yet applied on the switch.
--enableautosave    Enables Auto saving of running configuration to startup
                    configuration whenever an automatic zoning change is
                    done. Allowed with the --enable argument and --update
                    argument respectively.
--disableautosave   Disables Auto saving of running configuration to startup
                    configuration whenever an automatic zoning change is
                    done.. To save any automatic zoning changes to startup,
                    "copy running-config startup-config" must be manually
                    executed.
--vsan VSAN         Please provide VSAN between 1-4093
```

Autozone Example Scenario

Let us consider a topology where two devices—Initiator 1 and Target 1—are logged on to a Cisco MDS switch. We configure the Autozone feature on the switch and verify the zone configuration for these devices. Then, we introduce two new devices—Initiator 2 and Target 2—to this network and verify if they are automatically configured in the zone.

Figure 8: Autozone Example Topology



1. Verify the existing zone configuration using the **show zoneset active vsan 1** command:

```
switch# show zoneset active vsan 1
Zoneset not present
```

2. Verify the existing device logins using the **show fcns database** command:

```
switch# show fcns database
VSAN 1:
--------------------------------------------------------------------------
FCID        TYPE    PWWN          (VENDOR)           FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0xee0000    N     20:00:00:11:0d:97:00:00          scsi-fcp:init
0xee0020    N     20:01:00:11:0d:97:01:00          scsi-fcp:target
0xee0400    N     10:00:00:de:fb:74:e8:31 (Cisco) ipfc
Total number of entries = 2
```

3. Create zones and a zoneset on VSAN 1 automatically and schedule a timer to check for new device logins on the Cisco MDS switch using the **autozone --enable** command:

```
switch# autozone --enable
This command will create and activate single-initiator and single-target zones for all
end-devices are already logged-in automatically; that may lead to more tcam entries and
 also RSCN load on network. Please use AutoZone judiciously.
AutoZone feature is enabled
Device with pwwn 10:00:00:de:fb:74:e8:31 is not registered with FC4-type Init or Target.
 Hence, it will be ignored for AutoZone configuration.
Configuring zones for vsan 1
                AUTOZONE_JPG21190082_1
Configuring zoneset for vsan 1
Activating the zoneset. Please wait...
```

```
Configured zoneset AUTOZONESET for vsan 1 successfully.
```

4. Verify the zone configuration using the **show zoneset active vsan 1** command:

```
switch# show zoneset active vsan 1
zoneset name AUTOZONESET vsan 1
  zone name AUTOZONE_JPG21190082_1 vsan 1
  * fcid 0xee0000 [pwwn 20:00:00:11:0d:97:00:00]
  * fcid 0xee0020 [pwwn 20:01:00:11:0d:97:01:00]
```

You can see that a new zoneset named *AUTOZONESET* is created and a new zone of the format *AUTOZONE_<SwitchSerialNumber>_<number>* is created and the devices are added to this zoneset:

5. Add Initiator 2 and Target 2 to the network:
6. Verify the new device logins using the **show fcns database** command:

```
switch# show fcns database
VSAN 1:
------------------------------------------------------------------------
FCID        TYPE  PWWN            (VENDOR)        FC4-TYPE:FEATURE
------------------------------------------------------------------------
0xee0000    N     20:00:00:11:0d:97:00:00         scsi-fcp:init
0xee0001    N     20:00:00:11:0d:97:00:01         scsi-fcp:init
0xee0020    N     20:01:00:11:0d:97:01:00         scsi-fcp:target
0xee0021    N     20:01:00:11:0d:97:01:01         scsi-fcp:target
0xee0400    N     10:00:00:de:fb:74:e8:31 (Cisco) ipfc
Total number of entries = 5
```

7. Verify the pending zone configurations using the **autozone --showpending** command:

```
switch# autozone --showpending
Feature AutoZone : Enabled
zoneset name AUTOZONESET vsan 1
        zone name AUTOZONE_JPG21190082_2 vsan 1
                 member pwwn 20:00:00:11:0d:97:00:01
                 member pwwn 20:01:00:11:0d:97:01:00
        zone name AUTOZONE_JPG21190082_3 vsan 1
                 member pwwn 20:00:00:11:0d:97:00:00
                 member pwwn 20:01:00:11:0d:97:01:01
        zone name AUTOZONE_JPG21190082_4 vsan 1
                 member pwwn 20:00:00:11:0d:97:00:01
                 member pwwn 20:01:00:11:0d:97:01:01
```

8. (Optional) You can execute Autozone manually for updating the zoning information each time a new device is logged into the switch using the **autozone --update** command:

```
switch# autozone --update
Device with pwwn 10:00:00:de:fb:74:e8:31 is not registered with FC4-type Init or Target.

Hence, it will be ignored for AutoZone configuration.
Configuring zones for vsan 1
                AUTOZONE_JPG21190082_1
                AUTOZONE_JPG21190082_2
                AUTOZONE_JPG21190082_3
                AUTOZONE_JPG21190082_4
Configuring zoneset for vsan 1
```

```
Activating the zoneset. Please wait...
Configured zoneset AUTOZONESET for vsan 1 successfully.
```

9. Verify the zone configuration for the new devices using the **autozone --show** command:

```
switch# autozone --show
Feature AutoZone : Enabled
AutoSave Configuration : Enabled
The possible zone/zoneset configuration with AutoZone feature for currently logged-in
devices is :
zoneset name AUTOZONESET vsan 1
        zone name AUTOZONE_JPG21190082_1 vsan 1
                member pwwn 20:00:00:11:0d:97:00:00
                member pwwn 20:01:00:11:0d:97:01:00
       zone name AUTOZONE_JPG21190082_2 vsan 1
                member pwwn 20:00:00:11:0d:97:00:01
                member pwwn 20:01:00:11:0d:97:01:00
        zone name AUTOZONE_JPG21190082_3 vsan 1
                member pwwn 20:00:00:11:0d:97:00:00
                member pwwn 20:01:00:11:0d:97:01:01
        zone name AUTOZONE_JPG21190082_4 vsan 1
                member pwwn 20:00:00:11:0d:97:00:01
                member pwwn 20:01:00:11:0d:97:01:01
```

Zone Configuration

Configuring a Zone

Tip Use a relevant display command (for example, **show interface** or **show flogi database**) to obtain the required value in hex format.

Tip Use the **show wwn switch** command to retrieve the sWWN. If you do not provide a sWWN, the software automatically uses the local sWWN.

Tip Expand Switches from the Physical Attributes pane to retrieve the sWWN. If you do not provide a sWWN, the software automatically uses the local sWWN.

Note Interface-based zoning only works with Cisco MDS 9000 Series switches. Interface-based zoning does not work if interop mode is configured in that VSAN.

When the number of zones configured has exceeded the maximum number of zones allowed across all VSANs, this message is displayed:

```
switch(config)# zone name temp_zone1 vsan 300
cannot create the zone; maximum possible number of zones is already configured
```

Note For configuration limits on configuring the number of zones, zone members and zone sets, refer to the Cisco MDS NX-OS Configuration Limits.

To configure a zone and assign a zone name, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **zone name Zone1 vsan 3**

Example:

```
switch(config-zone)#
```

Configures a zone called Zone1 for the VSAN called vsan3.

Note All alphanumeric characters or one of the following symbols ($, -, ^, _) are supported.

Step 3 switch(config-zone)# **member** *type value*

Example:

```
pWWN example:
```

Example:

```
switch(config-zone)# member pwwn 10:00:00:23:45:67:89:ab
```

Example:

```
Fabric pWWN example:
```

Example:

```
switch(config-zone)# member fwwn 10:01:10:01:10:ab:cd:ef
```

Example:

```
FC ID example:
```

Example:

```
switch(config-zone)# member fcid 0xce00d1
```

Example:

```
FC alias example:
```

Example:

```
switch(config-zone)# member fcalias Payroll
```

Example:

```
Domain ID example:
```

Example:

```
switch(config-zone)# member domain-id 2 portnumber 23
```

Example:

```
IPv4 address example:
```

Example:

```
switch(config-zone)# member ip-address 10.15.0.0 255.255.0.0
```

Example:

```
IPv6 address example:
```

Example:

```
switch(config-zone)# member ipv6-address 2001::db8:800:200c:417a/64
```

Example:

```
Local sWWN interface example:
```

Example:

```
switch(config-zone)# member interface fc 2/1
```

Example:

```
Remote sWWN interface example:
```

Example:

```
switch(config-zone)# member interface fc2/1 swwn 20:00:00:05:30:00:4a:de
```

Example:

```
Domain ID interface example:
```

Example:

```
switch(config-zone)# member interface fc2/1 domain-id 25
```

Example:

```
switch(config-zone)# member symbolic-nodename iqn.test
```

Configures a member for the specified zone (Zone1) based on the type (pWWN, fabric pWWN, FC ID, fcalias, domain ID, IPv4 address, IPv6 address, or interface) and value specified.

Caution You must only configure pWWN-type zoning on all MDS switches running Cisco SAN-OS if there is a Cisco MDS 9020 switch running FabricWare in the same fabric.

Note The Cisco MDS 9396S switch has 96 ports and the other Cisco MDS switches have lower ranges. Therefore while configuring a zone member based on interface type always select a fabric switch which potentially has the highest interface count in the fabric.

Filtering End Devices Based on Name, WWN or FC ID

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

Procedure

Step 1 Click the Zone icon in the toolbar (see Zone Icon).
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

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

Procedure

Step 1 Click the Zone icon in the toolbar (see Zone Icon).
Step 2 Use the Ctrl key to select multiple end devices.
Step 3 Right-click and then select **Add to Zone**.
Step 4 Use 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.

Zone Sets and FC Aliases

Zones provide a method for specifying access control, while 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 zoneset is in a configured VSAN).

Zoneset Distribution—You can distribute full zone sets using one of two methods: one-time distribution or full zoneset distribution.

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

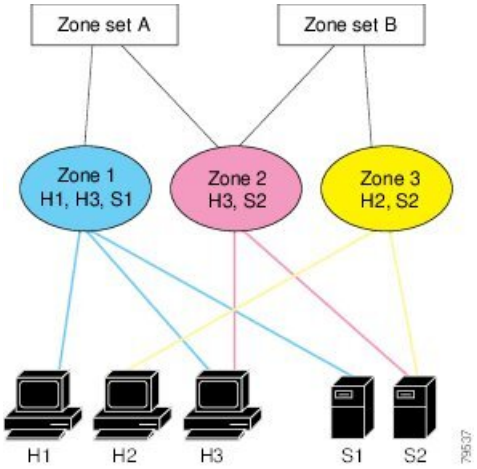
- To the full zoneset
- To a remote location (using FTP, SCP, SFTP, or TFTP)

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

ZoneSet Creation

In the figure, two separate sets are created, each with its own membership hierarchy and zone members.

Figure 9: Hierarchy of ZoneSets, Zones, and Zone Members



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

Tip Zonesets are configured with the names of the member zones and the VSAN (if the zoneset is in a configured VSAN).

Activating a Zoneset

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

Tip You do not have to issue the **copy running-config startup-config** command to store the active zoneset. However, you need to issue the **copy running-config startup-config** command to explicitly store full zone sets. If there is more than one switch in a fabric, the **copy running-config startup-config fabric** command should be issued. The **fabric** keyword causes the **copy running-config startup-config** command to be issued on all the switches in the fabric, and also saves the full zone information to the startup-config on all the switches in the fabric. This is important in the event of a switch reload or power cycle.

To activate or deactivate an existing zoneset, follow these steps:

Procedure

Step 1 switch# **config terminal**

Example:

```
switch(config)#
```

Enters configuration mode.

Step 2 switch(config)# **zoneset activate name Zoneset1 vsan 3**

Activates the specified zoneset.

If full zoneset distribution is configured for a VSAN, the zoneset activation also distributes the full zoning database to the other switches in the fabric.

If enhanced zoning is configured for a VSAN then the zoneset activation is held pending until the **zone commit vsan** *vsan-id* command is enabled. The **show zone pending-diff vsan** *vsan-id* displays the pending changes.

Note While activating a zoneset, if the zoneset overwrite-control vsan id command is enabled and the zoneset name is different from the current active zoneset, the activation will fail with an error message. For more information see Overwrite Control for an Active Zoneset, on page 62.

```
switch(config)# zoneset activate name Zoneset2 vsan 3

WARNING: You are trying to activate zoneset2, which is different from current active zoneset1. Do
you want to continue? (y/n) [n] y
```

Step 3 switch(config)# **no zoneset activate name Zoneset1 vsan 3**

Deactivates the specified zoneset.

Overwrite Control for an Active Zoneset

When activating a new zoneset, if users make a mistake while entering the zoneset name, and if this name already exists on the switch, it results in activation of the wrong zoneset and traffic loss. To avoid activating a wrong zoneset, the zoneset overwrite-control vsan id command is introduced.

Note Even when the zoneset overwrite-control vsan id command is enabled, the user can override it and continue with the activation of a new zoneset using the zoneset activate name zoneset name vsan *vsan* -id force command.

Procedure

Step 1 switch# **configure terminal**

Example:

```
switch(config)#
```

Enters configuration mode.

Step 2 switch(config)# **zoneset overwrite-control vsan 3**

Enables overwrite-control for the specified VSAN.

```
switch(config)# zoneset overwrite-control vsan 1

WARNING: This will enable Activation Overwrite control. Do you want to continue?
                                                                    (y/n) [n]
```

Note The zoneset overwrite-control vsan id command can be enabled only in enhanced zone mode.

Step 3 switch(config)# **show zone status vsan 3**

Displays the status of the VSAN, if overwrite-control is enabled or not.

What to do next

Displaying Zone Status

```
switch(config)# show zone status vsan 3
```

```
VSAN: 2 default-zone: deny distribute: full Interop: default
    mode: enhanced merge-control: allow
    session: none
    hard-zoning: enabled broadcast: unsupported
    smart-zoning: disabled
    rscn-format: fabric-address
    activation overwrite control: enabled
Default zone:
    qos: none broadcast: unsupported ronly: unsupported
Full Zoning Database :
    DB size: 348 bytes
    Zonesets:2  Zones:2 Aliases: 0 Attribute-groups: 1
Active Zoning Database :
    DB size: 68 bytes
    Name: hellset  Zonesets:1  Zones:1
Current Total Zone DB Usage: 416 / 2097152 bytes (0 % used)
Pending (Session) DB size:
    Full DB Copy size: 0 bytes
    Active DB Copy size: 0 bytes
SFC size: 0 / 2097152 bytes (0 % used)
Status: Commit completed at 15:19:49 UTC Jun 11 2015
```

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 zoneset 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 zoneset is active. When the default policy is configured as deny, the members of this zone are not explicitly enumerated when you issue the **show zoneset active** command.

Note The current default zoning policy is deny. The hidden active zoneset is d__efault__cfg in MDS. When there is a mismatch in default-zoning policies between two switches (permit on one side and deny on the other), zone merge will fail. The behavior is the same between two Brocade switches as well. The error messages will be as shown below.

The error messages will be as shown below:

Switch1 syslog:

switch(config-if)# 2014 Sep 2 06:33:21 hac15 %ZONE-2-ZS_MERGE_FAILED: %$VSAN 1%$ Zone merge failure, isolating interface fc2/10 received reason: Default zoning policy conflict. Received rjt from adjacent switch:[reason:0]

Switch2 syslog:

switch(config-if)# 2014 Sep 2 12:13:17 hac16 %ZONE-2-ZS_MERGE_FAILED: %$VSAN 1%$ Zone merge failure, isolating interface fc3/10 reason: Default zoning policy conflict.:[reason:0]

Configuring the Default Zone Access Permission

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

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **zone default-zone permit vsan 1**

Permits traffic flow to default zone members.

Step 3 switch(config)# **no zone default-zone permit vsan 1**

Denies (default) traffic flow to default zone members.

About FC Alias Creation

While the pWWN, fWWN, and so on of an end node or fabric port have to be specified to configure different features on a Cisco MDS switch, you must ensure to assign the correct value. An incorrect value, derived from a typo for example, may cause unexpected results. You can avoid this if you define a user-friendly name and use this name in all of the configuration commands, as required. These user-friendly names are referred to as *FC aliases* and they are defined according to naming conventions that are specific to each and every organization.

FC aliases are stored within the zone server database and the NX-OS software automatically converts FC aliases into their corresponding zone member types. A device alias name is a different type of alias and is

described in the Distributing Device Alias Services, on page 135 chapter. Device aliases can be assigned to FC aliases, but not vice-versa.

FC aliases are case sensitive and restricted to 64 alphanumeric characters. An FC alias name 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)

You can assign an FC alias name and configure an FC 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.
- 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.
- Device Alias—A device alias name is a different type of alias and it can be assigned as a member to a FC alias.

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

Creating FC Aliases

To create an alias, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **fcalias name AliasSample vsan 3**

```
switch(config-fcalias)#
```

Configures an alias name (AliasSample).

Step 3 switch(config-fcalias)# **member** *type value*

Configures a member for the specified fcalias (AliasSample) based on the type and value specified (pWWN, fabric pWWN, FC ID, domain ID, IPv4 address, IPv6 address, or interface).

```
Multiple members can be inserted for a single FC alias on multiple lines:
switch(config-fcalias)# member pwwn 10:00:00:23:45:67:89:ab
switch(config-fcalias)# member fwwn 10:01:10:01:10:ab:cd:ef
switch(config-fcalias)# member fcid 0x222222

pWWN example:
switch(config-fcalias)# member pwwn 10:00:00:23:45:67:89:ab

fWWN example:
switch(config-fcalias)# member fwwn 10:01:10:01:10:ab:cd:ef

FC ID example:
switch(config-fcalias)# member fcid 0x222222

Domain ID example:
switch(config-fcalias)# member domain-id 2 portnumber 23

IPv4 address example:
switch(config-fcalias)# member ip-address 10.15.0.0 255.255.0.0

IPv6 address example:
switch(config-fcalias)# member ipv6-address 2001::db8:800:200c:417a/64

Local sWWN interface example:
switch(config-fcalias)# member interface fc 2/1

Remote sWWN interface example:
switch(config-fcalias)# member interface fc2/1 swwn 20:00:00:05:30:00:4a:de

Domain ID interface example:
switch(config-fcalias)# member interface fc2/1 domain-id 25
```

Step 4 switch(config-fcalias)# **zone commit vsan** *id*

Commits the changes made to the specified VSAN.

Creating Zone Sets and Adding Member Zones

Tip You do not have to issue the **copy running-config startup-config** command to store the active zoneset. However, you need to issue the **copy running-config startup-config** command to explicitly store full zone sets. If there is more than one switch in a fabric, the **copy running-config startup-config fabric** command should be issued. The **fabric** keyword causes the **copy running-config startup-config** command to be issued on all the switches in the fabric, and also saves the full zone information to the startup-config on all the switches in the fabric. This is important in the event of a switch reload or power cycle.

Caution If you deactivate the active zoneset in a VSAN that is also configured for Inter-VSAN Routing (IVR), the active IVR zoneset (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 zoneset, check the active zone analysis for the VSAN (see the Zone and ZoneSet Analysis, on page 114). To reactivate the IVZS, you must reactivate the regular zoneset (refer to the Cisco MDS 9000 Series NX-OS Inter-VSAN Routing Configuration Guide).

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

To create a zoneset to include several zones, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **zoneset name Zoneset1 vsan 3**

Example:

```
switch(config-zoneset)#
```

Configures a zoneset called Zoneset1.

Tip To activate a zoneset, you must first create the zone and a zoneset.

Step 3 switch(config-zoneset)# **member Zone1**

Adds Zone1 as a member of the specified zoneset (Zoneset1).

Tip If the specified zone name was not previously configured, this command will return the Zone not present error message.

Step 4 switch(config-zoneset)# **zone name InlineZone1**

Example:

```
switch(config-zoneset-zone)#
```

Adds a zone (InlineZone1) to the specified zoneset (Zoneset1).

Tip Execute this step only if you need to create a zone from a zoneset prompt.

Step 5 switch(config-zoneset-zone)# **member fcid 0x111112**

Example:

```
switch(config-zoneset-zone)#
```

Adds a new member (FC ID 0x111112) to the new zone (InlineZone1).

Tip Execute this step only if you need to add a member to a zone from a zoneset prompt.

Adding Multiple Zones to Multiple Zone Sets

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

Procedure

Step 1 Click the Zone icon in the toolbar (see Zone Icon).
Step 2 From the tree view, select **Zoneset**.
Step 3 Use the Ctrl key to select multiple end devices.
Step 4 Right-click and then select **Add to Zoneset**.
Step 5 Use the Ctrl key to select multiple zones from the pop-up window displayed.
Step 6 Click **Add**.

Selected zones are added to the selected zone sets.

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 FCIDs 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 FCID 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 Series support both hard and soft zoning.

ZoneSet Distribution

You can distribute full zone sets using one of two methods: one-time distribution at the EXEC mode level or full zoneset distribution at the configuration mode level.

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

Table 6: Zone Set Distribution zoneset distribution Command Differences , on page 69 lists the differences between these distribution methods.

Table 6: Zone Set Distribution zoneset distribution Command Differences

One-Time Distribution zoneset distribute vsan Command (EXEC Mode)	Full Zone Set Distribution zoneset distribute full vsan Command (Configuration Mode)
Distributes the full zoneset immediately.	Does not distribute the full zoneset immediately.
Does not distribute the full zoneset information along with the active zoneset during activation, deactivation, or merge process.	Remembers to distribute the full zoneset information along with the active zoneset during activation, deactivation, and merge processes.



Tip You do not have to issue the **copy running-config startup-config** command to store the active zoneset. However, you need to issue the **copy running-config startup-config** command to explicitly store full zone sets. If there is more than one switch in a fabric, the **copy running-config startup-config fabric** command should be issued. The **fabric** keyword causes the **copy running-config startup-config** command to be issued on all the switches in the fabric, and also saves the full zone information to the startup-config on all the switches in the fabric. This is important in the event of a switch reload or power cycle.

Enabling Full Zoneset Distribution

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

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

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **zoneset distribute full vsan 33**

Enables sending a full zoneset along with an active zoneset.

Enabling a One-Time Distribution

Use the **zoneset distribute vsan** *vsan-id* command in EXEC mode to perform this distribution.

```
switch# zoneset distribute vsan 2
Zoneset distribution initiated. check zone status
```

This procedure command only distributes the full zoneset information; it does not save the information to the startup configuration. You must explicitly save the running configuration to the startup configuration issue the **copy running-config startup-config** command to save the full zoneset information to the startup configuration.

Note The **zoneset distribute vsan** *vsan-id* commandone-time distribution of the full zone set is supported in **interop 2** and **interop 3** modes, not in **interop 1** mode.

Use the **show zone status vsan** *vsan-id* command to check the status of the one-time zoneset distribution request.

```
switch# show zone status vsan 9
VSAN: 9 default-zone: deny distribute: full Interop: default
mode: enhanced merge-control: allow
session: none
hard-zoning: enabled broadcast: enabled
smart-zoning: disabled
rscn-format: fabric-address
activation overwrite control:disabled
Default zone:
qos: none broadcast: disabled ronly: disabled
Full Zoning Database :
DB size: 2002584 bytes
```

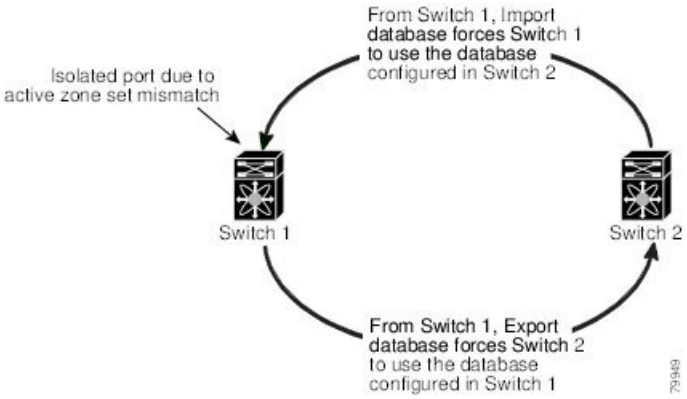
```
Zonesets:4 Zones:7004 Aliases: 0 Attribute-groups: 1
Active Zoning Database :
DB size: 94340 bytes
Name: zoneset-hac13-200 Zonesets:1 Zones:176
Current Total Zone DB Usage: 2096924 / 2097152 bytes (99 % used)
Pending (Session) DB size:
Full DB Copy size: 0 bytes
Active DB Copy size: 0 bytes
SFC size: 0 / 2097152 bytes (0 % used)
Status: Activation completed at 17:28:04 UTC Jun 16 2014
```

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 zoneset 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 zoneset database and replace the current active zoneset (see Figure 10: Importing and Exporting the Database , on page 71.
- Export the current database to the neighboring switch.
- Manually resolve the conflict by editing the full zoneset, activating the corrected zoneset, and then bringing up the link.

Figure 10: Importing and Exporting the Database



Importing and Exporting Zone Sets

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

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

Procedure

Step 1 switch# **zoneset import interface fc1/3 vsan 2**

Imports the zoneset from the adjacent switch connected through the fc 1/3 interface for VSAN 2.

Step 2 switch# **zoneset import interface fc1/3 vsan 2-5**

Imports the zoneset from the adjacent switch connected through the fc 1/3 interface for VSANs ranging from 2 through 5.

Step 3 switch# **zoneset export vsan 5**

Exports the zoneset to the adjacent switch connected through VSAN 5.

Step 4 switch# **zoneset export vsan 5-8**

Exports the zoneset to the adjacent switch connected through the range of VSANs 5 through 8.

Zoneset Duplication

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

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

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

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

Copying Zone Sets

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

Caution If the Inter-VSAN Routing (IVR) feature is enabled and if IVR zones exist in the active zoneset, then a zoneset 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 zoneset database before performing the copy operation. For more information on the IVR feature see the Cisco MDS 9000 Series NX-OS Inter-VSAN Routing Configuration Guide .

To make a copy of a zoneset, follow this step:

Procedure

Step 1 switch# **zone copy active-zoneset full-zoneset vsan 2**

Example:

```
Please enter yes to proceed.(y/n) [n]? y
```

Makes a copy of the active zoneset in VSAN 2 to the full zoneset.

Step 2 switch# **zone copy vsan 3 active-zoneset scp://guest@myserver/tmp/active_zoneset.txt**

Copies the active zone in VSAN 3 to a remote location using SCP.

Renaming Zones, Zone Sets, and Aliases

Note Backup option is available to switches that run Cisco NX-OS Release 4.1(3) or later.

To rename a zone, zone set, fcalias, or zone-attribute-group, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **zoneset rename oldname newname vsan 2**

Renames a zone set in the specified VSAN.

Step 3 switch(config)# **zone rename oldname newname vsan 2**

Renames a zone in the specified VSAN.

Step 4 switch(config)# **fcalias rename oldname newname vsan 2**

Renames a fcalias in the specified VSAN.

Step 5 switch(config)# **zone-attribute-group rename oldname newname vsan 2**

Renames a zone attribute group in the specified VSAN.

Step 6 switch(config)# **zoneset activate name newname vsan 2**

Activates the zone set and updates the new zone name in the active zone set.

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

To clone a zone, zoneset, fcalias, or zone-attribute-group, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **zoneset clone oldname newnamevsan 2**

Clones a zoneset in the specified VSAN.

Step 3 switch(config)# **zone clone oldname newname vsan 2**

Clones a zone in the specified VSAN.

Step 4 switch(config)# **fcalias clone oldname newnamevsan 2**

Clones a fcalias in the specified VSAN.

Step 5 switch(config)# **zone-attribute-group clone oldname newname vsan 2**

Clones a zone attribute group in the specified VSAN.

Step 6 switch(config)# **zoneset activate name newname vsan 2**

Activates the zoneset and updates the new zone name in the active zoneset.

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, use the following command:

```
switch# clear zone database vsan 2
```

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

Note After issuing a **clear zone database** command, you must explicitly issue the **copy running-config startup-config** to ensure that the running configuration is used when the switch reboots.

Note Clearing a zoneset 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.

Advanced Zone Attributes

About Zone-Based Traffic Priority

The zoning feature provides an additional segregation mechanism 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 Series Quality of Service Configuration Guide* for more information.

To use this feature, you need to obtain the ENTERPRISE_PKG license (refer to the *Cisco NX-OS Series Licensing Guide*) and you must enable QoS in the switch (refer to the *Cisco MDS 9000 Series NX-OS Quality of Service Configuration Guide*).

This feature allows SAN administrators to configure QoS in terms of 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.

Configuring Zone-Based Traffic Priority

To configure the zone priority, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **zone-attribute-group name priority vsan 2**

Example:

```
Switch(config-attribute-group)#
```

Configures a zone-attribute-group name (priority) and enters attribute configuration submode.

Step 3 Switch(config-attribute-group)# **qos priority high**

Example:

```
Switch(config-attribute-group)# qos priority high
Switch(config-attribute-group)# exit
```

Configures this zone-attribute to assign high priority QoS traffic to each frame matching

Step 4 switch(config-zone)# **attribute qos priority** {**high** | **low** | **medium**}

Configures this zone to assign QoS traffic to each frame matching this zone.

Step 5 Switch(config-zone)# **attribute-group priority**

Example:

```
Switch(config-zone)# attribute-group priority
Switch(config-zone)# exit
```

Configures this zone to assign high priority QoS traffic to each frame matching this zone in enhanced mode.

Step 6 switch(config)# **zoneset name QosZoneset vsan 2**

Example:

```
switch(config-zoneset)#
```

Configures a zoneset called QosZoneset for the specified VSAN (vsan 2) and enters zoneset configuration submode.

Tip To activate a zoneset, you must first create the zone and a zoneset.

Step 7 switch(config-zoneset)# **member QosZone**

Adds QosZone as a member of the specified zoneset (QosZoneset).

Tip If the specified zone name was not previously configured, this command will return the Zone not present error message.

Step 8 switch(config-zoneset)# **exit**

Example:

```
switch(config)#
```

Returns to configuration mode.

Step 9 switch(config)# **zoneset activate name QosZoneset vsan 2**

Activates the specified zoneset.

Disabling zone-based traffic priority

To disable zone-based QoS, the following specific sequence must be followed to ensure the configuration is applied correctly across the fabric:

Procedure

Step 1 Enter configuration mode.

Example:

```
switch# configure terminal
```

Step 2 Deactivate the active zoneset to prepare the fabric for the state change.

Example:

```
switch(config)# no zoneset activate name acl_zs vsan 1
```

Step 3 Disable the zone-based QoS feature. Note that while the feature status is updated in the configuration at this stage, the change is not yet active across the fabric.

Example:

```
switch(config)# zone name acl_zone vsan 1
switch(config-zone)# no attribute qos priority low
```

Step 4 Reactivate the zoneset. This step is critical as it acts as the "commit" operation that forces the updated zoning and QoS policies to distribute and synchronize throughout the fabric, ensuring that the changes take effect.

Example:

```
switch(config-zone)# zoneset activate name acl_zs vsan 1
```

Step 5 Verify the members of a specific zone within a designated VSAN:

Example:

```
switch# show zone name acl_zone vsan 1
zone name acl_zone vsan 1
  pwwn 20:00:00:00:00:00:ff:44
  pwwn 20:00:00:00:00:00:ff:45
```

Configuring Default Zone QoS Priority Attributes

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

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.

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

Procedure

Step 1 switch# **configure terminal**

Example:

```
switch(config)#
```

Enters configuration mode.

Step 2 switch(config)# **zone default-zone vsan 1**

Example:

```
switch(config-default-zone)#
```

Enters the default zone configuration submode.

Step 3 switch(config-default-zone)# **attribute qos priority high**

Sets the QoS priority attribute for frames matching these zones.

Step 4 switch(config-default-zone)# **no attribute qos priority high**

Removes the QoS priority attribute for the default zone and reverts to default low priority.

About Smart Zoning

Smart zoning implements hard zoning of large zones with fewer hardware resources than was previously required. The traditional zoning method allows each device in a zone to communicate with every other device in the zone. The administrator is required to manage the individual zones according to the zone configuration guidelines. Smart zoning eliminates the need to create a single initiator to single target zones. By analyzing device-type information in the FCNS, useful combinations can be implemented at the hardware level by the Cisco MDS NX-OS software, and the combinations that are not used are ignored. For example, initiator-target pairs are configured, but not initiator-initiator. The device is treated as unknown if:

- The FC4 types are not registered on the device.
- During Zone Convert, the device is not logged into the fabric.
- The zone is created, however, initiator, target, or initiator and target is not specified.

The device type information of each device in a smart zone is automatically populated from the Fibre Channel Name Server (FCNS) database as host, target, or both. This information allows more efficient utilisation of switch hardware by identifying initiator-target pairs and configuring those only in hardware. In the event of a special situation, such as a disk controller that needs to communicate with another disk controller, smart zoning defaults can be overridden by the administrator to allow complete control.

Note

- Smart Zoning can be enabled at VSAN level but can also be disabled at zone level.
- Smart zoning is not supported on VSANs that have DMM, IOA, or SME applications enabled on them.

Smart Zoning Member Configuration

Table displays the supported smart zoning member configurations.

Table 7: Smart Zoning Configuration

Feature	Supported
PWWN	Yes
FCID	Yes
FCalias	Yes
Device-alias	Yes
Interface	No
IP address	No
Symbolic nodename	No
FWWN	No
Domain ID	No

Enabling Smart Zoning on a VSAN

To configure the **smart zoning** for a VSAN, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **zone smart-zoning enable vsan 1**

Enables smart zoning on a VSAN.

Step 3 switch(config)# no **zone smart-zoning enable vsan 1**

Disables smart zoning on a VSAN.

Setting Default Value for Smart Zoning

To set the default value, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# system default zone smart-zone enable

Enables smart zoning on a VSAN that are created based on the specified default value.

Step 3 switch(config)# no system default zone smart-zone enable

Disables smart zoning on a VSAN.

Converting Zones Automatically to Smart Zoning

To fetch the device-type information from nameserver and to add that information to the member, follow the steps below: This can be performed at zone, zoneset, FCalias, and VSAN levels. After the zoneset is converted to smart zoning, you need to activate zoneset.

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# zone convert smart-zoning fcalias name <alias-name> vsan <vsan no>

Fetches the device type information from the nameserver for the fcalias members.



Note When the zone convert command is run, the FC4-Type should be SCSI-FCP. The SCSI-FCP has bits which determines whether the device is an initiator or target. If initiator and target are both set, the device is treated as both.

Step 3 switch(config)# zone convert smart-zoning zone name <zone name> vsan <vsan no>

Fetches the device type information from the nameserver for the zone members.

Step 4 switch(config)# zone convert smart-zoning zoneset name <zoneset name> vsan <vsan no>

Fetches the device type information from the nameserver for all the zones and fcalias members in the specified zoneset.

Step 5 switch(config)# zone convert smart-zoning vsan <vsan no>

Fetches the device type information from the nameserver for all the zones and fcalias members for all the zonesets present in the VSAN.

Step 6 switch(config)# show zone smart-zoning auto-conv status vsan 1

Displays the previous auto-convert status for a VSAN.

Step 7 switch(config)# show zone smart-zoning auto-conv log errors

Displays the error-logs for smart-zoning auto-convert.

What to do next

Use the show fcns database command to check if the device is initiator, target or both:

```
switch# show fcns database
VSAN 1:
------------------------------------------------------------------------
FCID TYPE PWWN (VENDOR) FC4-TYPE:FEATURE
------------------------------------------------------------------------
0x9c0000 N 21:00:00:e0:8b:08:96:22 (Company 1) scsi-fcp:init
0x9c0100 N 10:00:00:05:30:00:59:1f (Company 2) ipfc
0x9c0200 N 21:00:00:e0:8b:07:91:36 (Company 3) scsi-fcp:init
0x9c03d6 NL 21:00:00:20:37:46:78:97 (Company 4) scsi-fcp:target
```

Configuring Device Types for Zone Members

Note When device types are explicitly configured in smart zoning, any device must be configured with the same type in all zones of which the device is a member. A zone member must not be configured as initiator in some zones and target in other zones.

To configure the device types for zone members, follow these step:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config-zoneset-zone)# **member device-alias** *name* **both**

Configures the device type for the device-alias member as both. For every supported member-type, init, target, and both are supported.

Step 3 switch(config-zoneset-zone)# **member pwwn** *number* **target**

Configures the device type for the pwwn member as target. For every supported member-type, init, target, and both are supported.

Step 4 switch(config-zoneset-zone)# **member fcid** *number*

Configures the device type for the FCID member. There is no specific device type that is configured. For every supported member-type, init, target, and both are supported.



Note When there is no specific device type configured for a zone member, at the backend, zone entries that are generated are created as device type both.

Removing Smart Zoning Configuration

To remove the smart zoning configuration, follow this steps:

Procedure

Step 1 switch(config)# **clear zone smart-zoning fcalias name** *alias-name* **vsan** *number*

Removes the device type configuration for all the members of the specified fcalias.

Step 2 switch(config)# **clear zone smart-zoning zone name** *zone name* **vsan** *number*

Removes the device type configuration for all the members of the specified zone.

Step 3 switch(config)# **clear zone smart-zoning zoneset name** *zoneset name* **vsan** *number*

Removes the device type configuration for all the members of the zone and fcalias for the specified zoneset.

Step 4 switch(config)# **clear zone smart-zoning vsan** *number*

Removes the device type configuration for all the members of the zone and fcalias of all the specified zonesets in the VSAN.

Disabling Smart Zoning at Zone Level in the Basic Zoning Mode

To disable smart zoning at the zone level for a VSAN in basic zoning mode, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **zone name zone1 vsan 1**

Configures a zone name.

Step 3 switch(config-zone)# **attribute disable-smart-zoning**

Disables Smart Zoning for the selected zone.



Note This command only disables the smart zoning for the selected zone and does not remove the device type configurations.

Disabling Smart Zoning at Zone Level for a VSAN in the Enhanced Zoning Mode

To disable smart zoning at the zone level for a VSAN in enhanced zoning mode, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **zone-attribute-group name disable-sz vsan 1**

Creates an enhanced zone session.

Step 3 switch(config-attribute-group)#**disable-smart-zoning**

Disables Smart Zoning for the selected zone.



Note This command only disables the smart zoning for the selected zone and does not remove the device type configurations.

Step 4 switch(config-attribute-group)# **zone name prod vsan 1**

Configures a zone name.

Step 5 switch(config-zone)# **attribute-group disable-sz**

Configures to assign a group-attribute name for the selected zone.

Step 6 switch(config-zone)# **zone commit vsan 1**

Commits zoning changes to the selected VSAN.

Displaying Zone Information

You can view any zone information by using the **show** command. If you request information for a specific object (for example, a specific zone, zoneset, VSAN, or alias, or keywords such as **brief** or **active**), only information for the specified object is displayed. If you do not request specific information, all available information is displayed.

Displays Zone Information for All VSANs

```
switch# show zone
zone name Zone3 vsan 1
  pwwn 21:00:00:20:37:6f:db:dd
  pwwn 21:00:00:20:37:9c:48:e5
zone name Zone2 vsan 2
  fwwn 20:41:00:05:30:00:2a:1e
  fwwn 20:42:00:05:30:00:2a:1e
  fwwn 20:43:00:05:30:00:2a:1e
zone name Zone1 vsan 1
  pwwn 21:00:00:20:37:6f:db:dd
  pwwn 21:00:00:20:37:a6:be:2f
  pwwn 21:00:00:20:37:9c:48:e5
  fcalias Alias1
zone name Techdocs vsan 3
  ip-address 10.15.0.0 255.255.255.0
zone name Zone21 vsan 5
  pwwn 21:00:00:20:37:a6:be:35
  pwwn 21:00:00:20:37:a6:be:39
  fcid 0xe000ef
  fcid 0xe000e0
  symbolic-nodename iqn.test
  fwwn 20:1f:00:05:30:00:e5:c6
  fwwn 12:12:11:12:11:12:12:10
  interface fc1/5 swwn 20:00:00:05:30:00:2a:1e
  ip-address 12.2.4.5 255.255.255.0
  fcalias name Alias1 vsan 1
    pwwn 21:00:00:20:37:a6:be:35
zone name Zone2 vsan 11
  interface fc1/5 pwwn 20:4f:00:05:30:00:2a:1e
zone name Zone22 vsan 6
  fcalias name Alias1 vsan 1
    pwwn 21:00:00:20:37:a6:be:35
zone name Zone23 vsan 61
  pwwn 21:00:00:04:cf:fb:3e:7b lun 0000
```

Displays Zone Information for a Specific VSAN

```
switch# show zone vsan 1
zone name Zone3 vsan 1
  pwwn 21:00:00:20:37:6f:db:dd
  pwwn 21:00:00:20:37:9c:48:e5
zone name Zone2 vsan 1
    fwwn 20:4f:00:05:30:00:2a:1e
    fwwn 20:50:00:05:30:00:2a:1e
    fwwn 20:51:00:05:30:00:2a:1e
    fwwn 20:52:00:05:30:00:2a:1e
    fwwn 20:53:00:05:30:00:2a:1e
zone name Zone1 vsan 1
  pwwn 21:00:00:20:37:6f:db:dd
  pwwn 21:00:00:20:37:a6:be:2f
```

```
  pwwn 21:00:00:20:37:9c:48:e5
  fcalias Alias1
```

Use the **show zoneset** command to view the configured zonesets.

Displays Configured Zoneset Information

```
switch# show zoneset vsan 1
zoneset name ZoneSet2 vsan 1
  zone name Zone2 vsan 1
    fwwn 20:4e:00:05:30:00:2a:1e
    fwwn 20:4f:00:05:30:00:2a:1e
    fwwn 20:50:00:05:30:00:2a:1e
    fwwn 20:51:00:05:30:00:2a:1e
    fwwn 20:52:00:05:30:00:2a:1e
  zone name Zone1 vsan 1
    pwwn 21:00:00:20:37:6f:db:dd
    pwwn 21:00:00:20:37:a6:be:2f
    pwwn 21:00:00:20:37:9c:48:e5
    fcalias Alias1
zoneset name ZoneSet1 vsan 1
  zone name Zone1 vsan 1
    pwwn 21:00:00:20:37:6f:db:dd
    pwwn 21:00:00:20:37:a6:be:2f
    pwwn 21:00:00:20:37:9c:48:e5
    fcalias Alias1
```

Displays Configured Zoneset Information for a Range of VSANs

```
switch# show zoneset vsan 2-3
zoneset name ZoneSet2 vsan 2
  zone name Zone2 vsan 2
    fwwn 20:52:00:05:30:00:2a:1e
    fwwn 20:53:00:05:30:00:2a:1e
    fwwn 20:54:00:05:30:00:2a:1e
    fwwn 20:55:00:05:30:00:2a:1e
    fwwn 20:56:00:05:30:00:2a:1e
  zone name Zone1 vsan 2
    pwwn 21:00:00:20:37:6f:db:dd
    pwwn 21:00:00:20:37:a6:be:2f
    pwwn 21:00:00:20:37:9c:48:e5
    fcalias Alias1
zoneset name ZoneSet3 vsan 3
  zone name Zone1 vsan 1
    pwwn 21:00:00:20:37:6f:db:dd
    pwwn 21:00:00:20:37:a6:be:2f
    pwwn 21:00:00:20:37:9c:48:e5
    fcalias Alias1
```

Use the **show zone name** command to display members of a specific zone.

Displays Members of a Zone

```
switch# show zone name Zone1
zone name Zone1 vsan 1
  pwwn 21:00:00:20:37:6f:db:dd
  pwwn 21:00:00:20:37:a6:be:2f
  pwwn 21:00:00:20:37:9c:48:e5
  fcalias Alias1
```

Use the **show fcalias** command to display fcalias configuration.

Displays fcalias Configuration

```
switch# show fcalias vsan 1
fcalias name Alias2 vsan 1
fcalias name Alias1 vsan 1
  pwwn 21:00:00:20:37:6f:db:dd
  pwwn 21:00:00:20:37:9c:48:e5
```

Use the **show zone member** command to display all zones to which a member belongs using the FC ID.

Displays Membership Status

```
switch# show zone member pwwn 21:00:00:20:37:9c:48:e5
           VSAN: 1
zone Zone3
zone Zone1
fcalias Alias1
```

Use the **show zone statistics** command to display the number of control frames exchanged with other switches.

Displays Zone Statistics

```
switch# show zone statistics
Statistics For VSAN: 1
**********************************
Number of Merge Requests Sent: 24
Number of Merge Requests Recvd: 25
Number of Merge Accepts Sent: 25
Number of Merge Accepts Recvd: 25
Number of Merge Rejects Sent: 0
Number of Merge Rejects Recvd: 0
Number of Change Requests Sent: 0
Number of Change Requests Recvd: 0
Number of Change Rejects Sent: 0
Number of Change Rejects Recvd: 0
Number of GS Requests Recvd: 0
Number of GS Requests Rejected: 0
Statistics For VSAN: 2
**********************************
Number of Merge Requests Sent: 4
Number of Merge Requests Recvd: 4
Number of Merge Accepts Sent: 4
Number of Merge Accepts Recvd: 4
Number of Merge Rejects Sent: 0
Number of Merge Rejects Recvd: 0
Number of Change Requests Sent: 0
Number of Change Requests Recvd: 0
Number of Change Rejects Sent: 0
Number of Change Rejects Recvd: 0
Number of GS Requests Recvd: 0
Number of GS Requests Rejected: 0
```

Displays LUN Zone Statistics

```
switch# show zone statistics lun-zoning
LUN zoning statistics for VSAN: 1
**********************************************************
```

```
S-ID: 0x123456, D-ID: 0x22222, LUN: 00:00:00:00:00:00:00:00
----------------------------------------------------------
Number of Inquiry commands received:            10
Number of Inquiry data No LU sent:               5
Number of Report LUNs commands received:        10
Number of Request Sense commands received:       1
Number of Other commands received:               0
Number of Illegal Request Check Condition sent:  0
S-ID: 0x123456, D-ID: 0x22222, LUN: 00:00:00:00:00:00:00:01
----------------------------------------------------------
Number of Inquiry commands received:             1
Number of Inquiry data No LU sent:               1
Number of Request Sense commands received:       1
Number of Other commands received:               0
Number of Illegal Request Check Condition sent:  0
```

Displays LUN Zone Statistics

```
Need the latest output
switch# show zone statistics read-only-zoning
Read-only zoning statistics for VSAN: 2
**********************************************************
S-ID: 0x33333, D-ID: 0x11111, LUN: 00:00:00:00:00:00:00:64
----------------------------------------------------------
Number of Data Protect Check Condition Sent:    12
```

Displays Active Zone Sets

```
switch# show zoneset active
zoneset name ZoneSet1 vsan 1
  zone name zone1 vsan 1
    fcid 0x080808
    fcid 0x090909
    fcid 0x0a0a0a
  zone name zone2 vsan 1
  * fcid 0xef0000 [pwwn 21:00:00:20:37:6f:db:dd]
  * fcid 0xef0100 [pwwn 21:00:00:20:37:a6:be:2f]
```

Displays Brief Descriptions of Zone Sets

```
switch# show zoneset brief
zoneset name ZoneSet1 vsan 1
  zone zone1
  zone zone2
```

Displays Active Zones

```
switch# show zone active
zone name Zone2 vsan 1
* fcid 0x6c01ef [pwwn 21:00:00:20:37:9c:48:e5]
zone name IVRZ_IvrZone1 vsan 1
  pwwn 10:00:00:00:77:99:7a:1b
* fcid 0xce0000 [pwwn 10:00:00:00:c9:2d:5a:dd]
zone name IVRZ_IvrZone4 vsan 1
* fcid 0xce0000 [pwwn 10:00:00:00:c9:2d:5a:dd]
* fcid 0x6c01ef [pwwn 21:00:00:20:37:9c:48:e5]
zone name Zone1 vsan 1667
  fcid 0x123456
```

```
zone name $default_zone$ vsan 1667
```

Displays Active Zone Sets

```
switch# show zoneset active
zoneset name ZoneSet4 vsan 1
  zone name Zone2 vsan 1
  * fcid 0x6c01ef [pwwn 21:00:00:20:37:9c:48:e5]
  zone name IVRZ_IvrZone1 vsan 1
    pwwn 10:00:00:00:77:99:7a:1b
  * fcid 0xce0000 [pwwn 10:00:00:00:c9:2d:5a:dd]
zoneset name QosZoneset vsan 2
  zone name QosZone vsan 2
  attribute qos priority high
  * fcid 0xce0000 [pwwn 10:00:00:00:c9:2d:5a:dd]
  * fcid 0x6c01ef [pwwn 21:00:00:20:37:9c:48:e5]
Active zoneset vsan 1667
  zone name Zone1 vsan 1667
    fcid 0x123456
  zone name $default_zone$ vsan 1667
```

Displays Zone Status

```
switch(config)# show zone status
VSAN: 1 default-zone: deny distribute: active only Interop: default
mode: basic merge-control: allow
session: none
hard-zoning: enabled broadcast: disabled
smart-zoning: disabled
rscn-format: fabric-address
activation overwrite control:disabled
Default zone:
qos: none broadcast: disabled ronly: disabled
Full Zoning Database :
DB size: 4 bytes
Zonesets:0 Zones:0 Aliases: 0
Active Zoning Database :
Database Not Available
Current Total Zone DB Usage: 4 / 2097152 bytes (0 % used)
Pending (Session) DB size:
Full DB Copy size: n/a
Active DB Copy size: n/a
SFC size: 4 / 2097152 bytes (0 % used)
Status:
VSAN: 8 default-zone: deny distribute: full Interop: default
mode: basic merge-control: allow
session: none
hard-zoning: enabled broadcast: disabled
smart-zoning: disabled
rscn-format: fabric-address
Default zone:
qos: none broadcast: disabled ronly: disabled
Full Zoning Database :
DB size: 1946498 bytes
Zonesets:6 Zones:8024 Aliases: 0
Active Zoning Database :
DB size: 150499 bytes
Name: zoneset-1000 Zonesets:1 Zones:731
Current Total Zone DB Usage: 2096997 / 2097152 bytes (99 % used)
Pending (Session) DB size:
Full DB Copy size: n/a
```

```
Active DB Copy size: n/a
SFC size: 2096997 / 2097152 bytes (99 % used)
Status: Zoneset distribution failed [Error: Fabric changing Dom 33]:
at 17:05:06 UTC Jun 16 2014
VSAN: 9 default-zone: deny distribute: full Interop: default
mode: enhanced merge-control: allow
session: none
hard-zoning: enabled broadcast: enabled
smart-zoning: disabled
rscn-format: fabric-address
Default zone:
qos: none broadcast: disabled ronly: disabled
Full Zoning Database :
DB size: 2002584 bytes
Zonesets:4 Zones:7004 Aliases: 0 Attribute-groups: 1
Active Zoning Database :
DB size: 94340 bytes
Name: zoneset-hac13-200 Zonesets:1 Zones:176
Current Total Zone DB Usage: 2096924 / 2097152 bytes (99 % used)
Pending (Session) DB size:
Full DB Copy size: 0 bytes
Active DB Copy size: 0 bytes
SFC size: 0 / 2097152 bytes (0 % used)
Status: Activation completed at 17:28:04 UTC Jun 16 2014
VSAN: 12 default-zone: deny distribute: full Interop: default
mode: enhanced merge-control: allow
session: none
hard-zoning: enabled broadcast: enabled
smart-zoning: disabled
rscn-format: fabric-address
Default zone:
qos: none broadcast: disabled ronly: disabled
Full Zoning Database :
DB size: 84 bytes
Zonesets:0 Zones:1 Aliases: 0 Attribute-groups: 1
Active Zoning Database :
DB size: 144 bytes
Name: zs1 Zonesets:1 Zones:2
Current Total Zone DB Usage: 228 / 2097152 bytes (0 % used)
Pending (Session) DB size:
Full DB Copy size: 0 bytes
Active DB Copy size: 0 bytes
SFC size: 0 / 2097152 bytes (0 % used)
Status: Commit completed at 14:39:33 UTC Jun 27 201
```

Use the **show zone** command to display the zone attributes for all configured zones.

Displays Zone Statistics

```
switch# show zone
zone name lunSample vsan 1              <-----------------Read-write attribute
zone name ReadOnlyZone vsan 2
    attribute read-only                 <-----------------Read-only attribute
```

Use the **show running** and **show zone active** commands to display the configured interface-based zones.

Displays the Interface-Based Zones

```
switch# show running zone name if-zone vsan 1
      member interface fc2/15 swwn 20:00:00:0c:88:00:4a:e2
      member fwwn 20:4f:00:0c:88:00:4a:e2
```

```
      member interface fc2/1 swwn 20:00:00:05:30:00:4a:9e
      member pwwn 22:00:00:20:37:39:6b:dd
```

Displays the fWWNs and Interfaces in an Active Zone

```
switch# show zone active zone name if-zone vsan 1
  * fcid 0x7e00b3 [interface fc2/15 swwn 20:00:00:0c:88:00:4a:e2]
  * fcid 0x7e00b1 [interface fc2/15 swwn 20:00:00:0c:88:00:4a:e2]
  * fcid 0x7e00ac [interface fc2/15 swwn 20:00:00:0c:88:00:4a:e2]
  * fcid 0x7e00b3 [fwwn 20:4f:00:0c:88:00:4a:e2]
  * fcid 0x7e00b1 [fwwn 20:4f:00:0c:88:00:4a:e2]
  * fcid 0x7e00ac [fwwn 20:4f:00:0c:88:00:4a:e2]
    interface fc2/1 swwn 20:00:00:05:30:00:4a:9e
```

A similar output is also available on the remote switch (see the following example).

Displays the Local Interface Active Zone Details for a Remote Switch

```
switch# show zone active zone name if-zone vsan 1
   * fcid 0x7e00b3 [interface fc2/15 swwn 20:00:00:0c:88:00:4a:e2]
  * fcid 0x7e00b1 [interface fc2/15 swwn 20:00:00:0c:88:00:4a:e2]
  * fcid 0x7e00ac [interface fc2/15 swwn 20:00:00:0c:88:00:4a:e2]
  * fcid 0x7e00b3 [fwwn 20:4f:00:0c:88:00:4a:e2]
  * fcid 0x7e00b1 [fwwn 20:4f:00:0c:88:00:4a:e2]
  * fcid 0x7e00ac [fwwn 20:4f:00:0c:88:00:4a:e2]
    interface fc2/1 swwn 20:00:00:05:30:00:4a:9e
```

Displays the Zone Status for a VSAN

```
switch(config)# show zone status vsan 1
VSAN: 1 default-zone: deny distribute: active only Interop: default
mode: basic merge-control: allow
session: none
hard-zoning: enabled broadcast: disabled
smart-zoning: disabled
rscn-format: fabric-address
activation overwrite control:disabled
Default zone:
qos: none broadcast: disabled ronly: disabled
Full Zoning Database :
DB size: 4 bytes
Zonesets:0 Zones:0 Aliases: 0
Active Zoning Database :
Database Not Available
Current Total Zone DB Usage: 4 / 2097152 bytes (0 % used)
Pending (Session) DB size:
Full DB Copy size: n/a
Active DB Copy size: n/a
SFC size: 4 / 2097152 bytes (0 % used)
Status:
```

Displays the Zone Policy for a VSAN

```
switch# show zone policy vsan 1
Vsan: 1
   Default-zone: deny
   Distribute: full
   Broadcast: enable
   Merge control: allow
```

```
   Generic Service: read-write
   Smart-zone: enabled
```

Displays How to Create a Zone Attribute-Group to for a VSAN in the Enhanced Mode to Disable Smart Zoning at an Individual Zone Level

Note After the attribute-group is created, it needs to be applied to any zones requiring smart zoning to be disabled.

```
config# zone-attribute-group name <name> vsan 1
config-attribute-group# disable-smart-zoning
config-attribute-group# exit
config# zone commit vsan 1
```

Displays how to Auto-convert Zones

```
config# show zoneset vsan 1
zoneset name ZSv1 vsan 1
  zone name ddasZone vsan 1
    device-alias Init1
    device-alias Init2
    device-alias Init3
    device-alias Target1
config# zone convert smart-zoning vsan 1
smart-zoning auto_convert initiated. This operation can take few minutes. Please wait..
config# show zoneset vsan1
zoneset name ZSv1 vsan 1
  zone name ddasZone vsan 1
    device-alias Init1  init
    device-alias Init2  init
    device-alias Init3  init
    device-alias Target1 target
```

Displays how to Clear Device type Configuration for Members

```
config# show zoneset vsan 1
zoneset name ZSv1 vsan 1
  zone name ddasZone vsan 1
    device-alias Init1  init
    device-alias Init2  init
    device-alias Init3  init
    device-alias Target1 target
config# clear zone smart-zoning vsan1
config# show zoneset vsan 1
zoneset name ZSv1 vsan 1
  zone name ddasZone vsan 1
    device-alias Init1
    device-alias Init2
    device-alias Init3
    device-alias Target1
```

Enhanced Zoning

The zoning feature complies with the FC-GS-4 and FC-SW-3 standards. Both standards support the basic zoning functionalities explained in the previous section and the enhanced zoning functionalities described in this section.

About Enhanced Zoning

Table 8: Advantages of Enhanced Zoning , on page 92 lists the advantages of the enhanced zoning feature in all switches in the Cisco MDS 9000 Series.

Table 8: 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 zonesets, you create an instance of this zone in each zoneset.	References to the zone are used by the zonesets 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 zoneset. 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.

Changing from Basic Zoning to Enhanced Zoning

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

Procedure

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. By doing so, you will 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.

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

Procedure

Step 1 Verify that the active and full zoneset 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. By doing so, you will 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 on all switches in the Cisco MDS 9000 Series.

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

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **zone mode enhanced vsan** *id*

Enables enhanced zoning in the specified VSAN.

Step 3 switch(config)# **no zone mode enhanced vsan** *id*

Disables enhanced zoning in the specified VSAN.

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.

To commit or discard changes to the zoning database in a VSAN, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **zone commit vsan 2**

Applies the changes to the enhanced zone database and closes the session.

Step 3 switch(config)# **zone commit vsan 3 force**

Forcefully applies the changes to the enhanced zone database and closes the session created by another user.

Step 4 switch(config)# **no zone commit vsan 2**

Discards the changes to the enhanced zone database and closes the session.

Step 5 switch(config)# **no zone commit vsan 3 force**

Forcefully discards the changes to the enhanced zone database and closes the session created by another user.

Note You do not have to issue the **copy running-config startup-config** command to store the active zoneset. However, you need to issue the **copy running-config startup-config** command to explicitly store full zone sets. If there is more than one switch in a fabric, the **copy running-config startup-config fabric** command should be issued. The **fabric** keyword causes the **copy running-config startup-config** command to be issued on all the switches in the fabric, and also saves the full zone information to the startup-config on all the switches in the fabric. This is important in the event of a switch reload or power cycle.

Enabling Automatic Zone Pending Diff Display

To enable the display of pending-diff and subsequent confirmation on issuing a zone commit in enhanced mode, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# zone confirm-commit enable vsan vsan-id

Enables the confirm-commit option for zone database for a given VSAN.

Step 3 switch(config-zone)# zone commit vsan 12

If the zone confirm-commit command is enabled for a VSAN, on committing the pending database, the pending-diff is displayed on the console and the user is prompted for Yes or No. If the zone confirm-commit command is disabled, the pending-diff is not displayed and the user is not prompted for Yes or No.

Step 4 switch(config)# **no zone commit vsan 12**

If the zone confirm-commit command is enabled for a VSAN, on discarding the pending database, the pending-diff is displayed on the console and the user is prompted for Yes or No. If the zone confirm-commit command is disabled, the pending-diff is not displayed and the user is not prompted for Yes or No.

Releasing Zone Database Locks

To release the session lock on the zoning database on the switches in a VSAN, use the **no zone commit vsan** command from the switch where the database was initially locked.

```
switch# configure terminal
switch(config)# no zone commit vsan 2
```

If session locks remain on remote switches after using the **no zone commit vsan** command, you can use the **clear zone lock vsan** command on the remote switches.

```
switch# clear zone lock vsan 2
```

Note We recommend using the **no zone commit vsan** command first to release the session lock in the fabric. If that fails, use the **clear zone lock vsan** command on the remote switches where the session is still locked.

Creating Attribute Groups

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

To configure attribute groups, follow these steps:

Procedure

Step 1 Create an attribute group.

Example:

```
switch# confgure terminal
switch(config)# zone-attribute-group name SampleAttributeGroup vsan 2
switch(config-attribute-group)#
```

Step 2 Add the attribute to an attribute-group object.

Example:

```
switch(config-attribute-group)# readonly
switch(config-attribute-group)# broadcast
switch(config-attribute-group)# qos priority medium
readonly and broadcast commands are not supported from 5.2 release onwards.
```

Step 3 Attach the attribute-group to a zone.

Example:

```
switch(config)# zone name Zone1 vsan 2
switch(config-zone)# attribute-group SampleAttributeGroup
switch(config-zone)# exit
switch(config)#
```

Step 4 Activate the zoneset.

Example:

```
switch(config)# zoneset activate name Zoneset1 vsan 2
```

The attribute-groups are expanded and only the configured attributes are present in the active zoneset.

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

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 the Table 9: Database Zone Merge Status , on page 97.

Table 9: Database Zone Merge Status

Local Database	Adjacent Database	Merge Status	Results of the Merge
The databases contain zone sets with the same name 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.[1]	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 zoneset does not have a name in interop mode 1. The zoneset names are only present for full zone sets.

Merge Process

When two Fibre Channel (FC) switches that have already been configured with active zonesets and are not yet connected are brought together with an Extended ISL (EISL) link, the zonesets merge. However, steps must be taken to ensure zone consistency before configuring and activating new zones.

Best Practices

When a zone merge occurs, as long as there is not competing information, each switch learns the others zones. Each switch then has three configuration entities. The switches have:

- The saved configuration in NVRAM. This is the configuration as it was the last time the **copy running-configuration startup-configuration** command was issued.
- The running configuration. This represents the configuration brought into memory upon the last time the MDS was brought up, plus any changes that have been made to the configuration. With reference to the

zoning information, the running configuration represents the configurable database, known as the full database.

- The configured zoning information from the running configuration plus the zoning information learned from the zone merge. This combination of configured and learned zone information is the active zoneset.

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 zoneset and the full zoneset should be identical. Otherwise the link is isolated.
   b. If the setting is allow, then the merge rules are used to perform the merge.

When an MDS is booted, it comes up with the configuration previously saved in NVRAM. If you configured the switch after loading the configuration from NVRAM, there is a difference between the bootup and running configuration until the running configuration is saved to the startup configuration. This can be likened to having a file on the local hard drive of your PC. The file is saved and static, but if you open the file and edit, there exists a difference between the changed file and the file that still exists on saved storage. Only when you save the changes, does the saved entity look represent the changes made to the file.

When zoning information is learned from a zone merge, this learned information is not part of the running configuration. Only when the **zone copy active-zoneset full-zoneset vsan X** command is issued, the learned information becomes incorporated into the running configuration. This is key because when a zone merge is initiated by a new EISL link or activating a zoneset, the zoneset part is ignored by the other switch and the member zone information is considered topical.



Caution The **zone copy** command will delete all fcalias configuration.

Example

For example, you have two standalone MDS switches, already in place and each with their own configured zone and zoneset information. Switch 1 has an active zoneset known as set A, and Switch 2 has an active zoneset known as set B. Within set A on Switch 1 is zone 1, and on Switch 2, set B has member zone 2. When an ISL link is created between these two switches, each sends their zoneset including their zone information to the other switch. On a merge, the switch will select zoneset name with the higher ASCII value and then merge their zone member. After the merge, both switches will have a zoneset name set B with zone member zone 1 and zone 2.

Everything should be still working for all of the devices in zone 1 and zone 2. To add a new zone, you have to create a new zone, add the new zone to the zoneset, and then activate the zoneset.

Step-by-step, the switches are booted up and have no zoning information. You need to create the zones on the switches and add them to the zonesets.

Basic mode: When zones are in basic mode, refer to the sample command outputs below.

1. Create zone and zoneset. Activate on Switch 1.

```
Switch1# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.

Switch1#(config)# vsan database
Switch1#(config-vsan-db)# vsan 100
Switch1#(config-vsan-db)# exit

Switch1#(config)# zone name zone1 vsan 100
Switch1#(config-zone)# member pwwn 11:11:11:11:11:11:11:1a
Switch1#(config-zone)# member pwwn 11:11:11:11:11:11:11:1b
Switch1#(config-zone)# exit

Switch1#(config)# zoneset name setA vsan 100
Switch1#(config-zoneset)# member zone1
Switch1#(config-zoneset)# exit

Switch1#(config)# zoneset activate name setA vsan 100
Zoneset activation initiated. check zone status
Switch1#(config)# exit

Switch1# show zoneset active vsan 100
zoneset name setA vsan 100
zone name zone1 vsan 100
pwwn 11:11:11:11:11:11:11:1a
pwwn 11:11:11:11:11:11:11:1bSwitch1#
```

2. Create zone and zoneset. Activate on Switch 2.

```
Switch2# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.

Switch2#(config)# vsan database
Switch2#config-vsan-db)# vsan 100
Switch2#(config-vsan-db)# exit

Switch2#(config)# zone name zone2 vsan 100
Switch2#(config-zone)# member pwwn 22:22:22:22:22:22:22:2a
Switch2#(config-zone)# member pwwn 22:22:22:22:22:22:22:2b
Switch2#(config-zone)# exit

Switch2#(config)# zoneset name setB vsan 100
Switch2#(config-zoneset)# member zone2
Switch2#(config-zoneset)# exit

Switch2#(config)# zoneset activate name setB vsan 100
Zoneset activation initiated. check zone status
Switch2#(config)# exit

Switch2# show zoneset active vsan 100
zoneset name setB vsan 100
zone name zone2 vsan 100
pwwn 22:22:22:22:22:22:22:2a
pwwn 22:22:22:22:22:22:22:2b
```

3. Bring ISL link up and verify zone merge on Switch 1.

```
Switch1# configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
Switch1(config)# interface fc1/5
Switch1(config-if)# no shutdown
Switch1(config-if)# exit
Switch1(config)# exit
```

Note Note Ensure that vsan 100 is allowed on ISL.

```
Switch1# show zoneset active vsan 100
zoneset name setB vsan 100
zone name zone1 vsan 100
pwwn 11:11:11:11:11:11:11:1a
pwwn 11:11:11:11:11:11:11:1b

Switch1# show zoneset vsan 100
zoneset name setA vsan 100
zone name zone1 vsan 100
pwwn 11:11:11:11:11:11:11:1a
pwwn 11:11:11:11:11:11:11:1b
```

4. Bring ISL link up and verify zone merge on Switch 2.

```
Switch2# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch2(config)# int fc2/5
Switch2(config-if)# no shut
Switch2(config-if)# exit
Switch2(config)# exit

Switch2# show zoneset active vsan 100 zoneset name setB vsan 100
zone name zone1 vsan 100
pwwn 11:11:11:11:11:11:11:1a
pwwn 11:11:11:11:11:11:11:1b

zone name zone2 vsan 100
pwwn 22:22:22:22:22:22:22:2a
pwwn 22:22:22:22:22:22:22:2b

Switch2# show zoneset vsan 100zoneset name setB vsan 100
zone name zone2 vsan 100
pwwn 22:22:22:22:22:22:22:2a
pwwn 22:22:22:22:22:22:22:2b
```

Note The name of the newly merged zoneset will be the name of the zoneset with alphabetically higher value. In the given example, the active zoneset is setB. To avoid future zoneset activation problems, the **zone copy active-zoneset full-zoneset vsan** *100* command should be given, at this point on the switch. Examine if the command is given, and how the new zoning information is handled.

When the zone copy command is issued, it adds the learned zone information, zone 2 in this case, to the running configuration. If zone 2 has not been copied from residing in memory to copied into the running configuration, zone 2 information is not pushed back out.

Note The **zone copy** command will delete all fcalias configuration.

Running-Configuration of Switch1 (before issuing the **zone copy active-zoneset full-zoneset vsan** *100* command).

```
Switch1# show run | b "Active Zone Database Section for vsan 100"
!Active Zone Database Section for vsan 100
zone name zone1 vsan 100
pwwn 11:11:11:11:11:11:11:1a
pwwn 11:11:11:11:11:11:11:1b

zone name zone2 vsan 100
pwwn 22:22:22:22:22:22:22:2a
pwwn 22:22:22:22:22:22:22:2b

zoneset name setB vsan 100
member zone1
member zone2

zoneset activate name setB vsan 100
do clear zone database vsan 100
!Full Zone Database Section for vsan 100
zone name zone1 vsan 100
pwwn 11:11:11:11:11:11:11:1a
pwwn 11:11:11:11:11:11:11:1b

zoneset name setA vsan 100
member zone1
```

Running-Configuration of Switch1 (after issuing the "zone copy active-zoneset full-zoneset vsan 100" command)

```
Switch1# zone copy active-zoneset full-zoneset vsan 100
WARNING: This command may overwrite common zones in the full zoneset. Do you want to continue?
 (y/n) [n] y

Switch1# show run | b "Active Zone Database Section for vsan 100"
!Active Zone Database Section for vsan 100
zone name zone1 vsan 100
pwwn 11:11:11:11:11:11:11:1a
pwwn 11:11:11:11:11:11:11:1b

zone name zone2 vsan 100
pwwn 22:22:22:22:22:22:22:2a
pwwn 22:22:22:22:22:22:22:2b

zoneset name setB vsan 100
member zone1
member zone2

zoneset activate name setB vsan 100
do clear zone database vsan 100
!Full Zone Database Section for vsan 100
zone name zone1 vsan 100
pwwn 11:11:11:11:11:11:11:1a
pwwn 11:11:11:11:11:11:11:1b
```

```
zone name zone2 vsan 100
pwwn 22:22:22:22:22:22:22:2a
pwwn 22:22:22:22:22:22:22:2b

zoneset name setA vsan 100
member zone1

zoneset name setB vsan 100
member zone1
member zone2
```

Running-Configuration of Switch2 (before issuing the "zone copy active-zoneset full-zoneset vsan 100" command)

```
Switch2# show run | b "Active Zone Database Section for vsan 100"
!Active Zone Database Section for vsan 100
zone name zone2 vsan 100
pwwn 22:22:22:22:22:22:22:2a
pwwn 22:22:22:22:22:22:22:2b

zone name zone1 vsan 100
pwwn 11:11:11:11:11:11:11:1a
pwwn 11:11:11:11:11:11:11:1b

zoneset name setB vsan 100
member zone2
member zone1

zoneset activate name setB vsan 100
do clear zone database vsan 100
!Full Zone Database Section for vsan 100
zone name zone2 vsan 100
pwwn 22:22:22:22:22:22:22:2
apwwn 22:22:22:22:22:22:22:2b
zoneset name setB vsan 100
member zone2
```

Running-Configuration of Switch2 (after issuing the "zone copy active-zoneset full-zoneset vsan 100" command)

```
Switch2# zone copy active-zoneset full-zoneset vsan 100
WARNING: This command may overwrite common zones in the full zoneset. Do you want to continue?
 (y/n) [n] y

Switch2# show run | b "Active Zone Database Section for vsan 100"
!Active Zone Database Section for vsan 100
zone name zone2 vsan 100
pwwn 22:22:22:22:22:22:22:2a
pwwn 22:22:22:22:22:22:22:2b

zone name zone1 vsan 100
pwwn 11:11:11:11:11:11:11:1a
pwwn 11:11:11:11:11:11:11:1b

zoneset name setB vsan 100
member zone2
member zone1

zoneset activate name setB vsan 100
do clear zone database vsan 100
!Full Zone Database Section for vsan 100
```

```
zone name zone2 vsan 100
pwwn 22:22:22:22:22:22:22:2a
pwwn 22:22:22:22:22:22:22:2b

zone name zone1 vsan 100
pwwn 11:11:11:11:11:11:11:1a
pwwn 11:11:11:11:11:11:11:1b

zoneset name setB vsan 10
0member zone2
member zone1
```

Referring back to the three entities of configuration, they are as follows on zone 1 before the zone merge:

- Saved configuration: nothing since zone information has not been saved by issuing the copy run start command.
- Running configuration: consists of zone 1.
- Configured and learned information: consists of zone 1.

After the zone merge, the entities are:

- Saved configuration: nothing has been saved.
- Running configuration: consists of zone 1.
- Configured and learned information: consists of zone 1 and zone 2.

Zone 2 has not become part of the running configuration. Zone 2 has been learned, and is in the active zoneset. Only when the **zone copy active-zoneset full-zoneset vsan** *100* command is issued, zone 2 becomes copied from being learned to added to the running configuration. The configuration looks as follows after the command is issued:



Note The **zone copy** command will delete all fcalias configuration.

- Saved configuration: nothing has been saved.
- Running configuration: consists of zone 1 and zone 2.
- Configured and learned information: consists of zone 1 and zone 2.

Commands

By default zone in basic mode will only distribute active zoneset database only, this command was introduced in 1.0.4 SAN-OS will propagate active zoneset and full zoneset database:

zoneset distribute full vsan *vsan_id*

If the zone update or zoneset activation is going on, the above command must be explicitly enabled on each VSAN on every switch.

Enhanced mode: When zones are in enhanced mode, refer to the sample command outputs below.

1. Create zones and zoneset. Activate on Switch1.

```
Switch1# configure terminal
```

```
Enter configuration commands, one per line. End with CNTL/Z.
Switch1(config)# vsan database
Switch1(config-vsan-db)# vsan 200
Switch1(config-vsan-db)# zone mode enhanced vsan 200
WARNING: This command would distribute the zoning database of this switch throughout the
 fabric. Do you want to continue? (y/n) [n] y
Set zoning mode command initiated.
Check zone status
Switch1(config-vsan-db)# zone name zone1 vsan 200
Enhanced zone session has been created. Please 'commit' the changes when done.
Switch1(config-zone)# member pwwn 11:11:11:11:11:11:11:1a
Switch1(config-zone)# member pwwn 11:11:11:11:11:11:11:1b
Switch1(config-zone)# zoneset name SetA vsan 200
Switch1(config-zoneset)# member zone1
Switch1(config-zoneset)# zoneset activate name SetA vsan 200
Switch1(config)# zone commit vsan 200
Commit operation initiated. Check zone status
Switch1(config)# exit
Switch1# show zoneset activate vsan 200
zoneset name SetA vsan 200
zone name zone1 vsan 200
pwwn 11:11:11:11:11:11:11:1a
pwwn 11:11:11:11:11:11:11:1b
Switch1# show zoneset vsan 200
zoneset name SetA vsan 200
zone name zone1 vsan 200
pwwn 11:11:11:11:11:11:11:1a
pwwn 11:11:11:11:11:11:11:1b
```

2. Create zones and zoneset. Activate on Switch2.

```
Switch2# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch2(config)# vsan database
Switch2(config-vsan-db)# vsan 200
Switch2(config-vsan-db)# zone mode enhanced vsan 200
WARNING: This command would distribute the zoning database of this switch throughout the
 fabric. Do you want to continue? (y/n) [n] y
Set zoning mode command initiated. Check zone status
Switch2(config)# zone name zone2 vsan 200
Enhanced zone session has been created. Please 'commit' the changes when done.
Switch2(config-zone)# member pwwn 22:22:22:22:22:22:22:2a
Switch2(config-zone)# member pwwn 22:22:22:22:22:22:22:2b
Switch2(config-zone)# zoneset name SetB vsan 200
Switch2(config-zoneset)# member zone2
Switch2(config-zoneset)# zoneset act name SetB vsan 200
Switch2(config)# zone commit vsan 200
Commit operation initiated. Check zone status
Switch2(config)# exit
Switch2# show zoneset activate vsan 200
zoneset name SetB vsan 200
zone name zone2 vsan 200
pwwn 22:22:22:22:22:22:22:2a
pwwn 22:22:22:22:22:22:22:2b
Switch2# show zoneset vsan 200
zoneset name SetB vsan 200
zone name zone2 vsan 200
pwwn 22:22:22:22:22:22:22:2a
pwwn 22:22:22:22:22:22:22:2b
```

3. Bring ISL link up and verify zone merge on Switch1.

```
Switch1# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch1(config)# interface fc4/1
Switch1(config-if)# no shutdown
Switch1(config-if)# exit
Switch1(config)# exit

Switch1(config-if)# show zoneset activate vsan 200
zoneset name SetB vsan 200
zone name zone1 vsan 200
pwwn 11:11:11:11:11:11:11:1a
pwwn 11:11:11:11:11:11:11:1b
zone name zone2 vsan 200
pwwn 22:22:22:22:22:22:22:2a
pwwn 22:22:22:22:22:22:22:2b
Switch1(config-if)# show zoneset vsan 200
zoneset name SetA vsan 200
zone name zone1 vsan 200
pwwn 11:11:11:11:11:11:11:1a
pwwn 11:11:11:11:11:11:11:1b

zoneset name SetB vsan 200
zone name zone2 vsan 200
pwwn 22:22:22:22:22:22:22:2a
pwwn 22:22:22:22:22:22:22:2b
```



Note Unlike basic mode, the entire zone database is merged in the case of enhanced mode, wherein Switch1 has the information of zonesets originally configured in Switch2 and vice versa.

4. Bring ISL link up and verify zone merge on Switch2. After bringing up ISL between two switches:

```
Switch2# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch2(config)# interface fc4/1
Switch2(config-if)# no shutdown
Switch2(config-if)# exit
Switch2(config)# exit

Switch2(config-zoneset)# show zoneset activate vsan 200
zoneset name SetB vsan 200
zone name zone2 vsan 200
pwwn 22:22:22:22:22:22:22:2a
pwwn 22:22:22:22:22:22:22:2b
zone name zone1 vsan 200
pwwn 11:11:11:11:11:11:11:1a
pwwn 11:11:11:11:11:11:11:1b
Switch2(config-zoneset)# show zoneset vsan 200
zoneset name SetB vsan 200
zone name zone2 vsan 200
pwwn 22:22:22:22:22:22:22:2a
pwwn 22:22:22:22:22:22:22:2b

zoneset name SetA vsan 200
zone name zone1 vsan 200
pwwn 11:11:11:11:11:11:11:1a
pwwn 11:11:11:11:11:11:11:1b
```

5. Execute the **zone copy** command for enhanced zone.

   Switch 1

```
Switch1# zone copy active-zoneset full-zoneset vsan 200
WARNING: This command may overwrite common zones in the full zoneset. Do you want to
continue? (y/n) [n] y
Switch1(config-if)# show zoneset activate vsan 200
zoneset name SetB vsan 200
zone name zone1 vsan 200
pwwn 11:11:11:11:11:11:11:1a
pwwn 11:11:11:11:11:11:11:1b
zone name zone2 vsan 200
pwwn 22:22:22:22:22:22:22:2a
pwwn 22:22:22:22:22:22:22:2b
Switch1(config-if)# show zoneset vsan 200
zoneset name SetB vsan 200
zone name zone1 vsan 200
pwwn 11:11:11:11:11:11:11:1a
pwwn 11:11:11:11:11:11:11:1b
zone name zone2 vsan 200
pwwn 22:22:22:22:22:22:22:2a
pwwn 22:22:22:22:22:22:22:2b
```

   Switch 2

```
Switch2# zone copy active-zoneset full-zoneset vsan 200
WARNING: This command may overwrite common zones in the full zoneset. Do you want to
continue? (y/n) [n] y
Switch2(config-zoneset)# show zoneset activate vsan 200
zoneset name SetB vsan 200
zone name zone2 vsan 200
pwwn 22:22:22:22:22:22:22:2a
pwwn 22:22:22:22:22:22:22:2b
zone name zone1 vsan 200
pwwn 11:11:11:11:11:11:11:1a
pwwn 11:11:11:11:11:11:11:1b
Switch2(config-zoneset)# show zoneset vsan 200
zoneset name SetB vsan 200
zone name zone2 vsan 200
pwwn 22:22:22:22:22:22:22:2a
pwwn 22:22:22:22:22:22:22:2b
zone name zone1 vsan 200
pwwn 11:11:11:11:11:11:11:1a
pwwn 11:11:11:11:11:11:11:1b
```

Configuring Zone Merge Control Policies

To configure merge control policies, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **zone merge-control restrict vsan 4**

Configures a restricted merge control setting for this VSAN.

Step 3 switch(config)# **no zone merge-control restrict vsan 2**

Defaults to using the allow merge control setting for this VSAN.

Step 4 switch(config)# **zone commit vsan 4**

Commits the changes made to VSAN 4.

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

Preventing Zones From Flooding FC2 Buffers

By using the **zone fc2 merge throttle enable** command you can throttle the merge requests that are sent from zones to FC2 and prevent zones from flooding FC2 buffers. This command is enabled by default. This command can be used to prevent any zone merge scalability problem when you have a lot of zones. Use the **show zone status** command to view zone merge throttle information.

Permitting or Denying Traffic in the Default Zone

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

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **zone default-zone permit vsan 5**

Permits traffic flow to default zone members.

Step 3 switch(config)# **no zone default-zone permit vsan 3**

Denies traffic flow to default zone members and reverts to factory default.

Step 4 switch(config)# **zone commit vsan 5**

Commits the changes made to VSAN 5.

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.

Note broadcast command is not supported from 5.x release onwards.

Table 10: Broadcasting Requirements , on page 108 identifies the rules for the delivery of broadcast frames.

Table 10: Broadcasting Requirements

Active Zoning?	Broadcast Enabled?	Frames Broadcast?
Yes	Yes	Yes
No	Yes	Yes
Yes	No	No
Contains data.	Empty.	Successful.

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.

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

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **zone-attribute-group name BroadcastAttr vsan 2**

Configures the zone attribute group for the required VSAN.

Step 3 switch(config)# **no zone-attribute-group name BroadAttr vsan 1**

Removes the zone attribute group for the required VSAN.

Step 4 switch(config-attribute-group)# **broadcast**

Creates a broadcast attribute for this group and exits this submode.

Step 5 switch(config-attribute-group)# **no broadcast**

Removes broadcast attribute for this group and exits this submode.

Step 6 switch(config)# **zone name BroadcastAttr vsan 2**

Configures a zone named BroadcastAttr in VSAN 2.

Step 7 switch(config-zone)# **member pwwn 21:00:00:e0:8b:0b:66:56**

Adds the specified members to this zone and exits this submode.

Step 8 switch(config)# **zone commit vsan 1**

Applies the changes to the enhanced zone configuration and exits this submode.

Step 9 switch# **show zone vsan 1**

Displays the broadcast configuration

Configuring System Default Zoning Settings

You can configure default settings for default zone policies, full zone distribution, and generic service permissions for new VSANs on the switch. To configure switch-wide default settings, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **system default zone default-zone permit**

Configures permit as the default zoning policy for new VSANs on the switch.

Step 3 switch(config)# **system default zone distribute full**

Enables full zone database distribution as the default for new VSANs on the switch.

Step 4 switch(config)# **system default zone gs** {**read** | **read-write**}

Configures read only or read-write (default) as the default generic service permission for new VSANs on the switch.

Note Since VSAN 1 is the default VSAN and is always present on the switch, the **system default zone** commands have no effect on VSAN 1.

Configuring Zone Generic Service Permission Settings

Zone generic service permission setting is used to control zoning operation through generic service (GS) interface. The zone generic service permission can be read-only, read-write or none (deny).

To configure generic service (GS) settings, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **zone gs** {**read** | **read-write**} **vsan 3000**

Configures gs permission value as read only or read-write in the specified VSAN.

Displaying Enhanced Zone Information

You can view any zone information by using the **show** command.

Displays the Active Zoneset Information for a Specified VSAN

```
switch(config)# show zoneset active vsan 1
zoneset name qoscfg vsan 1
  zone name qos1 vsan 1
  * fcid 0xe80200 [pwwn 50:08:01:60:01:5d:51:11]
  * fcid 0xe60000 [pwwn 50:08:01:60:01:5d:51:10]
  * fcid 0xe80100 [pwwn 50:08:01:60:01:5d:51:13]

  zone name qos3 vsan 1
  * fcid 0xe80200 [pwwn 50:08:01:60:01:5d:51:11]
  * fcid 0xe60100 [pwwn 50:08:01:60:01:5d:51:12]
  * fcid 0xe80100 [pwwn 50:08:01:60:01:5d:51:13]

  zone name sb1 vsan 1
  * fcid 0xe80000 [pwwn 20:0e:00:11:0d:10:dc:00]
  * fcid 0xe80300 [pwwn 20:0d:00:11:0d:10:da:00]
  * fcid 0xe60200 [pwwn 20:13:00:11:0d:15:75:00]
  * fcid 0xe60300 [pwwn 20:0d:00:11:0d:10:db:00]
```

Displays the ZoneSet Information or a Specified VSAN

```
switch(config)# show zoneset vsan 1
zoneset name qoscfg vsan 1
  zone name qos1 vsan 1
    zone-attribute-group name qos1-attr-group vsan 1
    pwwn 50:08:01:60:01:5d:51:11
    pwwn 50:08:01:60:01:5d:51:10
    pwwn 50:08:01:60:01:5d:51:13

  zone name qos3 vsan 1
    zone-attribute-group name qos3-attr-group vsan 1
    pwwn 50:08:01:60:01:5d:51:11
    pwwn 50:08:01:60:01:5d:51:12
    pwwn 50:08:01:60:01:5d:51:13

  zone name sb1 vsan 1
    pwwn 20:0e:00:11:0d:10:dc:00
    pwwn 20:0d:00:11:0d:10:da:00
    pwwn 20:13:00:11:0d:15:75:00
```

```
    pwwn 20:0d:00:11:0d:10:db:00
```

Displays the Zone Attribute Group Information for a Specified VSAN

```
switch# show zone-attribute-group vsan 2
zone-attribute-group name $default_zone_attr_group$ vsan 2
  read-only
  qos priority high
  broadcast
zone-attribute-group name testattgp vsan 2
  read-only
  broadcast
  qos priority high
```

Displays the fcalias Information for the Specified VSAN

```
switch# show fcalias vsan 2
fcalias name testfcalias vsan 2
 pwwn 21:00:00:20:37:39:b0:f4
 pwwn 21:00:00:20:37:6f:db:dd
 pwwn 21:00:00:20:37:a6:be:2f
```

Displays the Zone Status for the Specified VSAN

```
switch(config)# show zone status vsan 1
VSAN: 1 default-zone: deny distribute: active only Interop: default
mode: basic merge-control: allow
session: none
hard-zoning: enabled broadcast: disabled
smart-zoning: disabled
rscn-format: fabric-address
activation overwrite control:disabled
Default zone:
qos: none broadcast: disabled ronly: disabled
Full Zoning Database :
DB size: 4 bytes
Zonesets:0 Zones:0 Aliases: 0
Active Zoning Database :
Database Not Available
Current Total Zone DB Usage: 4 / 2097152 bytes (0 % used)
Pending (Session) DB size:
Full DB Copy size: n/a
Active DB Copy size: n/a
SFC size: 4 / 2097152 bytes (0 % used)
Status:
```

Displays the Pending ZoneSet Information for the VSAN to be Committed

```
switch# show zoneset pending vsan 2
No pending info found
```

Displays the Pending Zone Information for the VSAN to be Committed

```
switch# show zone pending vsan 2
No pending info found
```

Displays the Pending Zone Information for the VSAN to be Committed

```
switch# show zone-attribute-group pending vsan 2
No pending info found
```

Displays the Pending Active ZoneSet Information for the VSAN to be Committed

```
switch# show zoneset pending active vsan 2
No pending info found
```

Displays the Difference Between the Pending and Effective Zone Information for the Specified VSAN

```
switch# show zone pending-diff vsan 2
zone name testzone vsan 2
   -  member pwwn 21:00:00:20:37:4b:00:a2
   +  member pwwn 21:00:00:20:37:60:43:0c
```

Exchange Switch Support (ESS) defines a mechanism for two switches to exchange various supported features.

Displays the ESS Information for All Switches in the Specified VSAN

```
switch# show zone ess vsan 2
ESS info on VSAN 2 :
    Domain : 210, SWWN : 20:02:00:05:30:00:85:1f, Cap1 : 0xf3, Cap2 : 0x0
```

Displays the Pending fcalias Information for the VSAN to be Committed

```
switch# show fcalias pending vsan 2
No pending info found
```

Controlling Zoning Configuration Sessions

In enhanced mode zoning, the switch in which a zoning session is started takes the fabric wide zoning configuration lock for the VSAN. This configuration lock prevents users on other switches in the fabric from making simultaneous, and possibly conflicting, configuration changes. However, by default, the same user is allowed to login multiple times on the switch where the configuration is locked and start multiple zoning configuration sessions. This also may result in conflicting or undesired zone configuration.

The single session option enforces a maximum of one zoning configuration session per VSAN at a time on the switch with the zone configuration fabric lock. This limit causes the switch to disallow any new zoning configuration sessions to be started in the same VSAN. This limit also applies to any configuration source such as another user, NX-API.

Note

- If the login session is disconnected for some reason, such as after a supervisor switchover, the zone session remains with the fabric wide lock and any pending changes. In this case, when the single session option is enabled, no further zone configuration from any other login to the switch will be allowed. Attempts to do so will be rejected with an error message displaying the old session owner information. This information can also be viewed using the **show zone status** command. To recover, the session lock must be cleared using the **clear zone lock** command from the switch where the session got locked. Clearing the session lock will delete any pending zoning configuration and the zone configuration changes will have to be reentered. Use the **show zone pending-diff** command to display any pending zoning configuration changes prior to clearing the zone lock.
- This option is available from Cisco MDS NX-OS Release 8.4(2).
- Ensure that you disable this option before downgrading to any earlier NX-OS release. Otherwise, the downgrade process will fail.

Configuring Zoning Session Limit

The command **zone mode enhanced vsan** *id* **single-session** is configured per switch, not across the entire fabric.

To configure zoning session limit in a VSAN, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **zone mode enhanced vsan** *id* **single-session**

Enables the single session option in the specified VSAN.

Step 3 switch(config)# **no zone mode enhanced vsan** *id* **single-session**

Disables the single session option in the specified VSAN and leaves the VSAN in enhanced zoning mode.

Compacting the Zone Database for Downgrading

Prior to Cisco SAN-OS Release 6.2(7), only 8000 zones are supported per VSAN. If you add more than 8000 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 8000 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 6.2(5) or earlier. Perform this procedure for every VSAN on the switch with more than 8000 zones.

Note A merge failure occurs when a switch supports more than 8000 zones per VSAN but its neighbor does not. Also, zoneset activation can fail if the switch has more than 8000 zones per VSAN and not all switches in the fabric support more than 8000 zones per VSAN.

To delete zones and compact the zone database for a VSAN, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **no zone name ExtraZone vsan 10**

Deletes a zone to reduce the number of zones to 8000 or fewer.

Step 3 switch(config)# **zone compact vsan 10**

Compacts the zone database for VSAN 10 to recover the zone ID released when a zone was deleted.

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

Zone and ZoneSet Analysis

To better manage the zones and zone sets on your switch, you can display zone and zoneset information using the **show zone analysis** command.

Full Zoning Analysis

```
switch# show zone analysis vsan 1
Zoning database analysis vsan 1
 Full zoning database
   Last updated at: 15:57:10 IST Feb 20 2006
   Last updated by: Local [ CLI ]
   Num zonesets: 1
   Num zones: 1
   Num aliases: 0
   Num attribute groups: 0
   Formattted size: 36 bytes / 2048 Kb
 Unassigned Zones: 1
   zone name z1 vsan 1
```

Note The maximum size of the full zone database per VSAN is 4096 KB.

Active Zoning Database Analysis

```
switch(config-zone)# show zone analysis active vsan 1
Zoning database analysis vsan 1
  Active zoneset: qoscfg
    Activated at: 14:40:55 UTC Mar 21 2014
    Activated by: Local [ CLI ]
    Default zone policy: Deny
    Number of devices zoned in vsan: 8/8 (Unzoned: 0)
    Number of zone members resolved: 10/18 (Unresolved: 8)
    Num zones: 4
    Number of IVR zones: 0
    Number of IPS zones: 0
    Formatted size: 328 bytes / 4096 Kb
```



Note The maximum size of the zone database per VSAN is 4096 KB.

ZoneSet Analysis

```
switch(config-zone)# show zone analysis zoneset qoscfg vsan 1
Zoning database analysis vsan 1
  Zoneset analysis: qoscfg
    Num zonesets: 1
    Num zones: 4
    Num aliases: 0
    Num attribute groups: 1
    Formatted size: 480 bytes / 4096 Kb
```

Displays the Zone Status

```
switch(config-zone)# show zone status
VSAN: 1 default-zone: deny distribute: active only Interop: default
mode: basic merge-control: allow
session: none
hard-zoning: enabled broadcast: disabled
smart-zoning: disabled
rscn-format: fabric-address
activation overwrite control:disabled
Default zone:
qos: none broadcast: disabled ronly: disabled
Full Zoning Database :
DB size: 4 bytes
Zonesets:0 Zones:0 Aliases: 0
Active Zoning Database :
Database Not Available
Current Total Zone DB Usage: 4 / 2097152 bytes (0 % used)
Pending (Session) DB size:
Full DB Copy size: n/a
Active DB Copy size: n/a
SFC size: 4 / 2097152 bytes (0 % used)
Status:
VSAN: 8 default-zone: deny distribute: full Interop: default
mode: basic merge-control: allow
session: none
hard-zoning: enabled broadcast: disabled
smart-zoning: disabled
rscn-format: fabric-address
```

```
Default zone:
qos: none broadcast: disabled ronly: disabled
Full Zoning Database :
DB size: 1946498 bytes
Zonesets:6 Zones:8024 Aliases: 0
Active Zoning Database :
DB size: 150499 bytes
Name: zoneset-1000 Zonesets:1 Zones:731
Current Total Zone DB Usage: 2096997 / 2097152 bytes (99 % used)
Pending (Session) DB size:
Full DB Copy size: n/a
Active DB Copy size: n/a
SFC size: 2096997 / 2097152 bytes (99 % used)
Status: Zoneset distribution failed [Error: Fabric changing Dom 33]:
at 17:05:06 UTC Jun 16 2014
VSAN: 9 default-zone: deny distribute: full Interop: default
mode: enhanced merge-control: allow
session: none
hard-zoning: enabled broadcast: enabled
smart-zoning: disabled
rscn-format: fabric-address
Default zone:
qos: none broadcast: disabled ronly: disabled
Full Zoning Database :
DB size: 2002584 bytes
Zonesets:4 Zones:7004 Aliases: 0 Attribute-groups: 1
Active Zoning Database :
DB size: 94340 bytes
Name: zoneset-hac13-200 Zonesets:1 Zones:176
Current Total Zone DB Usage: 2096924 / 2097152 bytes (99 % used)
Pending (Session) DB size:
Full DB Copy size: 0 bytes
Active DB Copy size: 0 bytes
SFC size: 0 / 2097152 bytes (0 % used)
Status: Activation completed at 17:28:04 UTC Jun 16 2014
VSAN: 12 default-zone: deny distribute: full Interop: default
mode: enhanced merge-control: allow
session: none
hard-zoning: enabled broadcast: enabled
smart-zoning: disabled
rscn-format: fabric-address
Default zone:
qos: none broadcast: disabled ronly: disabled
Full Zoning Database :
DB size: 84 bytes
Zonesets:0 Zones:1 Aliases: 0 Attribute-groups: 1
Active Zoning Database :
DB size: 144 bytes
Name: zs1 Zonesets:1 Zones:2
Current Total Zone DB Usage: 228 / 2097152 bytes (0 % used)
Pending (Session) DB size:
Full DB Copy size: 0 bytes
Active DB Copy size: 0 bytes
SFC size: 0 / 2097152 bytes (0 % used)
Status: Commit completed at 14:39:33 UTC Jun 27 201
```

Displaying the System Defalult Zone

```
switch(config)# show system default zone
system default zone default-zone deny
system default zone distribute active only
system default zone mode basic
```

```
system default zone gs read-write
system default zone smart-zone disabled
```

See the Cisco MDS 9000 Series Command Reference for the description of the information displayed in the command output.

Zoning Best Practice

A Cisco Multilayer Director Switch (MDS) uses a special kind of memory called Ternary Content Addressable Memory (TCAM) on its Fibre Channel (FC) linecards. This special memory provides an Access Control List (ACL) type of function for Cisco MDS. The process that controls this functionality is called the ACLTCAM. The E/TE ports (Inter Switch Links - ISLs) and F (Fabric) ports have their own programming, which is unique to their respective port types.

TCAM Regions

TCAM is divided into several regions of various sizes. The main regions and the type of programming contained in each region are described in Table 11: TCAM Regions , on page 117:

Table 11: TCAM Regions

Region	Programming Type
Region 1 - TOP SYS	Fabric-Login, Port-Login, Diagnostics features (10%-20%)
Region 2 - SECURITY	Security, Interop-Mode-4 features, IVR ELS capture (5%-10%)
Region 3 - Zoning	Zoning features, including IVR and SAN Analytics (50%-75%)
Region 4 - Bottom[2]	PLOGI,ACC, and FCSP trap, ISL, ECHO-permit (10%-20%)

[2] When a hard-zoning failure occurs, Region 4 (bottom region) is used to program wildcard entries to allow any-to-any communication.

TCAM regions are automatically configured and cannot be changed. TCAM is allocated on a per-module and per-forwarding engine (fwd-eng) basis.

TCAM space on MDS 9148S and MDS 9250i fabric switches is significantly less than that on the director-class Fibre Channel modules and newer fabric switches such as MDS 9396S, MDS 9132T, and the switches that will be launched in the future.

When a port comes online, some amount of basic programming is needed on that port. This programming differs according to the port type. This basic programming is minimal and does not consume many TCAM entries. Typically, programming is performed on inputs such that frames entering the switch are subject to this programming and frames egressing the switch are not.

ACL TCAM Alerting

Starting from Cisco MDS NX-OS Release 8.3(1), ACL TCAM usage alerting syslog messages were introduced on all Cisco MDS switches except on Cisco MDS 9148S and MDS 9250i switches. From Cisco MDS NX-OS Release 8.3(2), ACL TCAM usage alerting syslog messages were introduced on Cisco MDS 9148S and MDS 9250i switches.

- When the TCAM usage crosses 80% in the module, direction, region, and forwarding engine listed, the following system message is generated. This system message does not indicate that TCAM was exhausted or any TCAM programming failed.

```
%ACLTCAM-SLOT1-4-REGION_RISING_THRESHOLD: ACL (region) (input | output) region usage
(num of in use entries of total entries) exceeded 80% on forwarding engine (num)
```

- When TCAM usage falls below the 80% threshold in the module, region, direction, and forwarding engine indicated, the following system message is generated. This system message does not indicate that TCAM was exhausted or any TCAM programming failed.

```
%ACLTCAM-SLOT1-4-REGION_FALLING_THRESHOLD: ACL (region) (input | output) region usage
(num of in use entries of total entries) fell below 80% on forwarding engine (num)
```

- When the overall TCAM usage indicated in the forwarding engine crosses 60% in the module, direction, and forwarding engine indicated, the following system message is generated:

```
%ACLTCAM-SLOT1-4-TOTAL_RISING_THRESHOLD: ACL total (input | output) usage (num of in
use entries of total entries) exceeded 60% on forwarding engine (num)
```

- When the overall TCAM usage indicated in the forwarding engine falls below 60% in the module, direction, and forwarding engine indicated, the following system message is generated:

```
%ACLTCAM-SLOT1-4-TOTAL_FALLING_THRESHOLD: ACL total (input | output) usage (num of in
use entries of total entries) fell below 60% on forwarding engine (num)
```

Use the **show system internal acl tcam-usage** command to display the ACLTCAM usage except for Cisco MDS 9148S and MDS 9250i switches. For Cisco MDS 9148S and MDS 9250i switches, use the **show system internal acltcam-soc tcam-usage** command.

If you see the TCAM usage alerting syslog messages, the zoning, port-channel port allocation, and analytics configurations may need to be investigated. If TCAM usage reaches 100%, it is likely that some devices will not be able to communicate with other devices that they are zoned with. Follow the recommendations listed in this section to reduce the TCAM usage.

Zoning Types

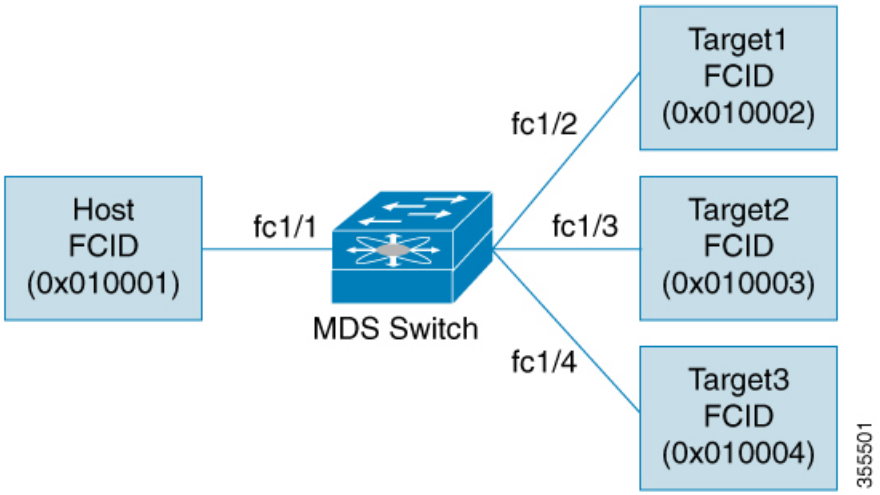
The Cisco MDS platform uses two types of zoning - 'Hard' and 'Soft' zoning.

Soft zoning - In this mode only control plane traffic is policed by the switch supervisor services. In particular, the Fibre Channel Name Server (FCNS) will limit the list of permitted devices in an FCNS reply to only those that are in the zone configuration. However, the end device data plane traffic is unpoliced. This means a rogue end device may connect to other devices it is not zoned with.

Hard zoning - In this mode both control plane and data plane traffic are policed. Control plane traffic is policed by the switch supervisor and data plane traffic is policed on each ingress port with hardware assistance. The policing rules are set by the zoneset which programmed into each linecard. The destination of each frame is checked by hardware and, if it is not permitted by zoning, it is dropped. In this mode any device can only communicate with end devices it is authorized to.

By default, both types of zoning are enabled, with hard zoning used in priority over soft zoning. In the event that the system is unable to use hard zoning due to hardware resource exhaustion it will be disabled and the system will fall back to use soft zoning

The following example shows how Cisco MDS programs TCAM on a port:



The following example shows a zone in the active zone set for a VSAN. This is the basic programming that exists on an interface because of Hard zoning.

```
zone1
member host (FCID 0x010001)
member target1 (FCID 0x010002)
```

In such a scenario, the following is the ACL programming:

```
fc1/1 - Host interface
Entry#    Source ID      Mask      Destination ID        Mask      Action
1         010001         ffffff    010002(target1)       ffffff    Permit
2         000000         000000    000000                000000    Drop
fc1/2 - Target1 interface
Entry#    Source ID      Mask      Destination ID        Mask       Action
1         010002         ffffff    010001(Host)          ffffff     Permit
2         000000         000000    000000                000000     Drop
```

Note In addition to what is provided here, additional programming exists. Moreover, any TCAM table is ended by a drop-all entry.

The mask indicates which parts of the FCIDs are matched with the input frame. So, when there is a mask 0xffffff, the entire FCID is considered when matching it to the ACL entry. If the mask is 0x000000, none of it is considered because, by default, it will match all the FCIDs.

In the above programming example, note that when a frame is received on fc1/1, and if it has a source ID(FCID) of 0x010001(the host) and a destination ID(FCID) of 0x010002(Target1), it will be permitted and routed to the destination. If it is any other end-to-end communication, it will be dropped.

The following example shows another scenario where zoning is changed:

```
zone1
member host (FCID 010001)
member target1 (FCID 010002)
member target2 (FCID 010003)
member target3 (FCID 010004)
```

In such a scenario, the following is the ACL programming:

```
fc1/1 Host interface
Entry#    Source ID      Mask      Destination ID          Mask    Action
1         010001         ffffff    010002(target1)         ffffff  Permit
2         010001         ffffff    010003(target2)         ffffff  Permit
3         010001         ffffff    010004(target3)         ffffff  Permit
4         000000         000000    000000                  000000  Drop
fc1/2 - Target1 interface
Entry#    Source ID      Mask      Destination ID          Mask    Action
1         010002         ffffff    010001(host)            ffffff  Permit
2         010002         ffffff    010003(target2)         ffffff  Permit
3         010002         ffffff    010004(target3)         ffffff  Permit
4         000000         000000    000000                  000000  Drop
fc1/3 - Target2 interface
Entry#    Source ID      Mask      Destination ID          Mask    Action
1         010003         ffffff    010001(host)            ffffff  Permit
2         010003         ffffff    010002(target1)         ffffff  Permit
3         010003         ffffff    010004(target3)         ffffff  Permit
4         000000         000000    000000                  000000  Drop
fc1/4 - Target3 interface
Entry#    Source ID      Mask      Destination ID          Mask    Action
1         010004         ffffff    010001(host)            ffffff  Permit
2         010004         ffffff    010002(target1)         ffffff  Permit
3         010004         ffffff    010003(target2)         ffffff  Permit
4         000000         000000    000000                  000000  Drop
```

The above example demonstrates that the number of TCAM entries consumed by a zone (N) is equal to N*(N-1). So, a zone with four members would have used a total of 12 TCAM entries (4*3 = 12). Note the drop-all entry does not count against the N*(N-1) rule.

The above example shows two entries in each of the target interfaces (fc1/2-fc1/4) that are probably not needed since it is usually not advantageous to zone multiple targets together. For example, in fc1/2, there is an entry that permits Target1 to communicate with Target2, and an entry that permits Target1 to communicate with Target3.

As these entries are not needed and could even be detrimental, they should be avoided. You can avoid the addition of such entries by using single-initiator or single-target zones (or use Smart Zoning).



Note If the same two devices are present in more than one zone in a zone set, TCAM programming will not be repeated.

The following example shows a zone that is changed to three separate zones:

```
zone1
member host (FCID 010001)
member target1 (FCID 010002)
zone2
member host (FCID 010001)
member target2 (FCID 010003)
zone3
member host (FCID 010001)
member target3 (FCID 010004)
```

In such a scenario, the following is the ACL programming:

```
fc1/1 - Host interface - This would look the same
```

```
Entry#    Source ID      Mask      Destination ID          Mask    Action
1         010001         ffffff    010002(target1)         ffffff  Permit
2         010001         ffffff    010003(target2)         ffffff  Permit
3         010001         ffffff    010004(target3)         ffffff  Permit
4         000000         000000    000000                  000000  Drop
fc1/2 - Target1 interface
Entry#    Source ID      Mask      Destination ID          Mask    Action
1         010002         ffffff    010001(host)            ffffff  Permit
2         000000         000000    000000                  000000  Drop
fc1/3 - Target2 interface
Entry#    Source ID      Mask      Destination ID          Mask    Action
1         010003         ffffff    010001(host)            ffffff  Permit
2         000000         000000    000000                  000000  Drop
fc1/4 - Target3 interface
Entry#    Source ID      Mask      Destination ID          Mask    Action
1         010004         ffffff    010001(host)            ffffff  Permit
2         000000         000000    000000                  000000  Drop
```

Note that in the above example, the target-to-target entries are not found, and that six of the 12 entries are no longer programmed. This results in lesser use of TCAM and better security (only the host can communicate with the three targets, and the targets themselves can communicate only with one host, and not with each other).

Best Practises for Forwarding Engines

Cisco MDS switches use Ternary Content Addressable Memory (TCAM) on its Fibre Channel modules. TCAM provides an Access Control List (ACL) type of function for Cisco MDS. The process that controls this functionality is called ACLTCAM. The E or TE ports (ISLs) and F (Fabric) ports have their own programming that is unique to their respective port types.

TCAM is allocated to individual forwarding engines and forwarding engines are assigned a group of ports. Director-class Fibre Channel modules have more TCAM space than fabric switches. The number of forwarding engines, the ports assigned to each forwarding engine, and the amount of TCAM allocated to each forwarding engine is hardware dependent.

The following example shows an output from Cisco MDS 9148S:

```
switch# show system internal acl tcam-soc tcam-usage
TCAM Entries:
=============
                    Region1   Region2    Region3     Region4   Region5   Region6
Mod Fwd   Dir       TOP SYS  SECURITY    ZONING      BOTTOM    FCC DIS   FCC ENA
    Eng           Use/Total Use/Total  Use/Total   Use/Total Use/Total Use/Total
--- ---  ------ ---------- --------- ------------ --------- --------- ---------
1   1    INPUT     19/407     1/407      1/2852 *    4/407     0/0       0/0
1   1    OUTPUT     0/25      0/25       0/140       0/25      0/12      1/25
1   2    INPUT     19/407     1/407      0/2852 *    4/407     0/0       0/0
1   2    OUTPUT     0/25      0/25       0/140       0/25      0/12      1/25
1   3    INPUT     19/407     1/407      0/2852 *    4/407     0/0       0/0
1   3    OUTPUT     0/25      0/25       0/140       0/25      0/12      1/25
-------------------------------------------------
* 1024 entries are reserved for LUN Zoning purpose.
```

The above example indicates the following:

- There are three forwarding engines, 1 through 3.
- Since there are 48 ports on Cisco MDS 9148 switches, each forwarding engine handles 16 ports.

- Each forwarding engine has 2852 entries in region 3 (the zoning region) for input. This is the main region used, and consequently, has the largest amount of available entries.
- Forwarding engine 3 has only one entry that is currently in use out of the total 2852 in the zoning region.

The following example shows the output from Cisco MDS 9710 switch with a 2/4/8/10/16 Gbps Advanced Fibre Channel Module (DS–X9448–768K9):

```
F241-15-09-9710-2# show system internal acl tcam-usage
TCAM Entries:
=============
                   Region1   Region2    Region3     Region4   Region5   Region6
Mod Fwd   Dir      TOP SYS  SECURITY    ZONING      BOTTOM    FCC DIS   FCC ENA
    Eng           Use/Total Use/Total  Use/Total   Use/Total Use/Total Use/Total
--- ---  ------ ---------- --------- ------------ --------- --------- ---------
1   0    INPUT     55/19664    0/9840     0/49136*   17/19664    0/0       0/0
1   0    OUTPUT    13/4075    0/1643     0/11467     0/4075    6/1649   21/1664
1   1    INPUT     52/19664    0/9840     2/49136*   14/19664    0/0       0/0
1   1    OUTPUT     7/4078    0/1646     0/11470     0/4078    6/1652    5/1651
1   2    INPUT     34/19664    0/9840     0/49136*   10/19664    0/0       0/0
1   2    OUTPUT     5/4078    0/1646     0/11470     0/4078    6/1652    1/1647
1   3    INPUT     34/19664    0/9840     0/49136*   10/19664    0/0       0/0
1   3    OUTPUT     5/4078    0/1646     0/11470     0/4078    6/1652    1/1647
1   4    INPUT     34/19664    0/9840     0/49136*   10/19664    0/0       0/0
1   4    OUTPUT     5/4078    0/1646     0/11470     0/4078    6/1652    1/1647
1   5    INPUT     34/19664    0/9840     0/49136*   10/19664    0/0       0/0
1   5    OUTPUT     5/4078    0/1646     0/11470     0/4078    6/1652    1/1647
...
```

The above example indicates the following:

- There are six forwarding engines, 0 through 5.
- Since there are 48 ports on a Cisco MDS DS–X9448–768K9 module, each forwarding engine handles eight ports.
- Each forwarding engine has 49136 entries in region 3 (the zoning region) for input. This is the main region that is used, and consequently, has the largest amount of available entries.
- Forwarding engine 2 has only two entries that are currently in use out of the total 49136 in the zoning region.

The following example shows the output from Cisco MDS 9396V switch with a 2/4/8/10/16/32/64 Gbps Advanced Fibre Channel Module (DS–X9448–768K9):

```
switch9396v# show system internal acl tcam-usage
Input TCAM Entries:
====================
                   Region1   Region2       Region3         Region4
Mod Fwd   Dir      TOP SYS  SECURITY       ZONING          BOTTOM
    Eng           Use/Total Use/Total   Use/Total(Anl)  Use/Total(Anl)
--- ---  ------ ---------- ---------- ---------------- --------------
1   0    INPUT    126/26208    0/13120     0/65536(0)    28/26208(0)
1   1    INPUT    122/26208    0/13120     2/65536(0)    27/26208(0)
1   2    INPUT    150/26208    0/13120     0/65536(0)    32/26208(0)
1   3    INPUT    126/26208    0/13120     0/65536(0)    28/26208(0)

Output TCAM Entries:
=====================
    Fwd           Region1   Region2       Region3         Region4      Region5   Region6
```

```
Mod Eng/  Dir      TOP SYS  SECURITY       ZONING          BOTTOM       FCC DIS   FCC ENA
    Port         Use/Total Use/Total   Use/Total(Anl)  Use/Total(Anl) Use/Total Use/Total
    Num
--- ---  ------ ---------- ---------- ---------------- -------------- --------- ---------
1   0    OUTPUT     4/51       0/51       0/281  (0)      0/51  (0)       4/25      3/51
1   1    OUTPUT     4/51       0/51       0/281  (0)      0/51  (0)       4/25      1/51
1   2    OUTPUT     4/51       0/51       0/281  (0)      0/51  (0)       4/25      1/51
1   3    OUTPUT     4/51       0/51       0/281  (0)      0/51  (0)       4/25      1/51
1   4    OUTPUT     4/51       0/51       0/281  (0)      0/51  (0)       4/25      1/51
.
..
...
1   94    OUTPUT     4/51       0/51       0/281  (0)      0/51  (0)       4/25      1/51
1   95    OUTPUT     4/51       0/51       0/281  (0)      0/51  (0)       4/25      1/51
Note: Analytics Entry Count(Anl) included in Use count.
```

The above example indicates the following:

- There are four forwarding engines, 0 through 3.
- Since there are 96 ports on a Cisco MDS DS-C9396V-K9-SUP module, each forwarding engine handles twenty-four ports.
- Each forwarding engine has 65536 entries in region 3 (the zoning region) for input. This is the main region that is used, and consequently, has the largest amount of available entries.
- Forwarding engine 2 has only two entries that are currently in use out of the total 65536 in the zoning region.

Note The commands that are used to view TCAM usage on fabric switches are different from the ones used for director–class switches. For MDS 9148, MDS 9148S, and MDS 9250i fabric switches, use the **show system internal acltcam-soc tcam-usage** command. For director class switches, MDS 9396V, MDS 9396S, and 32 Gbps fabric switches, use the **show system internal acl tcam-usage** command.

The following table provides information about ports to forwarding engines mapping:

Table 12: Ports to Forwarding Engines Mapping

Switch or Module	Forwarding Engines	Port Ranges	Forwarding Engine Number	Zoning Region Entries	Bottom Region Entries
MDS 9132T	2	1–16	0	49136	19664
		17–32	1	49136	19664
MDS 9148	3	fc1/25–36 and fc1/45–48	1	2852	407
		fc1/5–12 and fc1/37–44	2	2852	407
		fc1–4 and fc1/13–24	3	2852	407

Switch or Module	Forwarding Engines	Port Ranges	Forwarding Engine Number	Zoning Region Entries	Bottom Region Entries
MDS 9148S	3	fc1/1–16	1	2852	407
		fc1/17–32	2	2852	407
		fc1/33–48	3	2852	407
MDS 9148T	3	1–16	0	49136	19664
		17–32	1	49136	19664
		33–48	2	49136	19664
MDS 9250i	4	fc1/5–12 and eth1/1–8	1	2852	407
		fc1/1–4, fc1/13–20, and fc1/37–40	2	2852	407
		fc1/21–36	3	2852	407
		ips1/1–2	4	2852	407
MDS 9396S	12	fc1/1–8	0	49136	19664
		fc1/9–16	1	49136	19664
		fc1/17–24	2	49136	19664
		fc1/25–32	3	49136	19664
		fc1/33–40	4	49136	19664
		fc1/41–48	5	49136	19664
		fc1/49–56	6	49136	19664
		fc1/57–64	7	49136	19664
		fc1/65–72	8	49136	19664
		fc1/73–80	9	49136	19664
		fc1/81–88	10	49136	19664
		fc1/89–96	11	49136	19664

Switch or Module	Forwarding Engines	Port Ranges	Forwarding Engine Number	Zoning Region Entries	Bottom Region Entries
MDS 9396T	6	1–16	0	49136	19664
		17–32	1	49136	19664
		33–48	2	49136	19664
		49–64	3	49136	19664
		65-80	4	49136	19664
		81-96	5	49136	19664
DS–X9248–48K9	1	1–48	0	27168	2680
DS–X9248–96K9	2	1–24	0	27168	2680
		25–48	1	27168	2680
DS–X9224–96K9	2	1–12	0	27168	2680
		13–24	1	27168	2680
DS–X9232–256K9	4	1–8	0	49136	19664
		9–16	1	49136	19664
		17–24	2	49136	19664
		25–32	3	49136	19664
DS–X9248–256K9	4	1–12	0	49136	19664
		13–24	1	49136	19664
		25–36	2	49136	19664
		37–48	3	49136	19664
DS–X9448–768K9	6	1–8	0	49136	19664
		9–16	1	49136	19664
		17–24	2	49136	19664
		25–32	3	49136	19664
		33–40	4	49136	19664
		41–48	5	49136	19664

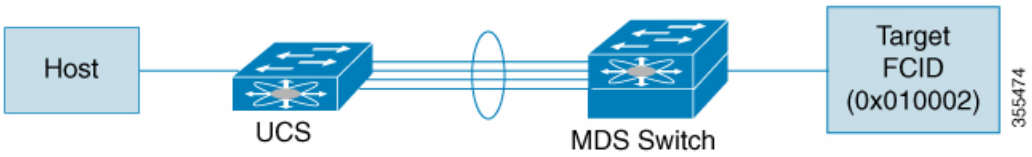
Switch or Module	Forwarding Engines	Port Ranges	Forwarding Engine Number	Zoning Region Entries	Bottom Region Entries
DS–X9334–K9	3	1–8	0	49136	19664
		9–16	1	49136	19664
		17–24	2	49136	19664
DS–X9648–1536K9	3	1–16	0	49136	19664
		17–32	1	49136	19664
		33–48	2	49136	19664
DS-C9124V-K9	1	1-24	0	65536	26208
DS-C9148V-24EK9	2	1-24	0	65536	26208
		25-48	1	65536	26208
DS-C9220I-K9	1	1-12	0	49136	19664
DS-X9748-3072-K9	2	1-24	0	65536	26208
		25-48	1	65536	26208
DS-C9396V-K9	4	1-24	0	65536	26208
		25-48	1	65536	26208
		49-72	2	65536	26208
		73-96	3	65536	26208

F, TF, NP, and TNP Port Channels

Note It is not recommended that you use interface, fWWN, or domain-ID based zoning for devices that are connected to the edge Cisco N-Port Virtualization (NPV) switches.

F port channels provide fault tolerance and performance benefits on connections to N-Port Virtualization (NPV) switches, including Cisco UCS Fabric Interconnects (FIs). F port channels present unique challenges to ACL TCAM programming. When F ports are aggregated into a port channel, ACL TCAM programming is repeated on each member interface. Consequently, these types of port channels multiply the amount of TCAM entries needed. Because of this, it is imperative that the member interfaces are allocated as optimally as possible, and that zoning best practices are also followed. Given that F port channels can contain 100+ host logins, TCAM can easily be exceeded, especially for fabric switches if best practices are not followed.

The following is a sample topology:



This example assumes that the port channel (PC) contains 8 interfaces, fc1/1-fc1/8.

In addition, the following two zones are active:

```
zone1
member host (host 0x010001)
member target1 (target1 0x010002)
zone2
member host (host 0x010001)
member target2 (target2 0x010003)
```

In such a scenario, the following ACL programming will be present on each member of the PC:

```
fc1/1(through fc1/8) (port-channel)
Entry#    Source ID     Mask         Destination ID         Mask             Action
1         010001       ffffff        010002(target1)       ffffff           Permit
2         010001       ffffff        010003(target2)       ffffff           Permit
3         000000       000000        000000                000000           Drop
```

The above example shows the ACL TCAM programming that will be duplicated on each member of the F port-channel.

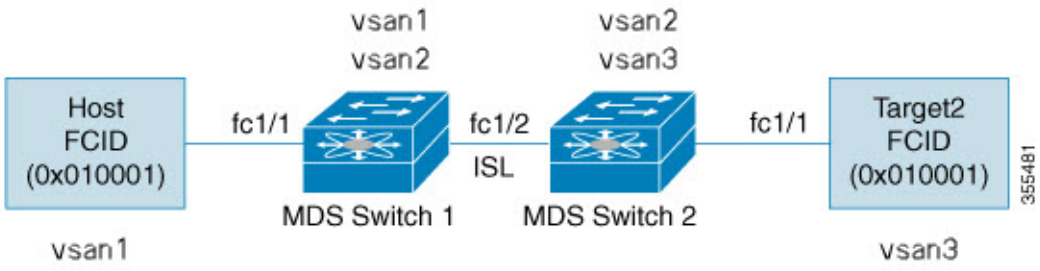
The following are the best practices for efficient use of TCAM with respect to F ports and F port-channels to optimize TCAM usage on a forwarding engine:

- Distribute port-channel member interfaces into different forwarding engines, especially on fabric switches.
- If TCAM usage is still too high in the case of port-channel with a large number of interfaces, then split the port-channel into two separate port-channels each with half the interfaces. This provides redundancy but reduces the number of FLOGIs per individual port-channel and thus reduces TCAM usage.
- Distribute member interfaces into separate linecards on director-class switches.
- Distribute member interfaces into forwarding engines with lower TCAM zoning region usage.
- Use single-initiator zones, single-target zones, or Smart Zoning.

Best Practises for E and TE Port Channels and IVR

Port channels provide Inter Switch Links (ISLs) between switches. Typically, there is minimal TCAM programming on these types of interfaces. When the Inter VSAN Routing(IVR) feature is being deployed, extensive TCAM programming can exist on ISLs because the IVR topology transitions from one VSAN to another. Most of the considerations that apply on F/TF port channels will be applicable here too.

The following is an example of a topology:



In this topology:

- Both Cisco MDS 9148S-1 and MDS 9148S-2 are in the IVR VSAN topology:

```
MDS9148S-1 vsan 1 and vsan 2
MDS9148S-2 vsan 2 and vsan 3
```

- IVR NAT is configured.
- VSAN 2 is the transit VSAN.

```
FCIDs per VSAN:
            VSAN 1  VSAN 2  VSAN 3
Host         010001  210001  550002
Target1     440002  360002  030001
```

Note Domains 0x44 in VSAN 1, 0x21 and 0x36 in VSAN 2, and 0x55 in VSAN 3 are virtual domains created by IVR NAT.

- The following is the IVR zoning topology:

```
ivr zone zone1
member host vsan 1
member target1 vsan3
```

- The following is the ACL TCAM programming for the IVR zoning topology:

```
MDS9148S-1  fc1/1(Host) - VSAN 1
Entry#    Source ID        Mask      Destination ID          Mask    Action
1         010001(host)     ffffff    440002(target1)         ffffff  Permit
           - Forward to fc1/2
       - Rewrite the following information:
         VSAN to 2
         Source ID to 210001
         Destination ID to 360002
2         000000         000000    000000                  000000  Drop
MDS9148S-1  fc1/2(ISL) - VSAN 2
Entry#    Source ID           Mask      Destination ID          Mask    Action
1         360002(Target1)     ffffff    210001(host)            ffffff  Permit
       - Forward to fc1/2
       - Rewrite the following information:
         VSAN to 1
```

```
          Source ID to 440002
          Destination ID to 010001
MDS9148S-2 fc1/2(ISL) - VSAN 2
Entry#    Source ID        Mask      Destination ID          Mask    Action
1          210001(host)     ffffff    360002(target1)         ffffff  Permit
       - Forward to fc1/2
       - Rewrite the following information:
          VSAN to 3
          Source ID to 550002
          Destination ID to 030001
MDS9148S-2  fc1/1(Target1) - VSAN 3
Entry#    Source ID            Mask      Destination ID          Mask    Action
1          030001(Target1)      ffffff    550002(host)            ffffff  Permit
       - Forward to fc1/2
       - Rewrite the following information:
          VSAN to 2
          Source ID to 360002
          Destination ID to 210001
2          000000          000000    000000                  000000  Drop
```

Note Besides the entries in this example, there are other entries that IVR adds to capture important frames such as PLOGIs, PRILIs, and ABTS.

The programming on the host and target1 ports is similar to the way it is without IVR, except that the FCIDs and VSANs are explicitly forwarded to an egress port and are rewritten to values that are appropriate for the transit VSAN (VSAN 2). These forwarding and rewrite entries are separate and are not included in the TCAM-usage values.

However,now, on the ISLs in both the switches, programming that did not exist earlier is present. When frames from Host to Target1 are received by Cisco MDS 9148S-2 fc1/2, they are rewritten to the values in VSAN3 where the target resides. In the reverse direction, when frames from Target1 to the Host are received by Cisco MDS 9148S-1 fc1/2, they are rewritten to the values in VSAN 1 where the Host resides. Therefore, for each VSAN transition on an ISL (that typically occurs across a transit VSAN) there is TCAM programming for each device in the IVR zone set.

Consequently,most of the best practices followed for the F and TF port channels should be followed to ensure that TCAM is utilized as efficiently as possible for the following purposes:

Note Unlike F and TF port-channels, the ACLTCAM programming on ISLs will be the same quantity regardless if the ISLs are part of a port-channel or not. If there are "n" ISLs between two MDS switches, then it doesn't matter if they are in one port-channel, two port-channels or just individual links. The ACLTCAM programming will be the same.

- Distribute port-channel member interfaces into different forwarding engines, especially on fabric switches.
- Distribute member interfaces into different linecards on director-class switches.
- Distribute member interfaces into forwarding engines with lower TCAM zoning region usage.
- Use single-initiator zones, single-target zones, or Smart Zoning.

Enhancing Zone Server Performance

Zone Server-Fibre Channel Name Server Shared Database

This options provides a shared database for the Zone Server and the Fibre Channel Name Sever (FCNS) to interact with one another. Sharing a database reduces the dependency of the FCNS on the zone server to manage soft zoning.

Note By default, the Zone Server- FCNS Shared Database option is enabled.

Enabling the Zone Server-FCNS Shared Database

To enable the Zone Server-FCNS shared database, perform the following steps:

Procedure

Step 1 Enter the configuration mode:

switch # **configure terminal**

Step 2 Enable database sharing for an active zone set in VSAN 1:

switch(config)# **zoneset capability active mode shared-db vsan 1**

Example

Enabling Zone Server-FCNS Shared Database

This example shows how to enable database sharing for the active zoneset in VSAN 1 only:

```
switch(config)# zoneset capability active mode shared-db vsan 1
SDB Activation success
```

Disabling Zone Server-FCNS shared database

To disable an active zone set in VSAN 1, perform the following step:

Procedure

Step 1 Enter global configuration mode:

switch# **configure terminal**

Step 2 Disable an active zone set in VSAN 1:

switch(config)# **no zoneset capability active mode shared-db vsan 1**

Example

Disabling Zone Server-FCNS Shared Database

This example shows how to disable database sharing for the active zone set in VSAN 1:

```
switch(config)# no zoneset capability active mode shared-db vsan 1
SDB Deactivation success
```

Zone Server SNMP Optimization

This option enables zone server-scaling enhancements for Simple Network Management Protocol (SNMP) operations, such that the zone server is not utilized for every zone query issued by the SNMP.

Note By default, the Zone Server-SNMP Optimization option is enabled..

Enabling Zone Server SNMP Optimization

To enable zone server-scaling enhancements for SNMP operations, perform the following procedure:

Procedure

Step 1 Enter the configuration mode:

switch # **configure terminal**

Step 2 Enable zone server-SNMP optimization:

switch(config)# **zone capability shared-db app snmp**

Step 3 Display the status of the configuration:

switch(config)# **show running | i shared-db**

Example

Enabling Zone Server- SNMP Optimizations

This example shows how to enable zone server-SNMP optimization:

```
switch(config)# zone capability shared-db app snmp
```

Disabling Zone Server SNMP Optimization

To disable zone server-SNMP optimizations, perform the following procedure:

Procedure

Step 1 Di the configuration mode:

switch # **configure terminal**

Step 2 Disable the zone server-SNMP optimizations:

switch(config)# **no zone capability shared-db app snmp**

Example

Disabling Zone Server- SNMP Optimizations

This example shows how to disable zone server-SNMP optimization:

```
switch(config)# no zone capability shared-db app snmp
```

Zone Server Delta Distribution

This feature helps distribute the difference in the zone changes between the existing zone database and the updated zone database across all the switches in a fabric. This distribution of delta changes helps avoid large payload distribution across switches whenever a zone database in modified.

Note

- By default, the Zone Server Delta Distribution feature is disabled and functions in enhanced mode only.
- All the switches in a fabric should have the Zone Server Delta Distribution feature enabled. If a switch is added to the fabric with Zone Server Delta Distribution feature disabled, it will disable the Zone Server Delta Distribution feature on all the switches in the fabric.
- The Zone Server Delta Distribution feature is supported only on Cisco MDS switches, beginning from Cisco MDS NX-OS Release 7.3(0)D1(1).
- The Zone Server Delta Distribution feature is not available on IVR enabled switches.

Enabling Zone Server Delta Distribution

To enable the distribution of data changes in a zone server, perform the following procedure:

Procedure

Step 1 Enter the configuration mode:

switch # **configure terminal**

Step 2 Enable the distribution of data changes in a zone in enhanced mode:

switch(config)# **zone capability mode enhanced distribution diffs-only**

Step 3 Display the status of delta distribution (changes in data) in a fabric:

switch(config)# **show running | include diffs-only**

Example

Enabling Zone Server Delta Distribution

This example shows how to enable distribution of changes in data in a Zone Server:

```
switch(config)# zone capability mode enhanced distribution diffs-only
```

Disabling Zone Server Delta Distribution

To disable the distribution of data changes in a zone server, perform the following procedure:

Procedure

Step 1 Enter the configuration mode:

switch # **configure terminal**

Step 2 Disable the distribution of data changes in a zone:

switch(config)# **no zone capability mode enhanced distribution diffs-only**

Example

Disabling Zone Server Delta Distribution

This example shows how to disable distribution of changes in data in a Zone Server:

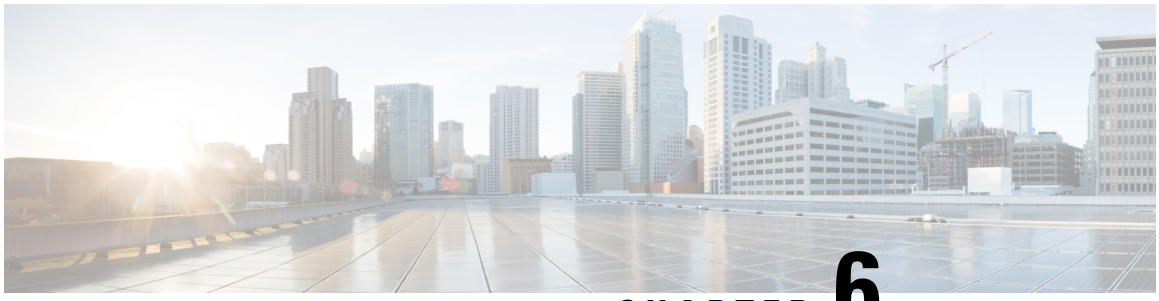
```
switch(config)# no zone capability mode enhanced distribution diffs-only
```

Default Settings

Table lists the default settings for basic zone parameters.

Table 13: Default Basic Zone Parameters

Parameter	Default
Default zone policy	Denied to all members.
Full zone set distribute	The full zone set is not distributed.
Zone-based traffic priority	Low.
Broadcast frames	Unsupported.
Enhanced zoning	Disabled.
Smart zoning	Disabled.



CHAPTER 6

Distributing Device Alias Services

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

This chapter includes the following sections:

- Understanding Device Aliases, on page 135
- Device Alias Modes, on page 135
- Device Alias Databases, on page 142
- About Legacy Zone Alias Configuration Conversion, on page 148
- Database Merge Guidelines, on page 150
- Device Alias Configuration Verification, on page 150
- Default Settings , on page 152
- Resolving Device Alias Merge Failures, on page 153

Understanding Device Aliases

While the port WWN (pWWN) of a device has to be specified to configure different features (zoning, QoS, and 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 may cause unexpected results. You can avoid this if you define a user-friendly name for a pWWN 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.

Device Alias Modes

Device-Alias Basic Mode and Enhanced Mode

The device alias feature supports two modes, basic mode and enhanced mode.

Note

- For applications such as NX-OS processes (zone, dpvm, ivr, and so on), the device-alias configurations get mapped to their PWWNs when device-alias is in basic mode. Whereas if device-alias is in enhanced mode, the device-alias configuration of the applications do not get mapped to their PWWNs immediately but will remain as configured in application which is referred as native form or format.
- From Cisco MDS NX-OS Release 8.5(1), the default device alias mode is enhanced mode.

When using device alias in basic mode, NX-OS processes such as zone, DPVM, IVR, and so on, immediately expand the device alias name into its associated pWWN in the configuration. For example, when a device alias member is added to a zone, it will be added as a pWWN member and not as a device alias member. Therefore, when you change the pWWN for the device alias entry any configuration (besides device alias) does not get automatically updated. You should manually edit the zones containing that device alias by removing the old entry and reconfiguring the zones and any other configuration where the PWWN is used by removing the old PWWN entry and adding it back with the same device alias name that now has the updated PWWN. After that is done, the configuration should be made active in whatever method is appropriate for the change. For example, if a zone was modified then the zoneset should be reactivated and committed, if necessary.

When using device alias in enhanced mode, NX-OS processes such as zone, DPVM, IVR, and so on, store the device alias names natively in the configuration as specified instead of expanding to pWWNs. The applications track the device alias database changes and take the necessary actions to enforce any changes made (for example like renaming a device alias).

In this mode, since the configuration is accepted in the native form, when the pWWN for the device alias is changed, the zones or other configuration containing that device alias are automatically updated.

Guidelines and Limitations

A native device alias configuration is not accepted in the interop mode VSAN. IVR zone set activation fails in the interop mode VSANs if the corresponding twilight zones being injected are native device alias members.

Device-Alias Mode Default

From Cisco MDS NX-OS Release 8.5(1), the default device alias mode is enhanced mode. Prior to Cisco MDS NX-OS Release 8.5(1), the default device alias mode was basic mode. After an upgrade from an earlier release to Cisco MDS NX-OS Release 8.5(1) or later releases, the device alias mode is set to enhanced mode only when there are no device alias entries configured and the device alias mode is basic. If device alias entries are present or the device alias mode is already in enhanced mode, then the device alias mode is not changed. When a switch initially boots up Cisco MDS NX-OS Release 8.5(1) or later releases, the default device alias mode is set to enhanced mode. If the switch is being downgraded from Cisco MDS NX-OS Release 8.5(1) or later releases to a Cisco MDS NX-OS Release 8.4(2b) or earlier releases and there are no device alias entries configured and the device alias mode is not set, the default alias mode reverts to basic mode. If device alias entries are present or the device alias mode is set, the device alias mode is not changed.

Downgrading from Cisco MDS NX-OS Release 8.5(1) to Release 8.4(2c) is a disruptive operation. Therefore, the device alias configuration is not retained on the switch and the default device alias mode is changed to basic mode, which is the default device alias mode in the Release 8.4(2c), after downgrading.

When the default has been set to enhanced mode the following syslog message will be displayed:

```
%DEVICE-ALIAS-2-DDAS_DEFAULT_MODE: Device alias mode has been set to enhanced mode
```

Note If a new switch running Cisco MDS NX-OS Release 8.5(1) or later releases is being introduced into an existing fabric running in device alias basic mode, then the device alias mode must be configured to basic mode in the new switch or the device alias mode can be set to enhanced mode for the switches in the existing fabric.

Changing Mode Settings

When a device alias mode is changed from basic mode to enhanced mode, the corresponding applications are informed about the change. The applications then 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 the device alias mode cannot be changed to basic.

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

Device Alias Mode Distribution

If device alias distribution is turned on, it is distributed to the other switches in the network whenever there is a change in the mode.

Device Alias Diffs-Only Distribution

From the Cisco MDS NX-OS Release 7.3(0)D1(1), the Device Alias Diffs-Only Distribution feature is supported on the Cisco MDS switches.

When this feature is enabled on all the switches in a fabric, only the session commands are sent across the fabric instead of the entire database, which helps ensure better scalability.

DDAS supports 20,000 entries when all the switches in a fabric have the Device Alias Diffs-Only Distribution feature enabled. This feature is enabled by default.

Note Ensure that all the switches in a fabric are running a minimum of Cisco MDS NX-OS Release 7.3(0)D1(1) with the Device Alias Diffs-Only feature enabled.

Configuring Device Alias Diffs-Only Distribution

To configure the Device Alias Diff-Only Distribution feature, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **device-alias distribute diffs-only**

Enable the distribution of diffs only on the switch.

This example shows how to enable and display the Device Alias Diffs-Only Distribution feature status on a switch:

Example:

```
switch(config)# device-alias distribute diffs-only
switch(config)# show device-alias status
Fabric Distribution: Enabled
Diffs-only Distribution: Enabled
Database:- Device Aliases 1  Mode: Basic
   Checksum: 0x43a9fe35852e91354543d712c3ec9d3
```

Displaying Device Alias Diffs-Only Distribution Status

This example shows the device alias status during an active session when the Device Alias Diffs-Only Distribution feature is enabled on a switch and in a fabric:

Example:

```
switch(config-device-alias-db)# show device-alias status
Fabric Distribution: Enabled
Diffs-only Distribution: Disabled
Database:- Device Aliases 0 Mode: Basic
Checksum: 0xf6bd6b3389b87233d462029172c8612
Locked By:- User "CLI/SNMPv3:admin" SWWN 20:00:54:7f:ee:1c:2d:40
Pending Database:- Device Aliases 1 Mode: Basic
Diffs-only Distribution capability in the fabric: Enabled

Diffs-only distribution in Session: Enabled
```

This example shows the device alias status during an active session when the Device Alias Diff-Only Distribution feature is disabled on a switch and in a fabric:

Example:

```
switch(config-device-alias-db)# show device-alias status
Fabric Distribution: Enabled
Diffs-only Distribution: Disabled
Database:- Device Aliases 0 Mode: Basic
Checksum: 0xf6bd6b3389b87233d462029172c8612
Locked By:- User "CLI/SNMPv3:admin" SWWN 20:00:54:7f:ee:1c:2d:40
Pending Database:- Device Aliases 1 Mode: Basic
Diffs-only Distribution capability in the fabric: Disabled
SWWN which doesnot support Diffs-only Distribution:
20:00:54:7f:ee:1c:2d:40
20:00:54:7f:e1:1c:2c:40
Diffs-only distribution in Session: Disabled
```

Note The status of *Diffs-only distribution in session* does not change during a session.

Step 3 switch(config)# **no device-alias distribute diffs-only**

Disables Device Alias Diffs-Only Distribution

This example shows how to disable and display the Device Alias Diffs-Only Distribution feature status on a switch:

Example:

```
switch(config)# no device-alias distribute diffs-only
switch(config)# show device-alias status
Fabric Distribution: Enabled
Diffs-only Distribution: Disabled
Database:- Device Aliases 1  Mode: Basic
           Checksum: 0x43a9fe35852e91354543d712c3ec9d3
```

Merging Device Alias with the Diffs-Only Distribution Feature Enabled

Device alias merge failure occurs in the following scenarios:

- When a switch configured with more than 12,000 entries and enabled with the Device Alias Diffs-Only Distribution feature is added to a fabric, that does not support the feature.
- When a switch with disabled Device Alias Diff-Only Distribution feature is added to a fabric, that is configured with more than 12,000 entries and enabled with the Device Alias Diffs-Only feature.

Displaying Merge Failure

This example displays device alias merge failure when one of the fabrics does not support more than 12,000 entries:

```
switch(config)# show cfs merge status name device-alias
Physical-fc Merge Status: Failed [ Wed Jan 20 10:00:34 2016 ]
Failure Reason: One of the merging fabrics cannot support more than 12Kdevice-al
iases
```

Note The Diffs-Only Distribution feature should be enabled on all the switches in a fabric for the device alias entries (more than 12,000) to be supported. If the Diffs-Only Distribution feature is not enabled on all the switches in a fabric, we recommend that you do not configure more than 12,000 entries.

Merging Device Alias in Different Modes

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

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 an application running on an enhanced fabric has a native device alias configuration, the application must fail the merge even if 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. After the device alias merge issue is resolved, each application must be fixed accordingly.

The following issue occurs when there is a device alias database mismatch in the switches that are a part of the same fabric:

The device alias associated to a pWWN is present in the port security/DPVM database even if the respective device alias member is not present in the switch. The device alias associated to a pWWN is missing in the port security/DPVM database even if the respective device alias member is present in the switch.

Resolving Merge Failure 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. 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.



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 will be added when the enhanced mode is turned on and removed when it is turned off. Applications should add confcheck if they have a device alias configuration in the native format, and remove it after the configuration is removed.

Device Alias Features

Device aliases have the following features:

- Device alias information is independent of your VSAN configuration.
- 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.
- 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, you will automatically see that the device aliases are 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.
- Prior to Cisco MDS NX-OS Release 9.2(2), device-alias names were restricted to 64 alphanumeric characters. From Cisco MDS NX-OS Release 9.2(2), device-alias names are restricted to 63 alphanumeric characters. Device-alias names 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)



Note For releases prior to Cisco MDS NX-OS Release 9.2(2), if the device-alias name was 64 characters in length, the DPVM and other application databases do not update properly. Restrict the number of characters in the device-alias name to 63.

Zone Aliases Versus Device Aliases

Table 14: Comparison Between Zone Aliases and Device Aliases , on page 141 compares the configuration differences between zone-based alias configuration and device alias configuration.

Table 14: 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.

Zone-Based Aliases	Device Aliases
FC aliases are not displayed with the associated WWNs in the show command outputs like show zoneset active, show flogi database, and show fcns database.	Device aliases are displayed with the associated WWNs in the show command outputs like show zoneset active, show flogi database, and show fcns database.
FC aliases are not distributed as part of active zoneset and are only distributed as part of full zone database as per the FC standards.	Device Aliases are distributed through CFS.

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.
- 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.

This section includes the following sections:

Creating Device Aliases

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

Procedure

Step 1 switch# **config t**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **device-alias database**

switch(config-device-alias-db)#

Enters the pending database configuration submode.

Step 3 switch(config-device-alias-db)# **device-alias name Device1 pwwn 21:01:00:e0:8b:2e:80:93**

Specifies a device name (Device1) for the device that is identified by its pWWN. Starts writing to the pending database and simultaneously locks the fabric as this is the first-issued device alias configuration command.

Step 4 switch(config-device-alias-db)# **no device-alias name Device1**

Removes the device name (Device1) for the device that is identified by its pWWN.

Step 5 switch(config-device-alias-db)# **device-alias rename Device1 Device2**

Renames an existing device alias (Device1) with a new name (Device2).

To display the device alias configuration, use the **show device-alias name** command.

```
switch# show device-alias name x
device-alias name x pwwn 21:01:00:e0:8b:2e:80:93
```

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.

Displays a Failed Status

```
switch# show
 device-alias status
Fabric Distribution: Disabled
Database:- Device Aliases 25
Status of the last CFS operation issued from this switch:
============================================================
Operation: Commit
Status: Failed (Reason: Operation is not permitted as the fabric distribution is
 currently disabled.)
```



Note From the Cisco MDS NX-OS Release 6.2.9 onwards, the ASCII configuration replay takes longer time for DDAS (Distributing Device Alias Services) without the write erase command.

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.

About Device Alias Configuration Best Practices

As a part of the device-alias configuration best practices, the following guidelines need to be adopted within a device-alias session:

If a device-alias name is reused while configuring a rename command, then the command fails and gets moved to the rejected list.

Displays the rejected device-alias command

```
switch(config-device-alias-db)# device-alias name dev10 pwwn 10:10:10:10:10:10:10:10
switch(config-device-alias-db)# device-alias rename dev10 new-dev10
Command rejected. Device-alias reused in current session :dev10
Please use 'show device-alias session rejected' to display the rejected set of commands and
 for the device-alias best-practices recommendation.
switch(config-device-alias-db)#
```

If a PWWN is reused while configuring an add or delete command, then the command fails and gets moved to the rejected list.

Displays the rejected device-alias command

```
switch(config-device-alias-db)# device-alias name dev11 pwwn 11:11:11:11:11:11:11:11
switch(config-device-alias-db)# no device-alias name dev11
Command rejected. Pwwn reused in current session: 11:11:11:11:11:11:11:11 is mapped to
device-alias dev11
Please use 'show device-alias session rejected' to display the rejected set of commands and
 for the device-alias best-practices recommendation.
switch(config-device-alias-db)#
```

If a device-alias name is reused in an add command which was earlier being renamed in a rename command, the command fails and gets moved to the rejected list.

```
switch(config-device-alias-db)# device-alias rename da3 new-da3
switch(config-device-alias-db)# device-alias name da3 pwwn 2:2:2:2:3:3:3:3
Command rejected. Device-alias name reused in current session: da3
Please use 'show device-alias session rejected' to display the rejected set of commands and
 for the device-alias best-practices recommendation.
switch(config-device-alias-db)#
```

Displays the rejected device-alias command

The rejected set of commands can be displayed using the show device-alias session rejected command.

```
switch(config-device-alias-db)# show device-alias session rejected
To avoid command rejections, within a device alias session
Do not reuse:
a) a device alias name while configuring a rename command
b) a PWWN while configuring an add or delete command
c) a device alias name already renamed while configuring add command

Rejected commands must be committed in a separate device alias session
which may cause traffic interruption for those devices. Plan accordingly.
Refer to this command in the NX-OS Command Reference Guide
for more information about device alias configuration best practices

Rejected Command List
---------------------
device-alias rename dev10 new-dev10
no device-alias name dev11
device-alias name da3 pwwn 02:02:02:02:03:03:03:03
switch(config-device-alias-db)# #
```

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.

To commit the changes, follow these steps:

Procedure

Step 1 switch# **config terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **device-alias commit**

Commits the changes made to the currently active session.

Whenever a switch in the fabric gets locked and goes for a blank commit, the following warning is displayed:

```
WARNING: Device-alias DB is empty in this switch.
Initiating a commit from this switch will clear [wipe out] Device-alias DB across all the
switches in the fabric, losing Device-alias full DB config permanently.
Do you want to continue? (y/n) [n]
```

Note After the **device-alias commit** is complete, the running configuration is modified on all switches participating in device-alias distribution. You can then use the **copy running-config startup-config fabric** command to save the running-config to the startup-config on all the switches in the fabric.

Step 3 switch(config)# **device-alias commit force**

Commits the changes forcefully and overwrite the changes made to the currently active session.

Enabling the Device Alias Pending Diff Display

To enable the display of the pending-diff and the subsequent confirmation on issuing a device-alias commit, follow these steps:

Procedure

Step 1 switch# **config t**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# device-alias confirm-commit

Enables the confirm commit option for device- alias.

Step 3 switch(config)# device-alias commit

```
The following device-alias changes are about to be committed
+ device-alias name Device1 pwwn 21:01:00:e0:8b:2e:80:93
Do you want to continue? (y/n) [n] y
```

If the device-alias confirm-commit command is enabled, on committing the pending database, the pending-diff is displayed on the console and user is prompted for Yes or No. If the device -alias confirm-commit command is disabled, the pending-diff is not displayed and the user is not prompted for Yes or No.

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.

To discard the device alias session, perform this task:

Procedure

Step 1 switch# **config terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **device-alias abort**

Discards the currently active session.

To display the status of the discard operation, use the show **device alias status** command.

```
switch# show
 device-alias status
Fabric Distribution: Enabled
Database:- Device Aliases 24
Status of the last CFS operation issued from this switch:
=========================================================
Operation: Abort
Status: Success
```

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.

To clear device-alias session, use the **clear device-alias session** command in CONFIGURATION mode.

```
switch(config)# clear device-alias session
```

To verify the status of the clear operation, use the **show device-alias session status** command.

```
switch(config)# show device-alias session status
Last Action Time Stamp     : None
Last Action                : None
Last Action Result         : None
Last Action Failure Reason : none
```

Clearing Database Content

To clear all the database content, use the **clear device-alias database** command in CONFIGURATION mode.

```
switch(config)# clear device-alias database
To verify the status of the clear device-alias database
command, use the show device-alias database
command.
switch(config)# show device-alias database
```

Clearing Statistics

To clear all the statistics, use the **clear device-alias statistics** command in CONFIGURATION mode.

```
switch# clear device-alias statistics
```

Disabling and Enabling Device Alias Distribution

To disable or enable the device alias distribution, follow these steps:

Procedure

Step 1 switch# **config t**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **no device-alias distribute**

Disables the distribution.

Step 3 switch(config)# **device-alias distribute**

Enables the distribution (default).

To display the status of device alias distribution, use the **show device-alias status** command (see the following examples).

Displays Device Alias Status When Distribution Is Enabled

Displays Device Alias Status When Distribution Is Disabled

```
switch# show
device-alias status
Fabric Distribution: Enabled <-------------------------------Distribution is enabled
Database:-Device Aliases 24
Locked By:-User "Test" SWWN 20:00:00:0c:cf:f4:02:83<-Lock holder's user name and switch ID
Pending Database:- Device Aliases 24
Status of the last CFS operation issued from this switch:
=========================================================
Operation: Enable Fabric Distribution
Status: Success
```

```
switch# show
 device-alias status
Fabric Distribution: Disabled
Database:- Device Aliases 24
Status of the last CFS operation issued from this switch:
=========================================================
Operation: Disable Fabric Distribution
Status: Success
```

About Legacy Zone Alias Configuration Conversion

You can import legacy zone alias configurations to use this feature without losing data, if they satisfy 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.

Tip Ensure to 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.

This section includes the following topics:

Importing a Zone Alias



Note Device alias does not allow importing and manually adding of device alias entries to the database in the same session.

To import the zone alias for a specific VSAN, follow these steps:

SUMMARY STEPS

1. switch# **config t**
2. switch(config)# **device-alias import fcalias vsan 3**

DETAILED STEPS

Procedure

	Command or Action	Purpose

Step 1	switch# **config t**
	Example:
	`switch# config t`
	`switch(config)#`
		Enters configuration mode.

Step 2	switch(config)# **device-alias import fcalias vsan 3**	Imports the fcalias information for the specified VSAN. To display device alias information in zone sets, use the **show zoneset** command (see the following examples).

Displays the Device Aliases in the Zone Set Information

```
switch# show zoneset
zoneset name s1 vsan 1
  zone name z1 vsan 1
    pwwn 21:01:00:e0:8b:2e:80:93 [x] <---------------Device alias displayed for each pWWN.
    pwwn 21:00:00:20:37:39:ab:5f [y]
zone name z2 vsan 1
    pwwn 21:00:00:e0:8b:0b:66:56 [SampleName]
    pwwn 21:00:00:20:37:39:ac:0d [z]
```

Example: Displays the Device Aliases in the Active Zone Set

```
switch# show zoneset active
zoneset name s1 vsan 1
  zone name z1 vsan 1
  * fcid 0x670100 [pwwn 21:01:00:e0:8b:2e:80:93] [x]
    pwwn 21:00:00:20:37:39:ab:5f [y]
  zone name z2 vsan 1
  * fcid 0x670200 [pwwn 21:00:00:e0:8b:0b:66:56] [SampleName]
    pwwn 21:00:00:20:37:39:ac:0d [z]
```

Device Alias Statistics Cleanup

Use the **clear device-name statistics** command to clear device alias statistics (for debugging purposes):

```
switch# clear device-alias statistics
```

Database Merge Guidelines

For information about CFS merge support, refer to the *Cisco MDS 9000 Series 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.
- Ensure the device -alias mode is similar for the both the fabrics being merged.

Device Alias Configuration Verification

You can view device alias information by using the **show device-alias** command. See the following examples.

Displays All Configured Device Aliases from the Effective Database

```
switch# show
device-alias database
device-alias name SampleName pwwn 21:00:00:e0:8b:0b:66:56
device-alias name x pwwn 21:01:00:e0:8b:2e:80:93
Total number of entries = 2
```

Displays the Pending Database with No Modifications

```
switch# show
device-alias database pending
There are no pending changes
```

Displays the Pending Database with Modifications

```
switch# show
device-alias database pending
device-alias name x pwwn 21:01:00:e0:8b:2e:80:93
device-alias name SampleName pwwn 21:00:00:e0:8b:0b:66:56
device-alias name y pwwn 21:00:00:20:37:39:ab:5f
device-alias name z pwwn 21:00:00:20:37:39:ac:0d
Total number of entries = 4
```

Displays the Specified Device Name in the Pending Database

```
switch# show
device-alias name x pending
device-alias name x pwwn 21:01:00:e0:8b:2e:80:93
```

Displays the Specified pWWN in the Pending Database

```
switch# show
device-alias pwwn 21:01:00:e0:8b:2e:80:93 pending
device-alias name x pwwn 21:01:00:e0:8b:2e:80:93
```

Displays the Difference Between the Pending and Effective Databases

```
switch# show
device-alias database pending-diff
- device-alias name Doc pwwn 21:01:02:03:00:01:01:01
+ device-alias name SampleName pwwn 21:00:00:e0:8b:0b:66:56
```

Displays the Specified pWWN

```
switch# show
device-alias pwwn 21:01:01:01:01:11:01:01
device-alias name Doc pwwn 21:01:01:01:01:11:01:01
```

Displays the Device Alias in the FLOGI Database

```
switch# show flogi database
--------------------------------------------------------------------------------
INTERFACE  VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------------
fc2/9      1     0x670100  21:01:00:e0:8b:2e:80:93  20:01:00:e0:8b:2e:80:93
                           [x
] <-----------------------------------------Device alias name
fc2/12     1     0x670200  21:00:00:e0:8b:0b:66:56  20:00:00:e0:8b:0b:66:56
                           [SampleName
] <--------------------------------Device alias name
Total number of flogi = 2
```

Displays the Device Alias in the FCNS Database

```
switch# show fcns database
VSAN 1:
--------------------------------------------------------------------------
```

```
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPE:FEATURE
------------------------------------------------------------------------------
0x670100    N     21:01:00:e0:8b:2e:80:93 (Qlogic)        scsi-fcp:init
                  [x
]
0x670200    N     21:00:00:e0:8b:0b:66:56 (Qlogic)        scsi-fcp:init
                  [SampleName
]
Total number of entries = 2
```

Displays the fcping Statistics for the Specified Device Alias

```
switch# fcping device-alias x vsan 1
28 bytes from 21:01:00:e0:8b:2e:80:93 time = 358 usec
28 bytes from 21:01:00:e0:8b:2e:80:93 time = 226 usec
28 bytes from 21:01:00:e0:8b:2e:80:93 time = 372 usec
```

Displays the fctrace Information for the Specified Device Alias

```
switch# fctrace device-alias x vsan 1
Route present for : 21:01:00:e0:8b:2e:80:93
20:00:00:05:30:00:4a:e2(0xfffc67)
```

Where available, device aliases are displayed regardless of a member being configured using a **device-alias** command or a zone-specific **member pwwn** command.

Displays Statistics for the Device Alias Application

```
switch# show
device-alias statistics
        Device Alias Statistics
========================================
Lock requests sent: 2
Database update requests sent: 1
Unlock requests sent: 1
Lock requests received: 1
Database update requests received: 1
Unlock requests received: 1
Lock rejects sent: 0
Database update rejects sent: 0
Unlock rejects sent: 0
Lock rejects received: 0
Database update rejects received: 0
Unlock rejects received: 0
Merge requests received: 0
Merge request rejects sent: 0
Merge responses received: 2
Merge response rejects sent: 0
Activation requests received: 0
Activation request rejects sent: 0
Activation requests sent: 2
Activation request rejects received: 0
```

Default Settings

Table 15: Default Device Alias Parameters , on page 153 lists the default settings for device alias parameters.

Table 15: 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.

Resolving Device Alias Merge Failures

The most common device-alias merge failure issues occur when merging databases. When a device-alias merge fails, we recommend that you review the syslog messages on the switch in which the merge was initiated in order to identify the issues. The application server in each fabric that is responsible for the merge is indicated by the term Merge Master in the messages.

In this example, the syslog messages indicate that the merge failed as a result of a database mismatch:

```
2007 Apr  9 15:52:42 switch-1 %CFS-3-MERGE_FAILED: Merge failed for app device-alias, local
 switch wwn 20:00:00:0d:ec:2f:c1:40,ip 172.20.150.38, remote switch wwn
20:00:00:0d:ec:04:99:40, ip 172.20.150.30
2007 Apr  9 15:52:42 switch-1 %DEVICE-ALIAS-3-MERGE_FAILED: Databases could not be merged
due to mismatch.
```

Note Use the **device-alias distribute** command to initiate a merge or remerge of device-alias databases. Use the **device-alias commit** command to *push* a switch's device-alias database to all the other switches in a fabric. If the switches whose device-alias databases are not merged (more than one merge master is shown in the output of the **show cfs merge status name device-alias** command), then the **device-alias commit** command causes the device-alias databases that are not merged to be overwritten.

Connecting an MDS switch running a version of NX-OS that is prior to Cisco MDS NX-OS Release 9.2(2) with an MDS switch running Release 9.2(2) may encounter a device-alias and zone merge failure. Consider two switches where switch 1 is running Cisco MDS NX-OS Release 9.2(2) or later release and switch 2 is running a release prior to Cisco MDS NX-OS Release 9.2(2). Both the switches are in device alias enhanced mode. Switch 2 has one or more device-alias names that are configured with 64 alphanumeric characters. Both zone and device alias merge will fail between switch 1 and switch 2 because switch 2 is using a device-alias name that is configured using 64 alphanumeric characters. In such a case, we recommend that all device-alias names configured with 64 alphanumeric characters be reconfigured to 63 alphanumeric characters or less. After that is done, the device alias database should be remerged using the no device-alias distribute command followed by the device-alias distribute command. The zoneset then can be remerged and the VSAN be brought out of isolation on the ISL by removing the VSAN from the ISL, readding it or shutting down the ISL, and bringing it back up again if it is only transporting a single VSAN.

Device Alias Best Practices

This section lists the best practices that you should follow when creating and using device aliases:

- Device aliases should be used to simplify the management of world wide names (WWNs) whenever possible. It is easier to identify devices with aliases rather than with WWNs. Hence, you should assign aliases to WWNs to easily identify the WWNs.
- Device-alias names are case-sensitive.
- Operate device aliases in Enhanced mode whenever possible. In Enhanced mode, applications accept a device-alias name in its *native* format, rather than expanding the alias to a port world wide name (pWWN). Because applications such as zone server, Inter-VSAN Routing (IVR), Port Security Manager (PSM), and Dynamic Port VSAN Membership automatically track and enforce device-alias membership changes, you have a single point of change.



Note Interop mode VSANs do not accept Enhanced mode configurations.

- Preplan device-alias configurations and implement a consistent naming convention.
- Keep documented backups of all device-alias configurations.
- Plan for what the final device-alias database should be after the merge, before attempting to resolve merge failures. This can prevent traffic disruptions caused by accidentally overwriting device-alias entries.



Caution Avoid performing a *blank commit* to resolve Cisco Fabric Services (CFS) merge failures. A blank commit overwrites the device-alias databases on all the switches with the device-alias database on the local switch.



Note A blank commit is a device-alias commit that is used when there are no changes (including mode changes), or when it is okay to overwrite the device-alias databases on the remote switches with the local switch's device-alias database.

Device alias mismatches might occur because of the following reasons:

- Duplicate Device-Alias Names—Same device-alias name, but different pWWNs. In such a scenario, the **show device-alias merge status** command displays the reason for the merge failure as `Reason: Another device-alias already present with the same name.`
- Duplicate pWWNs—Different device-alias names, but same pWWN. In such a scenario, the **show device-alias merge status** command displays the reason for the merge failure as `Reason: Another device-alias already present with the same pwwn.`

Note Each time device-alias changes are committed, the running configuration should be copied to the startup configuration on all the switches that were updated. Use the **copy running-config startup-config fabric** command to copy the running configuration to the startup configuration for all the switches in the fabric. If you do not copy the running configuration to the startup configuration after the device-alias changes are committed, and if the switch reloads, or loses power and restarts, the startup configuration will not have the correct device-alias database and merge failure will occur.

- You will be unable to upgrade to Cisco MDS NX-OS Release 9.2(2) or later releases if you have configured any device-alias names using 64 alphanumeric characters. For more information, see the *Cisco MDS 9000 NX-OS Software Upgrade and Downgrade Guide*.

Resolving Device Alias Mismatches

If a switch with an existing device-alias database is being added to an existing fabric, conflicts might arise because of the following reasons:

- The same device-alias name is used, but with different pWWNs.
- The same pWWN is used, but with different device-alias names.

To resolve duplicate device-alias names, perform these steps:

Procedure

Step 1 Run the **show cfs merge status name device-alias** command to review the CFS or device-alias merge failure syslogs to confirm that the merge failed:

```
switch-1# show cfs merge status name device-alias

Physical-fc Merge Status: Failed
[Sun Sep 25 14:45:55 2016]
Failure Reason: Another device-alias already present with the same pwwn

Local Fabric
-------------------------------------------------------------------------------
Switch WWN              IP Address
-------------------------------------------------------------------------------
20:00:54:7f:ee:1b:0e:b0 10.127.103.211        [Merge Master] <<< Merge Master#1
                         [switch-1]

Total number of switches = 1

Remote Fabric
-------------------------------------------------------------------------------
Switch WWN              IP Address
-------------------------------------------------------------------------------
20:00:54:7f:ee:1b:0e:50 10.197.111.54         [Merge Master] <<< Merge Master#2
```

```
Total number of switches = 1
```

Note A properly merged device-alias application should only show a single merge master. If there is more than one merge master, as shown in the above example, it indicates that the device-alias databases are not merged.

Step 2 Use the **no device-alias distribute** command on the switch in which the merge failure occurred in order to disable the device-alias distribution:

```
switch-1# configure terminal
switch-1(config)# no device-alias distribute
```

Step 3 Resolve merge failure on the switch. See Resolving Merge Failures, on page 156 section.

Resolving Merge Failures

This section provides information about how to resolve merge failures.

Resolving Duplicate Device Alias Names (Same Device Alias Name, Different pWWNs)

Note A device-alias name is considered to be duplicate when the same device-alias name is used to point to different pWWNs.

To verify if a duplicate device-alias name exists in fabrics, perform these steps:

Procedure

Step 1 Run the **show device-alias merge status** command to identify if the reason for the merge failure is a database mismatch:

```
switch# show device-alias merge status
     Result: Failure
     Reason: Another device-alias already present with the same name
```

Note A properly merged device-alias application should only show a single merge master. If there is more than one merge master, as shown in the above example, it indicates that the device-alias databases are not merged.

Step 2 Review the CFS or the device-alias merge failure syslog to confirm that the merge failed. Alternatively, run the **show cfs merge status name device-alias** command to view the status of the merge:

```
switch# show cfs merge status name device-alias
Physical-fc  Merge Status: Failed [ Mon Apr  9 15:57:58 2007 ] <===Merge status
      Local Fabric
     ---------------------------------------------------------------------
      Switch WWN                IP Address
     ---------------------------------------------------------------------
      20:00:00:0d:ec:2f:c1:40  172.20.150.38       [Merge Master] <<< Merge Master#1
                               switch-1
     Total number of switches = 1

      Remote Fabric
     ---------------------------------------------------------------------
      Switch WWN                IP Address
     ---------------------------------------------------------------------
      20:00:00:0d:ec:04:99:40  172.20.150.30       [Merge Master] <<< Merge Master#2
                               switch-2
     Total number of switches = 1
```

Step 3 Depending on the Cisco MDS NX-OS release your switch is using, run one of the following commands:

- Cisco MDS NX-OS Release 8.1(1) and later releases

  Run the **show device-alias merge conflicts** command to display the device alias and pWWNs that are causing the merge failure.

  Note Run the **show device-alias merge conflicts** command from a switch listed as a merge master.

  In the following example, the same device-alias name, A1, is assigned to two different pWWNs—a pWWN on a local switch and a pWWN on a peer switch:

```
switch-1# show device-alias merge conflicts
Merge Status : Failure
Peer Switch SWWN : 20:00:00:0d:ec:24:f5:00
Conflicts :
1. Conflicting Pwwns : 1
---------------------------------------------------------------------------
Local PWWN     Peer PWWN    Device-alias
---------------------------------------------------------------------------
pwwn 0:01:01:01:01:01:01:02  pwwn :01:01:01:01:01:01:03  A1
```

- Cisco MDS NX-OS Release 7.3 and earlier releases

  Compare the device-alias databases manually to identify the duplicate device-alias names.

  In the following example, the same device-alias name, A1, is assigned to two different pWWNs—a pWWN on a local switch and a pWWN on a peer switch.

  From merge master#1:

```
switch-1# show device-alias database
...output trimmed to show only mismatched device-alias
device-alias name A1 pwwn 21:01:01:01:01:01:01:02

switch-2# show device-alias database
...output trimmed to show only mismatched device-alias
```

```
device-alias name A1 pwwn 21:01:01:01:01:01:01:03
```

Step 4 Run the **device-alias name** *name* **pwwn** *id* command to change the pWWN on one of the switches to match the pWWN on the other switch.

Note Perform this step after device-alias distribution is disabled by running the **no device-alias distribute** command.

In the following example, the pWWN 21:01:01:01:01:01:01:02 on switch-1 is changed to match the pWWN 21:01:01:01:01:01:01:03 on switch-2:

```
switch-1# configure
Enter configuration commands, one per line. End with CNTL/Z.
switch-1(config)# device-alias database
switch-1(config-device-alias-db)# no device-alias name A1
switch-1(config-device-alias-db)# show device-alias database | i A1
switch-1(config-device-alias-db)# device-alias name A1 pwwn 21:01:01:01:01:01:01:03
switch-1(config-device-alias-db)# show device-alias database | i A1
device-alias name A1 pwwn 21:01:01:01:01:01:01:03
```

Step 5 If there are more duplicate device-alias names, perform step Step 3, on page 157 and step Step 4, on page 158 to resolve the duplicate device-alias names issue.

Step 6 Use the **device-alias distribute** command to enable the device-alias distribution and initiate a merge:

```
switch-1(config)# device-alias distribute
```

Step 7 Use the **show cfs merge status name device-alias** command to verify in the output if the merge was successful.

Resolving Duplicate pWWNs (Different Device Alias Names, Same pWWN)

To verify that the same pWWN is mapped to different device-alias names in fabrics, perform these steps:

Procedure

Step 1 Run the **show device-alias merge status** command to identify if the reason for the merge failure is a database mismatch:

```
switch# show device-alias merge status
     Result: Failure
Reason: Another device-alias already present with the same pwwn.
```

Note A properly merged device-alias application should only show a single merge master. If there is more than one merge master, as shown in the above example, it indicates that the device-alias databases are not merged.

Step 2 Review the CFS or the device-alias merge failure syslog to confirm that the merge failed. Alternatively, run the **show cfs merge status name device-alias** command to view the status of the merge:

```
switch# show cfs merge status name device-alias
Physical-fc  Merge Status: Failed [ Mon Apr  9 15:57:58 2007 ] <===Merge status
      Local Fabric
     ------------------------------------------------------------------------
      Switch WWN              IP Address
     ------------------------------------------------------------------------
      20:00:00:0d:ec:2f:c1:40  172.20.150.38       [Merge Master] <<< Merge Master#1
                               switch-1
     Total number of switches = 1

      Remote Fabric
     ------------------------------------------------------------------------
      Switch WWN              IP Address
     ------------------------------------------------------------------------
      20:00:00:0d:ec:04:99:40  172.20.150.30       [Merge Master] <<< Merge Master#2
                               switch-2
Total number of switches = 1
```

Step 3 Depending on the Cisco MDS NX-OS release your switch is using, run one of the following commands:

- Cisco MDS NX-OS Release 8.1(1) and later releases

  Use the **show device-alias merge conflicts** command to display the device alias and pWWNs that are causing a merge failure. Use the **no device-alias distribute** command, followed by the **device-alias distribute** command to update the information about the merge conflicts.

  Note Run the **show device-alias merge conflicts** command from a switch listed as a merge master.

  In the following example, the pWWN 21:01:01:01:01:01:01:02 is mapped to device-alias A3 on switch-1, and to device-alias A1 on switch-2:

```
switch-1# show device-alias merge conflicts
Merge Status : Failure
Peer Switch SWWN : 20:00:00:0d:ec:24:f5:00
Conflicts :
1. Conflicting Device-aliases : 1
------------------------------------------------------------------------
Local Device-alias  Peer Device-alias   PWWN
------------------------------------------------------------------------
A3  A1   pwwn    21:01:01:01:01:01:01:02
```

- Cisco MDS NX-OS Release 7.3 and earlier releases

  Compare the device-alias databases manually to identify the pWWNs that are causing a merge failure.

  On the switches where the merge failed in step Step 1, on page 158, use the **show device-alias database** command to identify if a pWWN that is mapped to two different device-alias names exists.

  In this example, the pWWN 21:01:01:01:01:01:01:02 is mapped to the device-alias A3 on switch-1 and to the device-alias A1 on switch-2:

```
switch-1# show device-alias database
device-alias name A3 pwwn 21:01:01:01:01:01:01:02
Total number of entries = 1

switch-2# show device-alias database
device-alias name A1 pwwn 21:01:01:01:01:01:01:02
```

Step 4 Run the **device-alias name** *name* **pwwn** *id* command to change the device-alias name on one of the switches to match the device-alias name on the other switch.

Note Perform this step after device-alias distribution is disabled by running the **no device-alias distribute** command.

In the following example, the device-alias name A3 on switch-1 is changed to match the device-alias name A1 on switch-2:

```
switch-1# configure
Enter configuration commands, one per line. End with CNTL/Z.
switch-1(config)# device-alias database
switch-1(config-device-alias-db)# no device-alias name A3
switch-1(config-device-alias-db)# device-alias name A1 pwwn 21:01:01:01:01:01:01:02
```

Step 5 If there are more duplicate device-alias names, perform step Step 3, on page 159 and step Step 4, on page 160 to resolve the duplicate device-alias names issue.

Step 6 Use the **device-alias distribute** command to enable the device-alias distribution and initiate a merge:

```
switch-1(config)# device-alias distribute
```

Step 7 Use the **show cfs merge status name device-alias** command to verify in the output if the merge was successful.

Resolving Mode Mismatch

The Device Alias feature can operate in either Basic or Enhanced mode. If the modes are different in two fabrics, CFS merge between the fabrics fails.

To verify that the device-alias mode is different in two fabrics, perform these steps:

Procedure

Step 1 Review the CFS or device-alias merge failure syslog to confirm that the merge failed. Alternatively, run the **show cfs merge status name device-alias** command to view the status of the merge.

```
switch# show cfs merge status name device-alias
Physical-fc  Merge Status: Failed [ Mon Apr  9 15:57:58 2007 ] <===Merge status
      Local Fabric
     ------------------------------------------------------------------------
      Switch WWN              IP Address
```

```
    ------------------------------------------------------------------------
     20:00:00:0d:ec:2f:c1:40  172.20.150.38       [Merge Master] <<< Merge Master#1
                             switch-1
    Total number of switches = 1
     Remote Fabric
    ------------------------------------------------------------------------
     Switch WWN               IP Address
    ------------------------------------------------------------------------
     20:00:00:0d:ec:04:99:40  172.20.150.30       [Merge Master] <<< Merge Master#2
                             switch-2
  Total number of switches = 1
```

Step 2 Use the **show device-alias merge status** command to verify that the reason for the merge failure is a mode mismatch. If there is a mode mismatch, the reason that is displayed in the output is either "`Databases could not be merged due to mode mismatch`" or "`One of the merging fabrics cannot support device-alias Enhanced mode.`"

```
switch# show device-alias merge status
     Result: Failure
     Reason: Databases could not be merged due to mode mismatch.
```

Step 3 Use the **show device-alias status** command to verify the device-alias mode for each of the fabric.

In this example, switch-1 is running in Enhanced mode, while switch-2 is running in Basic mode:

```
switch-1# show device-alias status
Fabric Distribution: Enabled
Database:- Device Aliases 2 Mode: Enhanced

switch-2# show device-alias status
     Fabric Distribution: Enabled
     Database:- Device Aliases 2 Mode: Basic
```

Step 4 Use the **no device-alias distribute** command to disable device-alias distribution after you detect mismatched device-alias modes.

Step 5 Depending on the mode you want to change in the switch, use the **device-alias mode enhanced** command to change to enhanced mode, or use the **no device-alias mode enhanced** command to change the switch mode to basic mode.

Note

- Prior to Cisco MDS NX-OS Release 8.5(1), the default device alias mode was basic mode. From Cisco MDS NX-OS Release 8.5(1), the default device alias mode is enhanced mode.
- If you want to change the device-alias mode from Enhanced to Basic, but an application contains a device-alias configuration in the native format, the device-alias mode cannot be changed until you explicitly remove all the native device-alias configurations or replace all the device-alias members with the corresponding pWWNs.

Step 6 Use the **device-alias distribute** command to enable the device-alias distribution and initiate a merge.

Resolving a Validation Failure

If the merger of device aliases takes place without any conflicts, the resultant device-alias database is validated with the registered applications on each switch in both the fabrics being merged. If an application fails the validation of the merged database for any reason, the device-alias merge fails.

To verify that the device-alias database merge is failing because of an application-validation failure, perform these steps:

Procedure

Step 1 Review the CFS or device-alias merge failure syslog to confirm that the merge failed. Alternatively, run the **show cfs merge status name device-alias** command to view the status of the merge.

Step 2 Use the **show device-alias merge status** command to verify that the reason for the merge failure is an application-validation failure:

```
switch# show device-alias merge status
     Result: Failure
     Reason: This is a non device-alias error.
```

Step 3 Examine the syslog messages. The syslog for the switch in which the validation is rejected and the syslog for the switch managing the merge show relevant error messages.

This example shows a sample message on a switch in which the validation is rejected:

```
2007 Apr 10 00:00:06 switch-2 %DEVICE-ALIAS-3-MERGE_VALIDATION_REJECTED:
Failed SAP: 110 Reason: inter-VSAN zone member cannot be in more than one
VSAN Expln:
```

This example shows the syslog message on a switch that is managing the merge, and in which the validation is rejected:

```
2007 Apr  9 16:41:22 switch-1 %DEVICE-ALIAS-3-MERGE_VALIDATION_FAILED: Failed
SWWN: 20:00:00:0d:ec:04:99:40 Failed SAP: 110 Reason: inter-VSAN zone member cannot be in more than
 one
VSAN Expln:
```

Step 4 Use the **show device-alias internal validation-info** command on the switch managing the merge, and examine the output.

This example shows that SAP 110 on switch 20:00:00:0d:ec:04:99:40 (switch-2) rejected the validation. The status message shows the reason for the failure along with the system application number:

```
switch# show device-alias internal validation-info
     Validation timer:    0s
 Per SAP Info Table:
    ==================
      SAPS:  0
     MTS Buffer Array Details:
     ========================
      Buffers:  0
     Local Status:
     =============
      Num Reqs Sent:  0 20:00:00:0d:ec:04:99:40
```

```
       Num SAPs Done:  0
       Failed SAP   :  0    Status: success    Expln:
     Remote Status:
     ==============
       CFS Resp Rcvd: TRUE
       Failed SWWN  : 20:00:00:0d:ec:04:99:40
SAP : 110 Status: inter-VSAN zone member cannot be in more than one VSAN <=== Status
       Expln:
```

Step 5 Use the **show system internal mts sup sap number description** command to find the application that rejected the configuration on the switch that rejected the validation.

In this example, the application that rejected the device-alias validation was the IVR process.

```
switch# show system internal mts sup sap 110 description
IVR-SAP
```

Step 6 Analyze the device-alias validation failure. This analysis is dependent on the application that failed the validation as well as the device-alias database configuration.

In this example, IVR is failing the validation. To troubleshoot this problem, begin by reviewing the device-alias databases that are being merged. Use the **show device-alias database** command from the switch managing the merge for each fabric.

```
switch# show device-alias database
device-alias name A1 pwwn 21:01:01:01:01:01:01:01
device-alias name A2 pwwn 21:01:01:01:01:01:01:02 => Pre-merge: A2 defined on switch-1
Total number of entries = 2

switch# show device-alias database
device-alias name A1 pwwn 21:01:01:01:01:01:01:01 => Pre-merge: A2 not defined on switch-2
Total number of entries = 1
Because IVR is enabled on switch-2, review the IVR zone set.
switch# show ivr zoneset
zoneset name s1
   zone name z1
       pwwn 21:01:01:01:01:01:01:02 vsan    1 autonomous-fabric-id  1
       device-alias A2              vsan    2 autonomous-fabric-id  1
```

Prior to the database merge, device-alias A2 is not defined on switch-2. Because of the merge between switch-1 and switch-2, device-alias A2 becomes available on switch-2, and A2 is mapped to pWWN 21:01:01:01:01:01:01:02.

The device alias-based member A2 in the IVR zone z1 is resolved and mapped to pWWN 21:01:01:01:01:01:01:02, and is a member of VSAN 2. However, pWWN 21:01:01:01:01:01:01:02 is already a member of VSAN 1. The mapping that occurs because of the device-alias merge makes the IVR configuration illegal. The same pWWN cannot be a member of multiple VSANs.

In the case when IVR configuration is illegal, the pWWN in VSAN 2 is defined using the device alias (A2), while the member in VSAN 1 is defined using the actual pWWN. The IVR detects this situation and rejects the device-alias validation. As a result, the device-alias merge fails.

Resolving Database Conflicts

If an entry in the device-alias database conflicts with the configuration of a registered application, the device-alias database commit fails the validation process. Correct either the device-alias database or the application configuration.

To determine the application that failed the validation and the reason for the failure, perform these steps:

Procedure

Step 1 Use the **device-alias commit** command to view the output.

This example shows that the commit failed because there is a conflict between the device-alias database and an application configuration:

```
switch# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# device-alias commit
inter-VSAN zone member cannot be in more than one VSAN ===> reason for commit failure
```

Step 2 Determine which application configuration is in conflict with the device-alias database by reviewing the syslogs for the switch in which the commit was issued.

This example shows that SAP 110 (IVR) on sWWN 20:00:00:0d:ec:04:99:40 (switch-2) has rejected the validation, and therefore, the device-alias commit has failed:

```
2007 Apr 10 11:54:24 switch-1 %DEVICE-ALIAS-3-VALIDATION_FAILED: Failed=>Validation Status
SWWN: 20:00:00:0d:ec:04:99:40 Failed SAP: 110 Reason: inter-VSAN zone ==>Switch and SAP member cannot
 be in more than one VSAN Expln:                          ==>Reason
2007 Apr 10 11:54:24 switch-1 %DEVICE-ALIAS-3-COMMIT_FAILED: Failed to  ==>Commit status commit the
 pending database: inter-VSAN zone member cannot be in more ==>Reason than one VSAN
```

Step 3 Review the syslog on the switch in which the validation is rejected.

This example shows that the following syslog is printed on switch-2:

```
2007 Apr 10 19:13:08 switch-2 %DEVICE-ALIAS-3-VALIDATION_REJECTED: Failed
SAP: 110 Reason: inter-VSAN zone member cannot be in more than one VSAN ==>SAP and reason
```

Step 4 Compare the existing device-alias database (including the desired changes) and the application configuration to find the conflict.

This example uses the **show device-alias database** and **show ivr zoneset** commands along with the console logs of the device-alias database changes made prior to the commit. The comparison shows that the definition of the new device-alias A2 results in the resolution of the enhanced device-alias member A2 in the IVR zone z1 to pWWN 21:01:01:01:01:01:01:02, which is already a member of zone z1. The pWWN is directly defined as a member of VSAN 1, while the enhanced device-alias A2 is defined as a member of VSAN 2. This configuration is not allowed in the IVR. The IVR detects the configuration problem and rejects the device-alias database validation.

```
switch# show device-alias database          ===> existing device alias database
device-alias name A1 pwwn 21:01:01:01:01:01:01:01
Total number of entries = 1
```

```
switch# show ivr zoneset                    ===> display existing IVR zone set
zoneset name s1
zone name z1
pwwn 21:01:01:01:01:01:01:02  vsan    1 autonomous-fabric-id  1
       device-alias A2               vsan    2 autonomous-fabric-id  1
switch# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# device-alias database
switch(config-device-alias-db)# device-alias name A2 pwwn 21:01:01:01:01:01:01:02
switch(config-device-alias-db)# exit
switch(config)# device-alias commit
inter-VSAN zone member cannot be in more than one VSAN
```

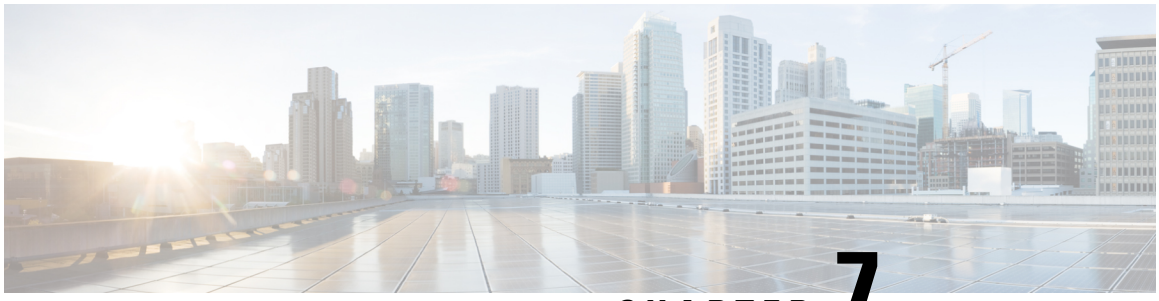
Step 5 Correct the conflict by making adjustments to the application configuration, or by making changes to the device-alias database, and running the **device-alias commit** command again.

Verifying the Device-Alias Database Status

This section provides information about verifying the device-alias database status.

Table 16: Verifying the Device-Alias Database Status

Command Name	Description
show cfs merge status name device-alias	Displays information about the status of the CFS merge for the device-alias database.
show device-alias database	Displays the entire device-alias database.
show device-alias internal validation info	Displays information about the status of the validation process (part of a commit or merge).
show device-alias merge conflicts	Displays the device-alias names or pWWNs causing a merge failure in Cisco MDS NX-OS Release 8.1(1) and later releases.
show device-alias merge status	Displays the result of the device-alias merge operation and the reason for the result.
show device-alias session status	Returns the status of the last CFS command, such as **clear**, **commit**, or **terminate**. The results of the last used CFS command and reason fields help identify the reason for the failure.
show device-alias status	Displays configuration information for the device-alias service, including whether fabric distribution is enabled, the number of device aliases in the database, lock information, and the database mode (Basic or Enhanced).



CHAPTER 7

Configuring 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. 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. Specifically, FSPF is used to:

- 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.

This chapter provides details on Fibre Channel routing services and protocols. It includes the following sections:

- About FSPF, on page 167
- FSPF Global Configuration, on page 170
- FSPF Interface Configuration, on page 172
- FSPF Routes, on page 179
- Load Balancing, on page 180
- In-Order Delivery, on page 186
- Flow Statistics Configuration, on page 190
- Default Settings, on page 196

About FSPF

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

- Supports multipath routing.
- Bases path status on a link state protocol.
- Routes hop by hop, based only on the domain ID.
- Runs only on E ports or TE ports and provides a loop free topology.
- Runs on a per VSAN basis. Connectivity in a given VSAN in a fabric is guaranteed only for the switches configured in that VSAN.

- Uses a topology database to keep track of 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.

FSPF Examples

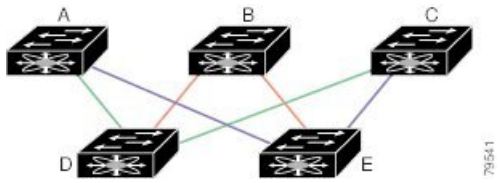
This section provides examples of topologies and applications that demonstrate the benefits of FSPF.

Note The FSPF feature can be used on any topology.

Fault Tolerant Fabric

Figure 11: Fault Tolerant Fabric, on page 168 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 11: Fault Tolerant Fabric



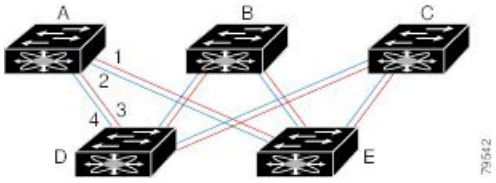
For example, if all links are of equal speed, the FSPF calculates two equal paths from A to C: A-D-C (green) and A-E-C (blue).

Redundant Links

To further improve on the topology in Figure 11: Fault Tolerant Fabric, on page 168, each connection between any pair of switches can be replicated; two or more links can be present between a pair of switches Figure 12: Fault Tolerant Fabric with Redundant Links, on page 169 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 12: 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 13: Failover Scenario Using Traffic Generators, on page 169. 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:

- Disabling the traffic link by physically removing the cable (see Table 17: Physically Removing the Cable for the SmartBits Scenario , on page 169).
- Shutting down the links in either switch 1 or switch 2 (see Table 18: Shutting Down the links in Switch for the SmartBits Scenario , on page 169).

Figure 13: Failover Scenario Using Traffic Generators

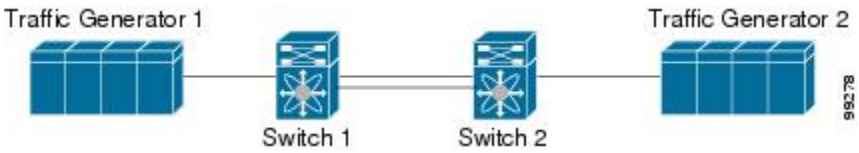


Table 17: Physically Removing the Cable for the SmartBits Scenario

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 18: Shutting Down the links in 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

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.

This section includes the following topics:

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 19: LSR Default Settings , on page 170 displays the default settings for switch responses.

Table 19: 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.

Configuring FSPF on a VSAN

To configure an FSPF feature for the entire VSAN, follow these steps:

Procedure

Step 1 switch# **config terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **fspf config vsan 1**

Enters FSPF global configuration mode for the specified VSAN.

Step 3 switch-config-(fspf-config)# **spf static**

Forces static SPF computation for the dynamic (default) incremental VSAN.

Step 4 switch-config-(fspf-config)# **spf hold-time 10**

Configures the hold time between two route computations in milliseconds (msec) for the entire VSAN. The default value is 0.

Note If the specified time is shorter, the routing is faster. However, the processor consumption increases accordingly.

Step 5 switch-config-(fspf-config)# **region 7**

Configures the autonomous region for this VSAN and specifies the region ID (7).

Resetting FSPF to the Default Configuration

To return the FSPF VSAN global configuration to its factory default, follow these steps:

Procedure

Step 1 switch# **config terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **no fspf config vsan 3**

Deletes the FSPF configuration for VSAN 3.

Enabling or Disabling FSPF

To enable or disable FSPF routing protocols, follow these steps:

Procedure

Step 1 switch# **config terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **fspf enable vsan 7**

Enables the FSPF routing protocol in VSAN 7.

Step 3 switch(config)# **no fspf enable vsan 5**

Disables the FSPF routing protocol in VSAN 5.

Clearing FSPF Counters for the VSAN

To clear the FSPF statistics counters for the entire VSAN, follow this step:

Procedure

switch# **clear fspf counters vsan 1**

Clears the FSPF statistics counters for the specified VSAN. If an interface reference is not specified, all counters are cleared.

FSPF Interface Configuration

Several FSPF commands are available on a per-interface basis. These configuration procedures apply to an interface in a specific VSAN.

This section includes the following topics:

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 30000. The default cost for 1 Gbps is 1000 and for 2 Gbps is 500.

Configuring FSPF Link Cost

To configure FSPF link cost, follow these steps:

Procedure

Step 1 switch# **config t**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **interface fc1/4**

switch(config-if)#

Configures the specified interface, or if already configured, enters configuration mode for the specified interface.

Step 3 switch(config-if)# **fspf cost 5 vsan 90**

Configures the cost for the selected interface in VSAN 90.

About FSPF Cost Multiplier

FSPF uses link costs to determine the shortest path between devices in a fabric. The default link costs become inefficient when calculating the cost of larger capacity port channels. Such paths may appear to have the same cost although they have different bandwidths leading to poor path selection by FSPF. The FSPF cost multiplier feature allows reassigning of links costs so that FSPF can calculate and select optimal high-speed paths.

Path cost calculation inefficiencies can occur when the total link bandwidth is over 128 Gbps. This feature should be configured when parallel paths above this threshold exist in a fabric so that FSPF selects paths as expected. A port channel can have a maximum of 16 member links so path inefficiencies may occur when port channels with as low as 9 x 16-Gbps members are present.

All switches in a fabric must use the same FSPF cost multiplier so that they all use the same basis for path cost calculations. This feature automatically distributes the configured FSPF cost multiplier to all Cisco MDS switches in the fabric with Cisco NX-OS versions that support the feature. If any switches are present in the fabric that do not support the feature, then the configuration fails and not applied to any switches. After the cost multiplier is accepted by all switches, a delay of 20 seconds occurs before being applied to ensure that all switches apply the update simultaneously. If the link costs do not change, there will not be any traffic disruption. However, if the update results in a different path selection by FSPF there may be a brief, one-time interruption to traffic as the new path is applied.

The link cost of an interface may also be manually changed in the default value. For more information, see About FSPF Link Cost, on page 172 section.

Setting up FSPF Cost Multiplier

The FSPF Cost Calculation Multiplier is configured to make the cost of the port-channel link optimal. The cost computation was not optimal for high speed port-channels (members of 16 Gbps speeds and later).The solution offers the following:

- FSPF Cost Calculation multiplier value 20 is configured to make the cost of the links optimal.
- FSPF Cost computation is optimal for port-channel with 16 members of up to 128-Gbps speed.
- Distribution of the FSPF cost calculation multiplier across the fabric for a given VSAN ensures all the links in the fabric for a VSAN are using the same factor for FSPF cost computation of a link.



Note The configuration of the FSPF Cost Multiplier is recommended to be done during a maintenance window, as there could be traffic impact due to change in routes based on the new link costs.

To set the cost admin factor, follow these steps.

Procedure

Step 1 switch# **config terminal**

Enters configuration mode.

Step 2 switch# **fspf config vsan**

switch(config-fspf-config)#

Enters Fabric Shortest Path First (FSPF) routing protocol.

Step 3 switch(config-fspf-config)# **cost-multiplier 20**

Sets the FSPF cost multiplier to 20

The following message is displayed.

This parameter will be distributed across all switches in the fabric. New routes will be computed after 20 seconds.

The following message is displayed when any switch in the fabric does not support the new cost computation admin factor value or the version is less than Cisco MDS NX-OS 9.3(1)

```
Unable to distribute fspf cost-multiplier due to one or more domains not supporting it. fspf
cost-multiplier supported on NX-OS 9.3(1) and later only.
VSAN 7
 FSPF cost multiplier is not supported on the following devices:
 Domain VSAN SWWN
 ------ ----------------------
 58 20:07:00:de:fb:b1:8d:e1
```

Displaying FSPF Cost Multiplier

This example show how to display the FSPF cost multiplier for VSAN 1:

Procedure

switch# **show fspf vsan1**

Displays the FSPF cost multiplier used for VSAN 1.

The following result of the command is displayed

```
switch(config)# show fspf vsan 1
FSPF routing for VSAN 1
FSPF routing administration status is enabled
FSPF routing operational status is UP
It is an intra-domain router
Autonomous region is 0
SPF hold time is 0 msec
MinLsArrival = 1000 msec , MinLsInterval = 2000 msec
Cost Multiplier = 1
Local Domain is 0x66(102)
Number of LSRs = 3, Total Checksum = 0x000198dd

Protocol constants :
   LS_REFRESH_TIME = 30 minutes (1800 sec)
   MAX_AGE         = 60 minutes (3600 sec)

Statistics counters :
   Number of LSR that reached MaxAge = 0
   Number of SPF computations        = 6
   Number of Checksum Errors         = 0
   Number of Transmitted packets :  LSU 30 LSA 32 Hello 984 Retransmitted LSU 0
   Number of received packets :  LSU 33 LSA 28 Hello 981 Error packets 3
```

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.

Configuring Hello Time Intervals

To configure the FSPF Hello time interval, follow these steps:

Procedure

Step 1 switch# **config t**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **interface fc1/4**

switch(config-if)#

Configures the specified interface, or if already configured, enters configuration mode for the specified interface.

Step 3 switch(config-if)# **fspf hello-interval 15 vsan 175**

switch(config-if)#

Specifies the hello message interval (15 seconds) to verify the health of the link in VSAN 175. The default is 20 seconds.

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.



Note This value must be the same in the ports at both ends of the ISL.

- An error is reported at the command prompt if the configured dead time interval is less than the hello time interval
- During a software upgrade, ensure that the fspf dead-interval is greater than the ISSU downtime (80 seconds). If the fspf dead-interval is lesser than the ISSU downtime, the software upgrade fails and the following error is displayed:

```
Service "fspf" returned error: Dead interval for interface is less than ISSU upgrade time.
```

Configuring Dead Time Intervals

To configure the FSPF dead time interval, follow these steps:

Procedure

Step 1 switch# **config terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **interface fc1/4**

switch(config-if)#

Configures the specified interface, or if already configured, enters configuration mode for the specified interface.

Step 3 switch(config-if)# **fspf dead-interval 25 vsan 7**

switch(config-if)#

Specifies the maximum interval for VSAN 7 before which a hello message must be received on the selected interface before the neighbor is considered lost. The default is 80 seconds.

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.

Configuring Retransmitting Intervals

To configure the FSPF retransmit time interval, follow these steps:

Procedure

Step 1 switch# **config terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **interface fc1/4**

switch(config-if)#

Configures the specified interface, or if already configured, enters configuration mode for the specified interface.

Step 3 switch(config-if)# **fspf retransmit-interval 15 vsan 12**

switch(config-if)#

Specifies the retransmit time interval for unacknowledged link state updates in VSAN 12. The default is 5 seconds.

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.

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.

To disable FSPF for a specific interface, follow these steps:

Procedure

Step 1 switch# **config terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **interface fc1/4**

switch(config-if)#

Configures a specified interface, or if already configured, enters configuration mode for the specified interface.

Step 3 switch(config-if)# **fspf passive vsan 1**

switch(config-if)#

Disables the FSPF protocol for the specified interface in the specified VSAN.

Step 4 switch(config-if)# **no fspf passive vsan 1**

switch(config-if)#

Reenables the FSPF protocol for the specified interface in the specified VSAN.

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.

Clearing FSPF Counters for an Interface

To clear the FSPF statistics counters for an interface, follow this step:

Procedure

switch# **clear fspf counters vsan 200 interface fc1/1**

Clears the FSPF statistics counters for the specified interface in VSAN 200.

FSPF Routes

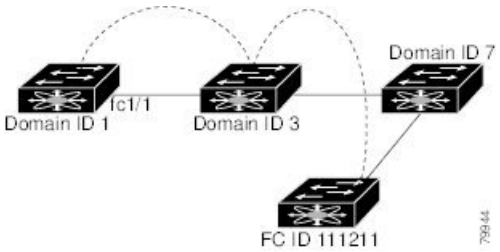
FSPF routes traffic across the fabric, based on entries in the FSPF database. These routes can be learned dynamically, or configured statically.

This section includes the following topics:

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 14: Fibre Channel Routes, on page 179).

Figure 14: Fibre Channel Routes



Note Other than in VSANs, runtime checks are not performed on configured and suspended static routes.

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 SAN-OS and NX-OS 4.1(1b) and later software 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.

Use the **mcast root lowest vsan** command to change the multicast root from the principal switch to lowest domain switch.

Setting the Multicast Root Switch

To use the lowest domain switch for the multicast tree computation, follow these steps:

Procedure

Step 1 switch# **config terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **mcast root lowest vsan 1**

Uses the lowest domain switch to compute the multicast tree.

Step 3 switch(config)# **mcast root principal vsan 1**

Defaults to using the principal switch to compute the multicast tree.

To display the configured and operational multicast mode and the selected root domain, use the **show mcast** command.

```
switch# show mcast vsan 1
Multicast root for VSAN 1
       Configured root mode : Principal switch
       Operational root mode : Principal switch
       Root Domain ID : 0xef(239)
```

Load Balancing

Load balancing is a forwarding mechanism that distributes traffic over equal-cost multipath (ECMP) and port channels. Load balancing uses a hash method to identify an egress link. The hash is a function that uses parameters in the frame header to identify a unique link to forward the frame to. The load balancing scheme used depends on both the type of ingress port and egress routing. If it is intended that traffic flow in both

directions on the same link, then ensure that the same load balancing scheme and hash method are used at both ends of the link.

Load Balancing Schemes

The following types of load balancing schemes are supported:

- Flow based—All frames between a given source FCID and destination FCID are transmitted on the same link. That is, whichever link is selected for the first exchange between the source-destination pair is used for all subsequent exchanges.
- Exchange based—The first frame in an exchange between a given source FCID and destination FCID is used to select an egress link and subsequent frames in the exchange are transmitted on the same link. However, subsequent exchanges between the source-destination pair will likely be transmitted on a different link. This provides more granular load balancing while preserving the order of frames within each exchange.

Figure 15: Flow Based Load Balancing, on page 181 illustrates how flow based load balancing works. In this example, when the first frame with a source FCID of sid1 and destination FCID of did1 is received for forwarding, port channel 2 is selected. Each subsequent frame in that flow is sent over the same port channel. No frame from sid1 to did1 utilizes port channel 1. Similarly, all frames with sid2 and did2 are sent over port channel 1. Exchange ID is not used with this type of load balancing.

Figure 15: Flow Based Load Balancing

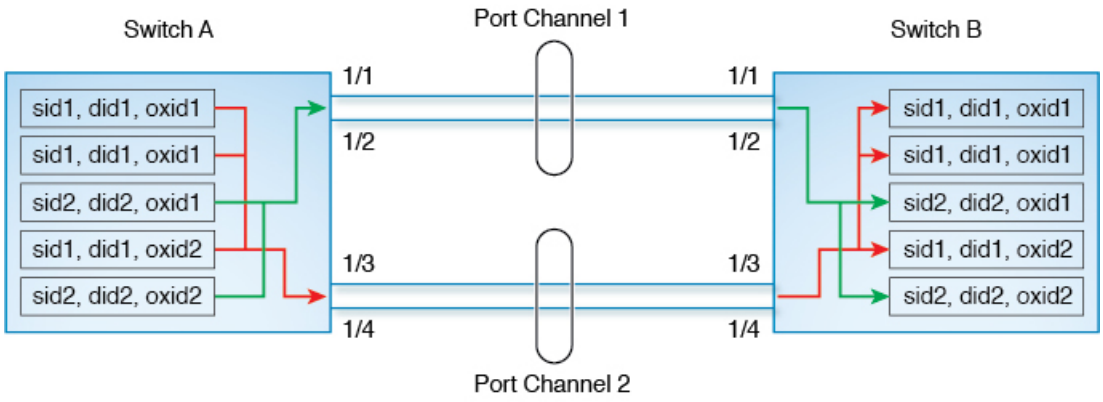
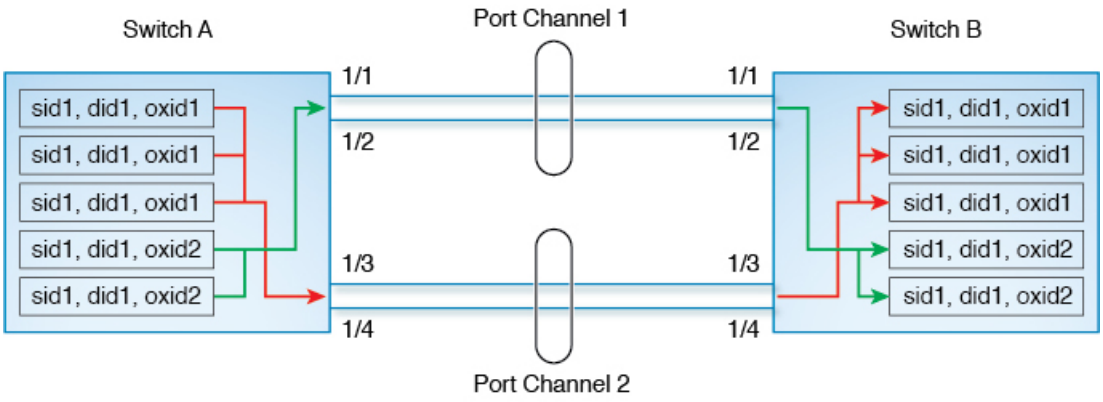


Figure 16: Exchange Based Load Balancing, on page 182 illustrates how exchange based load balancing works. In this example, when the first frame in an exchange between a source FCID sid1 and destination FCID did1 is received for forwarding, port channel 2 is selected. All remaining frames in that particular exchange are sent on the same port channel and none are sent on port channel 1. For the next exchange, the hash algorithm chooses port channel 1. So all frames in exchange 2 between the same source-destination pair are sent on port channel 1.

Figure 16: Exchange Based Load Balancing



Hashing Methods

Load balancing is applied to an ingress frame at two levels—At the first level, an ECMP hash is used to select an egress ECMP interface (this can be either a physical interface or logical interface such as a port channel interface) and at the second level, a port channel hash is used to select an egress port channel member.

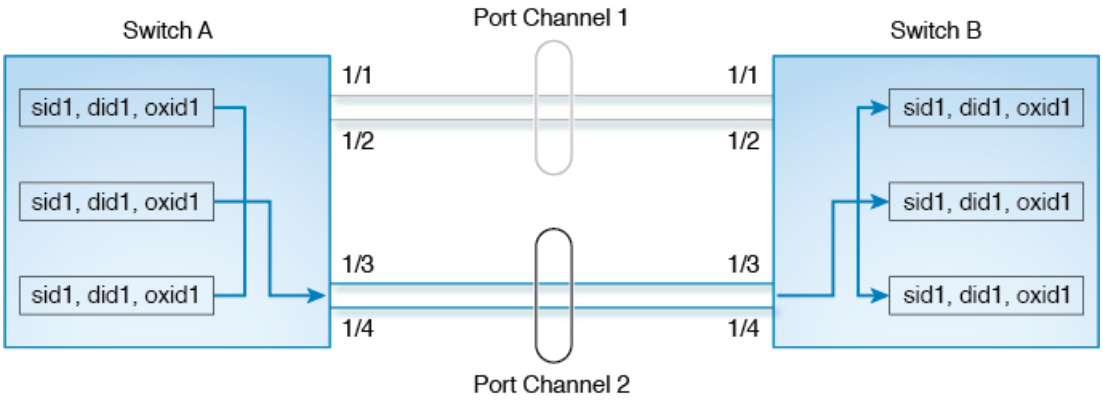
By default, the hash method that is used depends on the ingress hardware type. If either level of hash does not apply to the egress route, then no hash method is applied.

The following types of hashing methods are supported:

- ECMP Hashing Method—If multiple paths to a destination with equal cost exist in the switch, the FIB for the ingress port is updated with these paths for that destination. This hashing method is used to select one of such paths to send frames to.
- Port Channel Hashing Method—This hashing method is used to select an operational interface of an egress port channel.

Figure 17: ECMP Hashing Method, on page 182 illustrates how the ECMP hash method works. There are two port channels each including two equal speed links. Since the FSPF costs of the port channels are the same, both port channels are used for hashing. In this example, ECMP level hashing method selects port channel 2 as the egress port.

Figure 17: ECMP Hashing Method



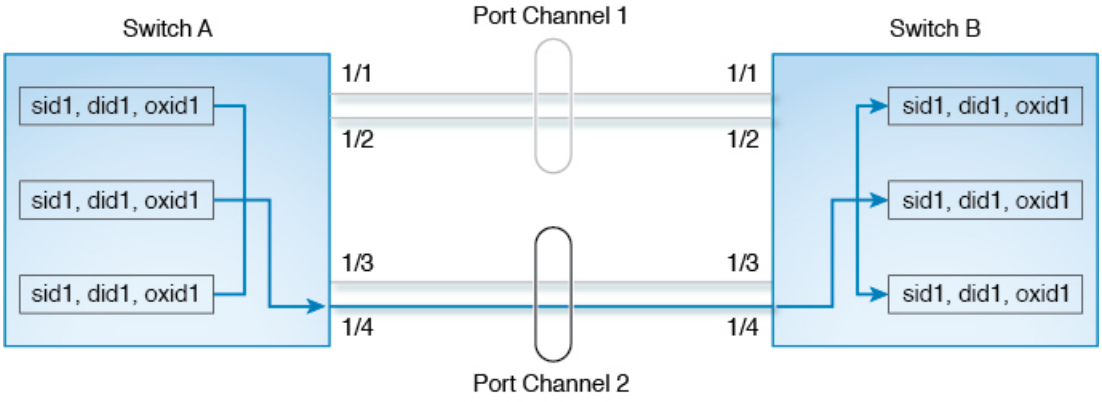
Depending on the type of ingress port, the following subtypes of ECMP hashing methods are supported:

- Type 1a
- Type 1b

For information on which hashing method is selected for a given ingress port, see Table 20: Hashing Matrix, on page 183.

Figure 18: Port Channel Hashing Method, on page 183 illustrates how port channel hashing method works. Continuing the Figure 17: ECMP Hashing Method, on page 182 example where port channel 2 was selected as the egress port, a port channel hash is subsequently applied to select an egress port within the port channel. In this example, the frames are transmitted by interface 1/4 of the selected port channel.

Figure 18: Port Channel Hashing Method



Depending on the type of ingress port, the following types of port channel hashing methods are supported:

- Type 2a
- Type 2b

For information on which hashing method is selected for a given ingress port, see Table 20: Hashing Matrix, on page 183.

Table 20: Hashing Matrix

Ingress Interface	Egress Interface	ECMP Hash Method	Port Channel Hash Method
Fibre Channel or FCIP port on Cisco MDS 9500 with Generation 3 or 4 module	Fibre Channel or FCIP ISL	Type 1a	Type 2b (only when at least one FCIP port is up)

Ingress Interface	Egress Interface	ECMP Hash Method	Port Channel Hash Method
Fibre Channel port on Cisco MDS 9500 with Generation 3 or 4 module	Fibre Channel ISL	Type 1a	Type 2a **Note** The hashing method changes to type 2b if FCIP tunnel were brought up in the switch. The hashing method will remain as type 2b even if the FCIP module is removed until the next switch reload.
Fibre Channel, FCIP, or FCoE port on Cisco MDS 9250i	Fibre Channel, FCIP, or FCoE ISL	Type 1a	Type 2b
Fibre Channel, FCIP, or FCoE port on Cisco MDS 9250i	FCIP ISL connected to Cisco MDS 24/10-Port SAN Extension Module with FCIP enhanced.	Type 1a	Type 1a
Fibre Channel port on Cisco MDS 9700	FCIP ISL	Type 1a	Type 1a
	Fibre Channel or FCoE ISL	Type 1a	Type 2a
FCIP port on Cisco MDS 24/10-Port SAN Extension Module	FCIP ISL	Type 1b	Type 1b
	Fibre Channel or FCoE ISL	Type 1b	Type 2a
FCoE port on Cisco MDS 9700	FCIP ISL	Type 1b	Type 1b
	Fibre Channel or FCoE ISL	Type 1b	Type 2a
Fibre Channel port on Cisco MDS 9148S Fibre Channel port on Cisco MDS 9396S Fibre Channel port on Cisco MDS 9132T Fibre Channel port on Cisco MDS 9396T and 9148T	Fibre Channel ISL	Type 1a	Type 2a

Configure VSAN hash-method

The Cisco MDS 9000 switch supports various hashing methods used for load balancing links within a VSAN. Hash methods influence how traffic is load balanced across parallel links. Common hash types include methods based on source and destination FCIDs or exchange IDs.

- The **vsan hash-method** command configures the hash method for a VSAN.
- Hash methods influence how traffic is load balanced across links.
- Common hash types include methods based on source and destination FCIDs or exchange IDs.
- You can verify the configured hash method with the **show vsan** command, which displays the load balancing scheme for the VSAN.

Procedure

Step 1 Enter the configuration mode.

Example:

```
switch# configure terminal
```

Step 2 Set the hash method for a specfic VSAN. Use the command **vsan** <*vsan-id*> **hash-method** <*method*>.

Supported hash methods are:

- 1a, 1b: ECMP hashing methods
- 2a, 2b: Port channel hashing methods
- Replace `<vsan-id>` with the VSAN number. VSAN ID range is from 1 to 4093
- Replace `<method>` with the desired hash method such as:

Example:

```
switch(config)# vsan 2 hash-method 2a
```

Step 3 Exit configuration mode and save the configuration.

What to do next

Verify the configuration using **show vsan** and **show loadbalancing hash-type** commands,

Display VSAN load balancing hash-type

Verify load balancing

The following example shows how to display the load balancing hash method for all VSANs.

```
switch# show vsan
vsan 1 information
         name:VSAN0001  state:active
```

```
        interoperability mode:default
        loadbalancing:src-id/dst-id/oxid
        operational state:up
vsan 2 information
        name:VmVSAN  state:active
        interoperability mode:default
        loadbalancing:src-id/dst-id/oxid
        operational state:up
```

Verify load balancing with hash-type

The following example shows how to display the link that will be selected for traffic between a given source and destination FCID using the specified Exchange ID when the hash type 2a is set for a particular VSAN.

```
switch# show loadbalancing hash-type 2a vsan 1000 0x830001 0x790000 4
Link : 3 will be chosen among 4 links for this hash method
```

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.

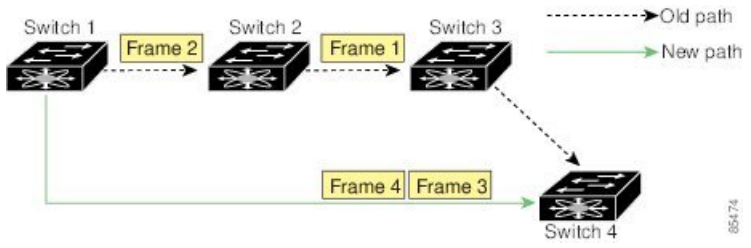
Tip If you enable the in-order delivery feature, the graceful shutdown feature is not implemented.

This section includes the following topics:

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.

Figure 19: Route Change Delivery



In Figure 19: Route Change Delivery, on page 186, 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.

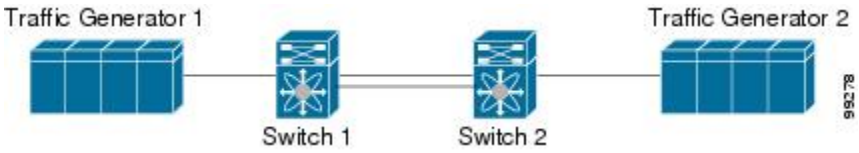
If the in-order guarantee feature is enabled, the frames within the network are treated as follows:

- 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 flow or the same initiator-target flow can switch from one path to another faster path.

Figure 20: Link Congestion Delivery



In Figure 20: Link Congestion Delivery, on page 187 , the port of the old path (black 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, known as Lossless IOD. 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.
- 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, on page 189.

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 Series.

Note Enabling or disabling the IOD feature does not disrupt traffic.

Tip We recommend that you only enable this feature when devices that cannot handle any out-of-order frames are connected to the fabric. Load-balancing algorithms within the Cisco MDS 9000 Series 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.

Enabling In-Order Delivery Globally

To ensure that the in-order delivery parameters are uniform across all VSANs on an MDS 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.

To enable in-order delivery for the switch, follow these steps:

Procedure

Step 1 switch# **config terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **in-order-guarantee**

Enables in-order delivery in the switch.

Step 3 switch(config)# **no in-order-guarantee**

Reverts the switch to the factory defaults and disables the in-order delivery feature.

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.

To use the lowest domain switch for the multicast tree computation, follow these steps:

Procedure

Step 1 switch# **config terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **in-order-guarantee vsan 3452**

Enables in-order delivery in VSAN 3452.

Step 3 switch(config)# **no in-order-guarantee vsan 101**

Reverts the switch to the factory defaults and disables the in-order delivery feature in VSAN 101.

Displaying the In-Order Delivery Status

Use the **show in-order-guarantee** command to display the present configuration status:

```
switch# show in-order-guarantee
global inorder delivery configuration:guaranteed
VSAN specific settings
vsan 1 inorder delivery:guaranteed
vsan 101 inorder delivery:not guaranteed
vsan 1000 inorder delivery:guaranteed
vsan 1001 inorder delivery:guaranteed
vsan 1682 inorder delivery:guaranteed
vsan 2001 inorder delivery:guaranteed
vsan 2009 inorder delivery:guaranteed
vsan 2456 inorder delivery:guaranteed
vsan 3277 inorder delivery:guaranteed
vsan 3451 inorder delivery:guaranteed
vsan 3452 inorder delivery:guaranteed
```

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.

To configure the network and the switch drop latency time, follow these steps:

Procedure

Step 1 switch# **configure terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **fcdroplatency network 5000**

Configures network drop latency time to be 5000 ms for the network. The valid range is 0 to 60000 ms. The default is 2000 ms.

Note The network drop latency must be computed as the sum of all switch latencies of the longest path in the network.

Step 3 switch(config)# **fcdroplatency network 6000 vsan 3**

Configures network drop latency time to be 6000 ms for VSAN 3.

Step 4 switch(config)# **no fcdroplatency network 4500**

Removes the current fcdroplatecy network configuration (4500) and reverts the switch to the factory defaults.

Displaying Latency Information

You can view the configured latency parameters using the **show fcdroplatency** command (see Displays Administrative Distance, on page 190).

Displays Administrative Distance

```
switch# show fcdroplatency

switch latency value:500 milliseconds
global network latency value:2000 milliseconds
VSAN specific network latency settings
vsan 1 network latency:5000 milliseconds
vsan 2 network latency:2000 milliseconds
vsan 103 network latency:2000 milliseconds
vsan 460 network latency:500 milliseconds
```

Flow Statistics Configuration

Flow statistics count the ingress traffic in the aggregated statistics table. You can collect two kinds of statistics:

- Aggregated flow statistics to count the traffic for a VSAN.
- Flow statistics to count the traffic for a source and destination ID pair in a VSAN.

This section includes the following topics:

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.

Counting Aggregated Flow Statistics

To count the aggregated flow statistics for a VSAN, follow these steps:

Procedure

Step 1 switch# config t

switch(config)#

Enters configuration mode.

Step 2 switch(config)# fcflow stats aggregated module 1 index 1005 vsan 1

switch(config)#

Enables the aggregated flow counter.

Step 3 switch(config)# no fcflow stats aggregated module 1 index 1005 vsan 1

switch(config)#

Disables the aggregated flow counter.

Counting Individual Flow Statistics

To count the flow statistics for a source and destination FC ID in a VSAN, follow these steps:

Procedure

Step 1 switch# config t

switch(config)#

Enters configuration mode.

Step 2 switch(config)# fcflow stats module 1 index 1 0x145601 0x5601ff 0xffffff vsan 1

switch(config)#

Enables the flow counter.

Note The source ID and the destination ID are specified in FC ID hex format (for example, 0x123aff). The mask can be one of 0xff0000 or 0xffffff.

Step 3 switch(config)# no fcflow stats aggregated module 2 index 1001 vsan 2

switch(config)#

Disables the flow counter.

Clearing FIB Statistics

Use the **clear fcflow stats** command to clear the aggregated flow counter (see Examples Clears Aggregated Flow Counters, on page 192 and Clears Flow Counters for Source and Destination FC IDs, on page 192).

Clears Aggregated Flow Counters

```
switch# clear fcflow stats aggregated module 2 index 1
```

Clears Flow Counters for Source and Destination FC IDs

```
switch# clear fcflow stats module 2 index 1
```

Displaying Flow Statistics

Use the **show fcflow stats** commands to view flow statistics (see Example Displays Aggregated Flow Details for the Specified Module, on page 192 to Displays Flow Index Usage for the Specified Module, on page 193).

Displays Aggregated Flow Details for the Specified Module

```
switch# show fcflow stats aggregated module 6
Idx  VSAN frames        bytes
---- ---- --------    -------
1  800   20185860     1211151600
```

Displays Flow Details for the Specified Module

```
switch# show fcflow stats module 6
Idx   VSAN   DID    SID    Mask        frames      bytes
---- -----  -------       ------     -----     -----   ------
2   800   0x520400    0x530260   0xffffff      20337793 1220267580
```

Displays Flow Index Usage for the Specified Module

```
switch# show fcflow stats usage module 6
Configured flows for module 6: 1-2
```

Displaying Global FSPF Information

Displays FSPF Information for a Specified VSAN, on page 193 displays global FSPF information for a specific VSAN:

- Domain number of the switch.
- Autonomous region for the switch.
- Min_LS_arrival: minimum time that must elapse before the switch accepts LSR updates.
- Min_LS_interval: minimum time that must elapse before the switch can transmit an LSR.

Tip If the Min_LS_interval is higher than 10 seconds, the graceful shutdown feature is not implemented.

- LS_refresh_time: interval time lapse between refresh LSR transmissions.
- Max_age: maximum time aa LSR can stay before being deleted.

Displays FSPF Information for a Specified VSAN

```
switch# show fspf vsan 1
FSPF routing for VSAN 1
FSPF routing administration status is enabled
FSPF routing operational status is UP
It is an intra-domain router
Autonomous region is 0
SPF hold time is 0 msec
MinLsArrival = 1000 msec , MinLsInterval = 5000 msec
Local Domain is 0x65(101)
Number of LSRs = 3, Total Checksum = 0x0001288b
Protocol constants :
   LS_REFRESH_TIME = 1800 sec
   MAX_AGE         = 3600 sec
Statistics counters :
   Number of LSR that reached MaxAge = 0
   Number of SPF computations        = 7
   Number of Checksum Errors         = 0
   Number of Transmitted packets :  LSU 65 LSA 55 Hello 474 Retranmsitted LSU 0
   Number of received packets :  LSU 55 LSA 60 Hello 464 Error packets 10
```

Displaying the FSPF Database

Displays FSPF Database Information (Prior to Cisco MDS NX-OS Release 9.4(1)), on page 194 displays a summary of the FSPF database for a specified VSAN. If other parameters are not specified, all LSRs in the database are displayed:

- LSR type

- Domain ID of the LSR owner
- Domain ID of the advertising router
- LSR age
- LSR incarnation member
- Number of links

You could narrow the display to obtain specific information by issuing additional parameters for the domain ID of the LSR owner. For each interface, the following information is also available:

- Domain ID of the neighboring switch
- E port index
- Port index of the neighboring switch
- Prior to Cisco MDS NX-OS Release 9.4(1), the Link type is numerical.
- From Cisco MDS NX-OS Release 9.4(1), the Link type is alphnumercial and the following types.

Table 21: Link Tyoe

Link Type	Description
P2P	Peer-to-peer interfaces connections
FCIP PC	Fibre Channel over IP Protocol (FCIP) connection
FC PC	Fibre Channel connections
VFC PC	virtual Fibre Channel connections

- Cost

Displays FSPF Database Information (Prior to Cisco MDS NX-OS Release 9.4(1))

```
switch# show fspf database vsan 1
FSPF Link State Database for VSAN 1 Domain 0x0c(12)
LSR Type                = 1
Advertising domain ID   = 0x0c(12)
LSR Age                 = 1686
LSR Incarnation number  = 0x80000024
LSR Checksum            = 0x3caf
Number of links         = 2
 NbrDomainId      IfIndex   NbrIfIndex    Link Type         Cost
--------------------------------------------------------------------------------
   0x65(101) 0x0000100e     0x00001081              1          500
   0x65(101) 0x0000100f     0x00001080              1          500
FSPF Link State Database for VSAN 1 Domain 0x65(101)
LSR Type                = 1
Advertising domain ID   = 0x65(101)
LSR Age                 = 1685
LSR Incarnation number  = 0x80000028
LSR Checksum            = 0x8443
Number of links         = 6
```

```
 NbrDomainId      IfIndex   NbrIfIndex    Link Type         Cost
-------------------------------------------------------------------------------
   0xc3(195) 0x00001085     0x00001095               1          500
   0xc3(195) 0x00001086     0x00001096               1          500
   0xc3(195) 0x00001087     0x00001097               1          500
   0xc3(195) 0x00001084     0x00001094               1          500
    0x0c(12) 0x00001081     0x0000100e               1          500
    0x0c(12) 0x00001080     0x0000100f               1          500
FSPF Link State Database for VSAN 1 Domain 0xc3(195)
LSR Type                = 1
Advertising domain ID   = 0xc3(195)
LSR Age                 = 1686
LSR Incarnation number  = 0x80000033
LSR Checksum            = 0x6799
Number of links         = 4
 NbrDomainId      IfIndex   NbrIfIndex    Link Type         Cost
-------------------------------------------------------------------------------
   0x65(101) 0x00001095     0x00001085               1          500
   0x65(101) 0x00001096     0x00001086               1          500
   0x65(101) 0x00001097     0x00001087               1          500
   0x65(101) 0x00001094     0x00001084               1          500
```

Displays FSPF Database Information (From Cisco MDS NX-OS Release 9.4(1))

```
switch# show fspf database vsan 1
FSPF Link State Database for VSAN 1 Domain 0xd8(216)
LSR Type                = 1
Advertising domain ID   = 0xd8(216)
LSR Age                 = 646
LSR Incarnation number  = 0x80001c06
LSR Checksum            = 0x0e03
Number of links         = 5
    NbrDomainId           IfIndex(Interface Name)  NbrIfIndex  Link Type   Cost
    ------------------------------------------------------------------------------
      0xe3(227)  0x00010312(              fc4/19)  0x00010011        P2P     62
      0xe3(227)  0x00010313(              fc4/20)  0x0001000e        P2P     62
      0xdb(219)  0x0004003b(      port-channel160)  0x0004003b    FCIP PC    100
      0xdb(219)  0x000400ff(     port-channel256)  0x000400ff      FC PC     31
      0x59(89)   0x00fb0200(           vfc-po513)  0x00fb0200     VFC PC     50
```

Displaying FSPF Interfaces

Displays FSPF Interface Information, on page 196 displays the following information for each selected interface.

- Link cost
- Timer values
- Neighbor’s domain ID (if known)
- Local interface number
- Remote interface number (if known)
- FSPF state of the interface
- Interface counters

Displays FSPF Interface Information

```
switch# show fspf vsan 1 interface fc1/1
FSPF interface fc1/1 in VSAN 1
FSPF routing administrative state is active
Interface cost is 500
Timer intervals configured, Hello 20 s, Dead 80 s, Retransmit 5 s
FSPF State is FULL
Neighbor Domain Id is 0x0c(12), Neighbor Interface index is 0x0f100000
Statistics counters :
   Number of packets received : LSU  8  LSA  8  Hello 118  Error packets 0
   Number of packets transmitted : LSU  8  LSA  8  Hello 119  Retransmitted LSU 0
   Number of times inactivity timer expired for the interface = 0
```

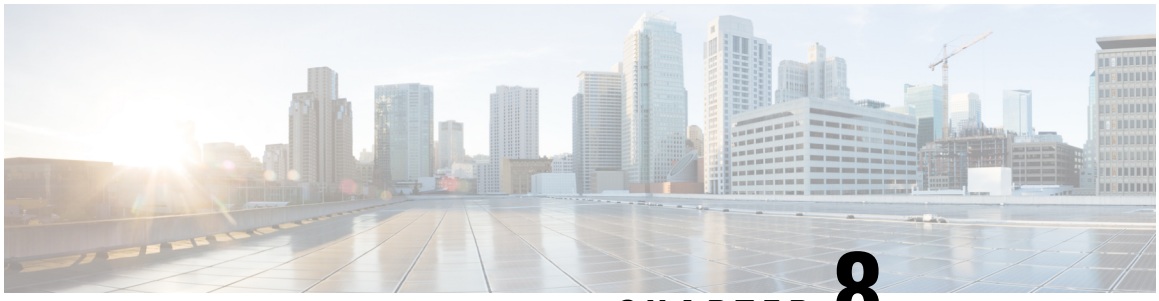
Default Settings

Table 22: Default FSPF Settings , on page 196 lists the default settings for FSPF features.

Table 22: 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.

Parameters	Default
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.



CHAPTER 8

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. It includes the following sections:

- About FLOGI, on page 199
- Name Server , on page 203
- FDMI, on page 209
- VMID, on page 213
- RSCN , on page 221
- Default Settings, on page 231
- Enabling Port Pacing , on page 231

About FLOGI

In a Fibre Channel fabric, each host or disk requires an Fibre Channel ID. Use the **show flogi database** command to verify if a storage device is displayed in the FLOGI table as in the next section. 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.

FLOGI Scale Optimization

The FLOGI scale optimization feature enables MDS switches to support an increased number of FLOGIs for module and chassis. The FLOGI scale optimization preloads the routing information for devices after a switch or module reload. This can speed up the Accepts to FLOGIs. This feature is supported on all MDS switches, except Cisco MDS 9250i Multiservice Fabric Switch and Cisco MDS 9148S 16G Multilayer Fabric Switch, and is enabled by default, starting from Cisco MDS NX-OS Release 8.1(1). As of Cisco MDS Release 8.2(2), the higher FLOGI scale limits are published only for MDS 9718. For more information see the "Cisco MDS NX-OS Configuration Limits" documentation for the FLOGI limits.

FLOGI Quiesce Timeout

The FLOGI Quiesce Timeout feature causes the FLOGI process to delay notifications to the other Fibre Channel services such as Routing Information and Fibre Channel Name Server when a device logs out from a fabric or when an interface goes down. If the device logs back to the fabric within the FLOGI quiesce timeout value, the FLOGI Accept can be immediately returned without the other Fibre Channel services being notified. This feature must be disabled by setting the timeout value to zero when there are devices in the fabric that can share a pWWN at different times by logging into different switches within the fabric during failover situations.

Restrictions

- Downgrading from Cisco MDS NX-OS Release 8.1(1) to an earlier release is not supported when FLOGI Scale Optimization is enabled. This feature must be disabled before downgrading. For more information on disabling this feature, see the Disabling FLOGI Scale Optimization section.
- In Cisco MDS NX-OS Releases 8.1 and Release 8.2, the default FLOGI quiesce timeout value is 2000 ms.

  However, starting with Cisco MDS NX-OS Release 8.3(1), the default FLOGI quiesce timeout value was changed from 2000 ms to 0 ms. Any configured FLOGI quiesce timeout value will be maintained on upgrading. If the FLOGI quiesce timeout value is not configured when upgrading to Cisco MDS NX-OS Releases 8.3(1) or later release, the new default value of 0 ms will be used.
- This feature is supported on all MDS switches, except Cisco MDS 9250i Multiservice Fabric Switch and Cisco MDS 9148S 16G Multilayer Fabric Switch.
- Cisco DCNM and SNMP support is not available for this feature.
- This feature is supported only on the Fibre Channel ports on the Cisco MDS 24/10 port SAN Extension Module.

Enabling FLOGI Scale Optimization and Quiesce Timeout

To enable FLOGI scale optimization and quiesce timeout, perform the following steps:

Procedure

Step 1 Enter global configuration mode:

switch# **configure terminal**

Step 2 Enable FLOGI scale optimization:

switch(config)# **flogi scale enable**

Step 3 Set the FLOGI quiesce timeout value to retain device login information:

switch(config)# **flogi quiesce timeout** *milliseconds*

For information on the default FLOGI quiesce timeout value, see the Restrictions section.

Step 4 Exit global configuration mode:

switch(config)# **exit**

Step 5 (Optional) Verify that FLOGI scale optimization is enabled:

switch# **show flogi internal info | i scale**

switch# **show flogi internal info | i quiesce**

Example: Enabling FLOGI Scale Optimization

The following running configuration shows how to enable FLOGI scale optimization and set the quiesce timeout value to 2000 milliseconds:

```
configure terminal
 flogi scale enable
 flogi quiesce timeout 2000
 exit
```



Note For more information on FLOGI scale numbers, see the Cisco MDS NX-OS Configuration Limits document.

The following sample outputs from the **show flogi internal info | i scale** and **show flogi internal info | i quiesce** commands display details about FLOGI scale optimization:

```
switch# show flogi internal info | i scale
  Stats: fs_flogi_scale_enabled: 1
switch# show flogi internal info | i quiesce
  Stats: fs_flogi_quiesce_timerval: 2000
```

Disabling FLOGI Scale Optimization and Quiesce Timeout

To disable FLOGI scale optimization and quiesce timeout, perform the following steps:

Procedure

Step 1 Enter global configuration mode:

switch# **configure terminal**

Step 2 Disable FLOGI scale optimization:

switch(config)# **no flogi scale enable**

Step 3 Set the FLOGI quiesce timeout value to 0:

switch(config)# **flogi quiesce timeout** *0*

The default quiesce timeout value is 2000 milliseconds.

Step 4 Exit global configuration mode:

switch(config)# **exit**

Step 5 (Optional) Verify that FLOGI scale optimization is disabled:

switch# **show flogi internal info | i scale**

switch# **show flogi internal info | i quiesce**

Example: Disabling FLOGI Scale Optimization

The following running configuration shows how to disable FLOGI scale optimization and set the quiesce timeout value to 0 milliseconds:

```
configure terminal
 no flogi scale enable
 flogi quiesce timeout 0
 exit
```

The following sample outputs from the **show flogi internal info | i scale** and **show flogi internal info | i quiesce** commands display details about FLOGI scale optimization:

```
switch# show flogi internal info | i scale
  Stats: fs_flogi_scale_enabled: 0
switch# show flogi internal info | i quiesce
  Stats: fs_flogi_quiesce_timerval: 0
```

Displaying FLOGI Details

To view the FLOGI database details, use the show flogi database command. See Examples Displays Details on the FLOGI Database , on page 202 to Displays the FLOGI Database by FC ID, on page 203.

Displays Details on the FLOGI Database

```
switch# show flogi database
--------------------------------------------------------------------------------
INTERFACE  VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------------
sup-fc0    2     0xb30100  10:00:00:05:30:00:49:63  20:00:00:05:30:00:49:5e
fc9/13     1     0xb200e2  21:00:00:04:cf:27:25:2c  20:00:00:04:cf:27:25:2c
fc9/13     1     0xb200e1  21:00:00:04:cf:4c:18:61  20:00:00:04:cf:4c:18:61
fc9/13     1     0xb200d1  21:00:00:04:cf:4c:18:64  20:00:00:04:cf:4c:18:64
fc9/13     1     0xb200ce  21:00:00:04:cf:4c:16:fb  20:00:00:04:cf:4c:16:fb
fc9/13     1     0xb200cd  21:00:00:04:cf:4c:18:f7  20:00:00:04:cf:4c:18:f7
Total number of flogi = 6.
```

Displays the FLOGI Database by Interface

```
switch# show flogi database interface fc1/11
INTERFACE   VSAN    FCID        PORT NAME                NODE NAME
----------  ------  ---------   -----------------------  ----------------------
fc1/11      1       0xa002ef    21:00:00:20:37:18:17:d2  20:00:00:20:37:18:17:d2
fc1/11      1       0xa002e8    21:00:00:20:37:38:a7:c1  20:00:00:20:37:38:a7:c1
fc1/11      1       0xa002e4    21:00:00:20:37:6b:d7:18  20:00:00:20:37:6b:d7:18
fc1/11      1       0xa002e2    21:00:00:20:37:18:d2:45  20:00:00:20:37:18:d2:45
fc1/11      1       0xa002e1    21:00:00:20:37:39:90:6a  20:00:00:20:37:39:90:6a
fc1/11      1       0xa002e0    21:00:00:20:37:36:0b:4d  20:00:00:20:37:36:0b:4d
fc1/11      1       0xa002dc    21:00:00:20:37:5a:5b:27  20:00:00:20:37:5a:5b:27
fc1/11      1       0xa002da    21:00:00:20:37:18:6f:90  20:00:00:20:37:18:6f:90
fc1/11      1       0xa002d9    21:00:00:20:37:5b:cf:b9  20:00:00:20:37:5b:cf:b9
fc1/11      1       0xa002d6    21:00:00:20:37:46:78:97  0:00:00:20:37:46:78:97
Total number of flogi = 10.
```

Displays the FLOGI Database by VSAN

```
switch# show flogi database vsan 1
--------------------------------------------------------------------------------
INTERFACE  VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------------
fc1/3      1     0xef02ef  22:00:00:20:37:18:17:d2  20:00:00:20:37:18:17:d2
fc1/3      1     0xef02e8  22:00:00:20:37:38:a7:c1  20:00:00:20:37:38:a7:c1
fc1/3      1     0xef02e4  22:00:00:20:37:6b:d7:18  20:00:00:20:37:6b:d7:18
fc1/3      1     0xef02e2  22:00:00:20:37:18:d2:45  20:00:00:20:37:18:d2:45
fc1/3      1     0xef02e1  22:00:00:20:37:39:90:6a  20:00:00:20:37:39:90:6a
fc1/3      1     0xef02e0  22:00:00:20:37:36:0b:4d  20:00:00:20:37:36:0b:4d
fc1/3      1     0xef02dc  22:00:00:20:37:5a:5b:27  20:00:00:20:37:5a:5b:27
fc1/3      1     0xef02da  22:00:00:20:37:18:6f:90  20:00:00:20:37:18:6f:90
fc1/3      1     0xef02d9  22:00:00:20:37:5b:cf:b9  20:00:00:20:37:5b:cf:b9
fc1/3      1     0xef02d6  22:00:00:20:37:46:78:97  20:00:00:20:37:46:78:97
Total number of flogi = 10.
```

Displays the FLOGI Database by FC ID

```
switch# show flogi database fcid 0xef02e2
--------------------------------------------------------------------------------
INTERFACE  VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------------
fc1/3      1     0xef02e2  22:00:00:20:37:18:d2:45  20:00:00:20:37:18:d2:45
Total number of flogi = 1.
```

For more information, see the Default Company ID List, on page 296 and refer to the “Loop Monitoring” section in the *Cisco MDS 9000 Family Troubleshooting Guide*.

Name Server

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.

This section includes the following topics:

Bulk Notification Sent from the Name Server

In order to improve the performance of the Fibre Channel protocols on the Cisco MDS 9000 switch, the name server optimizes the remote entry change notifications by sending multiple notifications in one MTS payload. Nearly 10 other components that receive this MTS notification would have to function on the single bulk notification instead of multiple notifications.

Enabling Name Server Bulk Notification

For NX-OS Release 6.2(1) to 6.2(7), bulk notification is disabled by default. Enabling this feature in one switch has no bearing on the other switches in the same fabric.

Note From NX-OS Release 6.2(9) onwards, bulk notification is enabled by default.

Restrictions

- Whenever the intelligent applications such as the DMM, IOA, and SME are enabled, the bulk notification feature is not supported.
- Any configuration present in the FC-Redirect, conflicts with the bulk notification feature.

Note The above restrictions are applicable only to release 6.2.7.

To enable the name server bulk notification, follow these steps for NX-OS Release 6.2(1) to 6.2(7):

Procedure

Step 1 switch# **config t**

Enters configuration mode.

Step 2 switch(config)# **fcns bulk-notify**

switch(config)#

Enables the transmission of multiple name server entry change notification in one Messaging and Transaction Services (MTS) payload.

Disabling Name Server Bulk Notification

To disable the name server bulk notification, follow these steps for NX-OS Release 6.2(1) to 6.2(7):

Procedure

Step 1 switch# **config t**

Enters configuration mode.

Step 2 switch(config)# **no fcns bulk-notify**

switch(config)#

Disables the transmission of multiple name server entry change notification in one Messaging and Transaction Services (MTS) payload.

Disabling Name Server Bulk Notification for NX-OS Release 6.2(9)

To disable the name server bulk notification, follow these steps for NX-OS Release 6.2(9) and later:

Procedure

Step 1 switch# **config t**

Enters configuration mode.

Step 2 switch(config)# **fcns no-bulk-notify**

switch(config)#

Disables the transmission of multiple name server entry change notification in one Messaging and Transaction Services (MTS) payload.

Re-enabling Name Server Bulk Notification

To re-enable once it is disabled already for NX-OS Release 6.2(9) and later, follow these steps:

Procedure

Step 1 switch# **config terminal**

Enters configuration mode.

Step 2 switch(config)# **no fcns no-bulk-notify**

switch(config)#

Re-enables the transmission of multiple name server entry change notification in one Messaging and Transaction Services (MTS) payload.

Name Server Proxy Registration

All name server registration requests are sent from the same port with a parameter that is registered or changed. If the port that does not have the parameter, then the request is rejected.

This authorization enables WWNs to register specific parameters for another node.

Registering Name Server Proxies

To register the name server proxy, follow these steps:

Procedure

Step 1 switch# **config terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **fcns proxy-port 21:00:00:e0:8b:00:26:d0 vsan 2**

Configures a proxy port for the specified VSAN.

About Rejecting Duplicate pWWN

By FC standard, NX-OS will accept a login on any interface of a pwwn that is already logged in on the same switch, same vsan and same fcdomain. To prevent the same pwwn from logging in the same switch on a different interface, use the port security feature.

By default, any future flogi (with duplicate pwwn) on different switch in the same vsan, will be rejected and previous FLOGI retained, which does not follow FC standards. If you disable this option, any future flogi (with duplicate pwwn) on different switch in the same VSAN, would be allowed to succeed by deleting previous FCNS entry

Rejecting Duplicate pWWNs

To reject duplicate pWWNs, follow these steps:

Procedure

Step 1 switch# **configure terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **fcns reject-duplicate-pwwn vsan 1**

Any future flogi (with duplicate pwwn) on different switch, will be rejected and previous FLOGI retained. (default)

Step 3 switch(config)# **no fcns reject-duplicate-pwwn vsan 1**

Any future flogi (with duplicate pwwn) on different switch, will be allowed to succeed by deleting earlier FCNS entry.

But you can still see the earlier entry in FLOGI database in the other switch.

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.

Optimizing Name Server Database Sync

If an end device doesn't register FC4 feature with Name Server database, VHBA (also called scsi-target) component would perform PRLI to the end device to discover FC4 feature and register with Name Server on behalf of end device. This discovery from VHBA was performed both for locally connected devices

as well as remotely connected devices. This discovery was unnecessary for remotely connected devices because, Name Server would get FC4 feature of remotely connected devices through regular Name Server sync protocol. So, the default behavior of VHBA component has been modified to discover only locally connected devices. To modify this behavior, follow these steps:

Procedure

Step 1 switch(config)# scsi-target discovery

Enables a switch to discover fc4-feature for

remote devices also. But this would not be

the default behavior if the users reload or switchover the switch.

Step 2 switch(config)# scsi-target discovery local-only

Switches back to the default behavior.

Verifying the Number of Name Server Database Entries

To Verify the number of name server database entries, follow these steps:

Procedure

Step 1 switch# show fcns internal info global

Displays the number of device entries in the name server database.

Step 2 switch# show fcns internal info

Displays the number of devices in the name server database at the end of the output.

Displaying Name Server Database Entries

Use the **show fcns** command to display the name server database and statistical information for a specified VSAN or for all VSANs (see Examples Displays the Name Server Database, on page 208 to Displays the Name Server Statistics, on page 209).

Displays the Name Server Database

```
switch# show fcns database
--------------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0x010000    N     50:06:0b:00:00:10:a7:80                   scsi-fcp fc-gs
0x010001    N     10:00:00:05:30:00:24:63 (Cisco)           ipfc
0x010002    N     50:06:04:82:c3:a0:98:52 (Company 1)       scsi-fcp 250
0x010100    N     21:00:00:e0:8b:02:99:36 (Company A)       scsi-fcp
0x020000    N     21:00:00:e0:8b:08:4b:20 (Company A)
0x020100    N     10:00:00:05:30:00:24:23 (Cisco)         ipfc
0x020200    N     21:01:00:e0:8b:22:99:36 (Company A)       scsi-fcp
```

Displays the Name Server Database for the Specified VSAN

```
switch# show fcns database vsan 1
VSAN 1:
--------------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0x030001    N     10:00:00:05:30:00:25:a3 (Cisco)         ipfc
0x030101    NL    10:00:00:00:77:99:60:2c (Interphase)
0x030200    N     10:00:00:49:c9:28:c7:01
```

```
0xec0001    NL    21:00:00:20:37:a6:be:14 (Seagate)       scsi-fcp
Total number of entries = 4
```

Displays the Name Server Database Details

```
switch# show fcns database detail
------------------------
VSAN:1     FCID:0x030001
------------------------
port-wwn (vendor)     :10:00:00:05:30:00:25:a3 (Cisco)
node-wwn              :20:00:00:05:30:00:25:9e
class                 :2,3
node-ip-addr          :0.0.0.0
ipa                   :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:ipfc
symbolic-port-name    :
symbolic-node-name    :
port-type             :N
port-ip-addr          :0.0.0.0
fabric-port-wwn       :00:00:00:00:00:00:00:00
hard-addr             :0x000000
------------------------
VSAN:1     FCID:0xec0200
------------------------
port-wwn (vendor)     :10:00:00:5a:c9:28:c7:01
node-wwn              :10:00:00:5a:c9:28:c7:01
class                 :3
node-ip-addr          :0.0.0.0
ipa                   :ff ff ff ff ff ff ff ff
fc4-types:fc4_features:
symbolic-port-name    :
symbolic-node-name    :
port-type             :N
port-ip-addr          :0.0.0.0
fabric-port-wwn       :22:0a:00:05:30:00:26:1e
hard-addr             :0x000000
Total number of entries = 2
```

Displays the Name Server Statistics

```
switch# show fcns statistics

registration requests received = 27
deregistration requests received = 0
queries received = 57
queries sent = 10
reject responses sent = 14
RSCNs received = 0
RSCNs sent = 0
```

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 Cisco 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

FDMI is compatible with both physical and virtual end devices. The number of registered virtual devices is as follows:

- Prior to Cisco MDS NX-OS 9.4(2), the maximum limit was 255 virtual devices per HBA end device.
- FromCisco MDS NX-OS 9.4(2), the maximum limit is 32 virtual devices per HBA end device.

Displaying FDMI

Use the **show fdmi** command to display the FDMI database information.

Displays All HBA Management Servers prior to Cisco MDS NX-OS Release 9.4(2)

```
switch# show fdmi database
Registered HBA List for VSAN 1
  10:00:00:00:c9:32:8d:77
  21:01:00:e0:8b:2a:f6:54
switch# show fdmi database detail
Registered HBA List for VSAN 1
--------------------------------
HBA-ID: 10:00:00:00:c9:32:8d:77
--------------------------------
Node Name        :20:00:00:00:c9:32:8d:77
Manufacturer     :Emulex Corporation
Serial Num       :0000c9328d77
Model            :LP9002
Model Description:Emulex LightPulse LP9002 2 Gigabit PCI Fibre Channel Adapter
Hardware Ver     :2002606D
Driver Ver       :SLI-2 SW_DATE:Feb 27 2003, v5-2.20a12
ROM Ver          :3.11A0
Firmware Ver     :3.90A7
OS Name/Ver      :Window 2000
CT Payload Len   :1300000
  Port-id: 10:00:00:00:c9:32:8d:77
--------------------------------
HBA-ID: 21:01:00:e0:8b:2a:f6:54
--------------------------------
Node Name        :20:01:00:e0:8b:2a:f6:54
Manufacturer     :QLogic Corporation
Serial Num       :\74262
Model            :QLA2342
Model Description:QLogic QLA2342 PCI Fibre Channel Adapter
Hardware Ver     :FC5010409-10
Driver Ver       :8.2.3.10 Beta 2 Test 1 DBG (W2K VI)
ROM Ver          :1.24
Firmware Ver     :03.02.13.
OS Name/Ver      :500
```

```
CT Payload Len   :2040
  Port-id: 21:01:00:e0:8b:2a:f6:54
```

Display all registered end devices including virtual end device information from Cisco MDS NX-OS Release 9.4(2)

```
switch# show fdmi database detail
Registered HBA List for VSAN 1
-------------------------------
HBA-ID: 10:00:00:10:9b:e0:ff:0c
-------------------------------
Node Name        :20:00:00:10:9b:e0:ff:0c
Manufacturer     :Emulex Corporation
Serial Num       :FP13662272
Model            :LPe36002-M64
Model Description:Emulex LPe36002-M64 2-Port 64Gb Fibre Channel Adapter
Hardware Ver     :0000000
Driver Ver       :12.6.0.2
ROM Ver          :12.8.351.47
Firmware Ver     :12.8.351.47
OS Name/Ver      :Linux 4.18.0-193.el8.x86_64 #1 SMP Fri Mar 27 14:35:58 UTC 2020
CT Payload Len   :245760
  Port-id: 10:00:00:10:9b:e0:ff:0c
    Supported FC4 types:1 scsi-fcp fc-gs NVMe
    Supported Speed    :16G 32G 64G
    Current Speed      :32G
    Maximum Frame Size :2048
    OS Device Name     :/sys/class/scsi_host/host13
    Host Name          :localhost.localdomain
-------------------------------
HBA-ID: 21:00:00:24:ff:7e:e6:14
-------------------------------
Node Name        :20:00:00:24:ff:7e:e6:14
Manufacturer     :QLogic Corporation
Serial Num       :RFD1604J61197
Model            :QLE2742
Model Description:Cisco QLE2742 Dual Port 32Gb FC to PCIe Gen3 x8 Adapter
Hardware Ver     :BK3210407-43 02
Driver Ver       :10.01.00.63.08.0-k
ROM Ver          :3.62
Firmware Ver     :8.08.05 (d0d5)
  Port-id: 21:00:00:24:ff:7e:e6:14
    Supported FC4 types:scsi-fcp
    Supported Speed    :8G 16G 32G
    Current Speed      :16G
    Maximum Frame Size :2048
    OS Device Name     :qla2xxx:host9
    Host Name          :localhost.localdomain
  Port-id: 21:04:00:24:ff:7e:e6:14
    Supported FC4 types:scsi-fcp
    Supported Speed    :8G 16G 32G
    Current Speed      :16G
    Maximum Frame Size :2048
    OS Device Name     :qla2xxx:host15
    Host Name          :localhost.localdomain
  Port-id: 21:05:00:24:ff:7e:e6:14
    Supported FC4 types:scsi-fcp
    Supported Speed    :8G 16G 32G
    Current Speed      :16G
    Maximum Frame Size :2048
    OS Device Name     :qla2xxx:host16
```

```
    Host Name          :localhost.localdomain
  Port-id: 21:06:00:24:ff:7e:e6:14
    Supported FC4 types:scsi-fcp
    Supported Speed    :8G 16G 32G
    Current Speed      :16G
    Maximum Frame Size :2048
    OS Device Name     :qla2xxx:host17
    Host Name          :localhost.localdomain
  Port-id: 21:07:00:24:ff:7e:e6:14
    Supported FC4 types:scsi-fcp
    Supported Speed    :8G 16G 32G
    Current Speed      :16G
    Maximum Frame Size :2048
    OS Device Name     :qla2xxx:host18
    Host Name          :localhost.localdomain
```

Displays HBA Details for a Specified VSAN

```
switch# show fdmi database detail vsan 1
Registered HBA List for VSAN 1
--------------------------------
HBA-ID: 10:00:00:00:c9:32:8d:77
--------------------------------
Node Name :20:00:00:10:9b:e0:ff:0c
Manufacturer :Emulex Corporation
Serial Num :FP13662272
Model :LPe36002-M64
Model Description:Emulex LPe36002-M64 2-Port 64Gb Fibre Channel Adapter
Hardware Ver :0000000
Driver Ver :12.6.0.2
ROM Ver :12.8.351.47
Firmware Ver :12.8.351.47
OS Name/Ver :Linux 4.18.0-193.el8.x86_64 #1 SMP Fri Mar 27 14:35:58 UTC 2020
CT Payload Len :245760
Port-id: 10:00:00:10:9b:e0:ff:0c
Supported FC4 types:1 scsi-fcp fc-gs NVMe
Supported Speed :16G 32G 64G
Current Speed :32G
Maximum Frame Size :2048
OS Device Name :/sys/class/scsi_host/host13
Host Name :localhost.localdomain
  Port-id: 10:00:00:00:c9:32:8d:77
--------------------------------
HBA-ID: 21:01:00:e0:8b:2a:f6:54
--------------------------------
Node Name        :20:00:00:24:ff:7e:e6:14
Manufacturer     :QLogic Corporation
Serial Num       :RFD1604J61197
Model            :QLE2742
Model Description:Cisco QLE2742 Dual Port 32Gb FC to PCIe Gen3 x8 Adapter
Hardware Ver     :BK3210407-43 02
Driver Ver       :10.01.00.63.08.0-k
ROM Ver          :3.62
Firmware Ver     :8.08.05 (d0d5)
  Port-id: 21:01:00:e0:8b:2a:f6:54
```

Displays Details for the Specified HBA Entry

```
switch# show fdmi database detail hba-id 21:01:00:e0:8b:2a:f6:54 vsan 1
```

```
Node Name        :20:01:00:e0:8b:2a:f6:54
Manufacturer     :QLogic Corporation
Serial Num       :\74262
Model            :QLA2342
Model Description:QLogic QLA2342 PCI Fibre Channel Adapter
Hardware Ver     :FC5010409-10
Driver Ver       :8.2.3.10 Beta 2 Test 1 DBG (W2K VI)
ROM Ver          :1.24
Firmware Ver     :03.02.13.
OS Name/Ver      :500
CT Payload Len   :2040
  Port-id: 21:01:00:e0:8b:2a:f6:54
```

VMID

Note The VMID feature is currently in preview (beta) status for use in non-production environment only. This preview (beta) status and restriction will change to regular production status in an upcoming release.

The switch-based Virtual Machine Identifier (VMID) feature allows identification of traffic sources at an individual virtual machine (VM) level by the SAN fabric infrastructure.

VMID on the MDS switch provides a range of identifiers to a host hypervisor. These identifiers can then be assigned to local VMs by the hypervisor. Supplemental information about the VM assigned to an identifier is reported back to the switch. The identifiers are then inserted into the CS_CTL field of traffic from the VMs by the hypervisor, allowing identification of the traffic source by the SAN fabric.

The VMID feature uses the following IDs:

- Virtual Entity (VE): Refers to any virtual device.
- Virtual Entities Manager (VEM): Refers to a hypervisor.
- Virtual Entity Identifier (VE ID): Refers to the different types of identifiers assigned to VEs. There are four types of VE IDs:
  - Local VE ID: Local VE ID is used to uniquely identify a VE within a VEM N_Port. Local VE IDs change when virtual machines come up, go down, or go through migration between VEMs.
  - Fabric VE ID: Fabric VE ID is used to uniquely identify a VE within a fabric. It is a combination of the VEM N_Port FCID and the Local VE ID.
  - Global VE ID: Global VE ID is used to uniquely identify a VE and is a 16-byte Universally Unique Identifier (UUID). The Global VE ID is assigned by a service outside the SAN fabric such as a VM management platform. Once a global VE ID is assigned, it does not expire.
  - VEM ID: VEM ID is used to uniquely identify a VEM and is a 16-byte UUID. The VEM ID is assigned by a service outside the SAN fabric such as a VM management platform.
- Fabric ports on a VEM comprise of the following N_Ports:
  - Physical Network Port (PN_Port): A physical network port of a hypervisor host bus adapter (HBA).
  - Virtual Network Port (VN_Port): An optional virtual network port that can be shared by a set of VEs. A PN_Port can have multiple VN_Ports. Each VN_Port is allocated a unique FCID.

- Physical Fabric Port (PF_Port): A physical fabric port of a switch.

When an HBA Port Comes Up

After an HBA driver has logged a physical or virtual HBA port into the fabric, the driver may request Local VE IDs from the fabric through the port. The Virtual Machine Identification Server (VMIS) on the locally attached switch provides the range of Local VE IDs (up to 255) in the response. The driver then assigns the identifiers to the port FCID.

When a VM Initially Accesses a Disk

Outside of a fabric, a VM is identified by a Global VE ID. Within a fabric, a VM is identified by a Fabric VE ID. When a VM initially accesses a virtual disk, the hypervisor starts accessing the corresponding physical disk through an HBA port. For each path to the physical disk an unused Local VE ID from the pool for the FCID is assigned. The FCID and Local VE ID are combined to create a unique Fabric VE ID by the HBA driver. The HBA driver then notifies the locally attached VMIS about the assigned VE ID to Global VE ID mapping. This mapping is done for each VM path to the fabric through the hypervisor and allows all VM traffic to be identified by SAN tools by path, such as Cisco MDS SAN Analytics.

When a VM Goes Down or Moves Within a Fabric

The difference between a Global VE ID and Local VE ID is seen when a VM migrates between VEMs. When a VM is deinstantiated or migrated between VEMs, the Local VE ID is returned to the pool by the HBA driver, but does not notify the switch VMIS. The switch will time out the Local VE ID or VM mapping after 1 to 4 hours of no traffic. If a VM reinstantiates on the same VEM, it may get a different Local VE ID from the pool of the same FCID because the previously assigned Local VE ID may have been allocated to another VM when this VM was down. When a VM migrates to a different VEM, the VM will likely utilize a different FCID and likely be assigned a different Local VE ID from the FCID's pool. Thus, when a VM reboots or migrates between VEMs, the Global VE ID remains constant but Local VE ID may change.

Figure 21: VMID Components, on page 215 show the components of VMID:

Figure 21: VMID Components

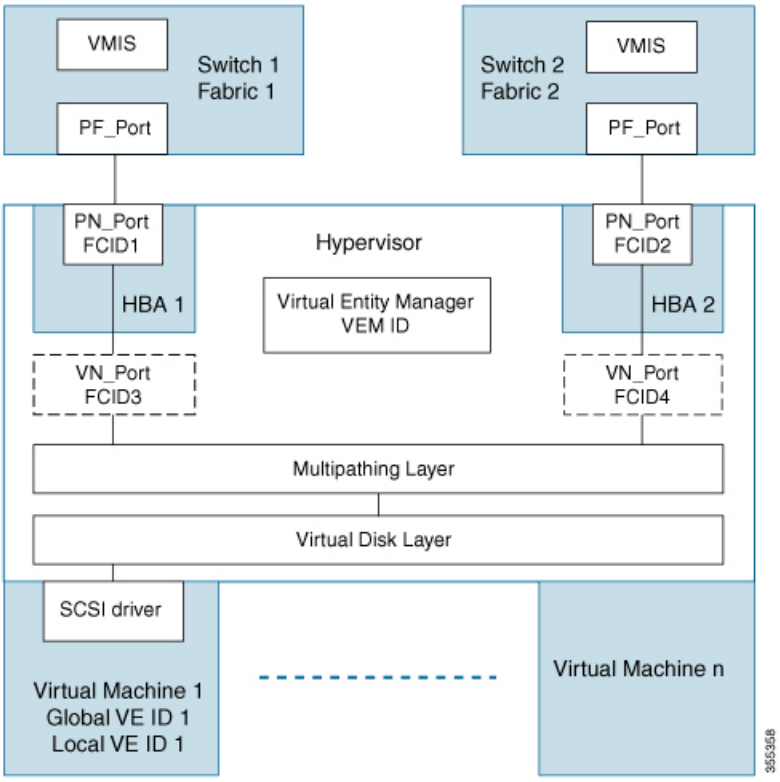
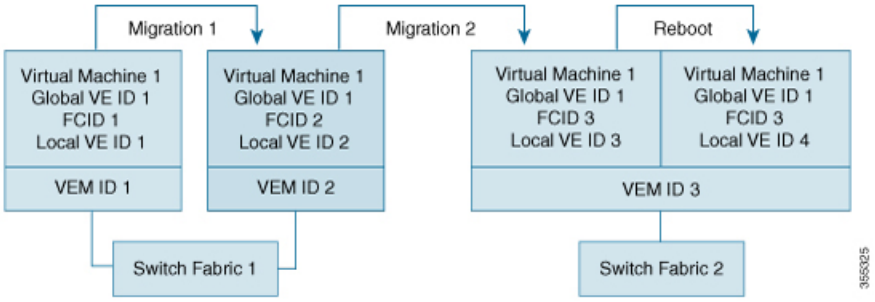


Figure 22: VE ID Life Cycle, on page 215 shows how VE ID changes during a VM life cycle:

Figure 22: VE ID Life Cycle



Guidelines and Limitations for VMID

- The VMID feature is not supported on Cisco N-Port Virtualizer (Cisco NPV) switches.
- There is no mechanism in the VMID protocol for the VMIS to notify the attached hypervisor HBA driver clients of a new VE ID range. For clients to detect a new range, they must query the VMIS again. To force the clients to query again after a range modification, the user must manually log the FCIDs out and back into the fabric. Consequently, local clients will continue to tag the VM traffic with the previous

range until this occurs. This restriction applies when enabling and disabling VMID, and changing the VE ID range of a VSAN.

- The Extended Receiver Ready (ER_RDY) feature uses CSCTL 1 to 15. The VMID feature uses CSCTL 16 to 255. If the VMID database has any interfaces that are configured in the VMIS range 1 to 15 and if you are upgrading to Cisco MDS NX-OS Release 9.2(1) or later releases, then you will prompted to change the range between 16 to 255 and flap the interfaces before upgrading.
- The VMID feature is not supported in any VSAN which has interoperability enabled. For more information about interoperability modes, see the Cisco MDS 9000 Series Switch-to-Switch Interoperability Configuration Guide.

Configuring the VMID Server

Enabling the VMID Server

To enable the VMID Server feature, perform these steps:

Procedure

Step 1 Enter global configuration mode:

switch# **configure terminal**

Step 2 Enable the VMID Server feature:

switch(config)# **feature vmis**

Disabling the VMID Server

To disable the VMID Server feature, perform these steps:

Procedure

Step 1 Enter global configuration mode:

switch# **configure terminal**

Step 2 Disable the VMID Server feature:

switch(config)# **no feature vmis**

Configuring a VMID Range

The VMID range is used to limit the Local VE IDs an HBA driver will use. By restricting the Local VE ID range to use a subset of bits in the CS_CTL field it can be partitioned and shared with future Fibre Channel features.

To configure a VMID range, perform these steps:

Procedure

Step 1 Enter global configuration mode:

switch# **configure terminal**

Step 2 Configure VE ID range for use within a VSAN:

switch(config)# **vmis range** *range* **vsan** *id*

Examples: Configuring the VMID Server

This example shows how to enable the VMID Server feature:

```
switch# configure terminal
switch(config)# feature vmis
```

This example shows how to disable the VMID Server feature:

```
switch# configure terminal
switch(config)# no feature vmis
```

This example shows how to configure multiple Local VE ID ranges for use by hypervisor HBA drivers in a VSAN:

```
switch# configure terminal
switch(config)# vmis range 3-45,51-70 vsan 1
```

Verifying a VMID Configuration

This example shows the FCIDs that are capable of using the VMID Server feature. The letter *M* under the **FLAGS** field indicates that the corresponding FCID is capable of using the VMID Server feature.

```
switch# show flogi database details
--------------------------------------------------------------------------------
INTERFACE        VSAN    FCID           PORT NAME               NODE NAME           FLAGS

--------------------------------------------------------------------------------
fc1/7            1     0xef0000  20:07:8c:60:4f:10:0f:e0 20:01:8c:60:4f:10:0f:e1     P
fc1/7            1     0xef0001  20:19:8c:60:4f:19:bf:25 21:00:00:20:38:de:c3:9f     VPM
```

```
Total number of flogi = 2.
```

This example shows all the entries in a VMIS database. This is the database of all IDs in the SAN fabric. Locally connected IDs show the connecting interface; remotely connected IDs show the interface name as "--" in the output.

```
switch# show vmis database
Total 17 entries
--------------------------------------------------------------------------------
INTERFACE         VSAN  FCID      LOCAL VEID           GLOBAL VEID
--------------------------------------------------------------------------------
fc1/7             1     0xef000a  0x01        9a07686b-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000a  0x02        66fb6a4e-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000a  0x03        325de425-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000a  0x04        0d509b51-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000a  0x05        b7d71b43-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000a  0x32        1b231602-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000b  0x01        e8e9161f-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000b  0x02        e7cd9011-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000b  0x03        8d43ef66-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000b  0x04        760f0e14-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000b  0x05        5a255233-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000b  0x1e        1b231602-0405-0607-0809-0a0b0c0d0e0f
--                10    0x4c0020  0x1e        ba581b3d-0405-0607-0809-0a0b0c0d0e0f
--                10    0x4c0020  0x1f        abd77e50-0405-0607-0809-0a0b0c0d0e0f
--                10    0x4c0020  0x20        f241b12e-0405-0607-0809-0a0b0c0d0e0f
--                10    0x4c0020  0x21        fb1eb741-0405-0607-0809-0a0b0c0d0e0f
--                10    0x4c0020  0x22        e3a9e279-0405-0607-0809-0a0b0c0d0e0f
```

This example shows a VMIS database entries of a specified local VSAN domain:

```
switch# show vmis database local vsan 1
Total 12 entries
--------------------------------------------------------------------------------
INTERFACE         VSAN  FCID      LOCAL VEID           GLOBAL VEID
--------------------------------------------------------------------------------
fc1/7             1     0xef000a  0x01        9a07686b-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000a  0x02        66fb6a4e-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000a  0x03        325de425-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000a  0x04        0d509b51-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000a  0x05        b7d71b43-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000a  0x32        1b231602-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000b  0x01        e8e9161f-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000b  0x02        e7cd9011-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000b  0x03        8d43ef66-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000b  0x04        760f0e14-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000b  0x05        5a255233-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000b  0x1e        1b231602-0405-0607-0809-0a0b0c0d0e0f
```

This example shows the entries in a VSAN filtered by the hosting domain:

```
switch# show vmis database domain 0xef vsan 1
Total 12 entries
--------------------------------------------------------------------------------
INTERFACE         VSAN  FCID      LOCAL VEID           GLOBAL VEID
--------------------------------------------------------------------------------
fc1/7             1     0xef000a  0x01        9a07686b-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000a  0x02        66fb6a4e-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000a  0x03        325de425-0405-0607-0809-0a0b0c0d0e0f
```

```
fc1/7             1     0xef000a  0x04        0d509b51-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000a  0x05        b7d71b43-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000a  0x32        1b231602-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000b  0x01        e8e9161f-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000b  0x02        e7cd9011-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000b  0x03        8d43ef66-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000b  0x04        760f0e14-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000b  0x05        5a255233-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000b  0x1e        1b231602-0405-0607-0809-0a0b0c0d0e0f
```

This example shows the entries in a VSAN filtered by an interface:

```
switch# show vmis database interface fc1/7 vsan 1
Total 12 entries
--------------------------------------------------------------------------------
INTERFACE         VSAN  FCID      LOCAL VEID           GLOBAL VEID
--------------------------------------------------------------------------------
fc1/7             1     0xef000a  0x01        9a07686b-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000a  0x02        66fb6a4e-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000a  0x03        325de425-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000a  0x04        0d509b51-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000a  0x05        b7d71b43-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000a  0x32        1b231602-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000b  0x01        e8e9161f-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000b  0x02        e7cd9011-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000b  0x03        8d43ef66-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000b  0x04        760f0e14-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000b  0x05        5a255233-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000b  0x1e        1b231602-0405-0607-0809-0a0b0c0d0e0f
```

This example shows the entries in a VSAN:

```
switch# show vmis database vsan 10
Total 5 entries
--------------------------------------------------------------------------------
INTERFACE         VSAN  FCID      LOCAL VEID           GLOBAL VEID
--------------------------------------------------------------------------------
--                10    0x4c0020  0x1e        ba581b3d-0405-0607-0809-0a0b0c0d0e0f
--                10    0x4c0020  0x1f        abd77e50-0405-0607-0809-0a0b0c0d0e0f
--                10    0x4c0020  0x20        f241b12e-0405-0607-0809-0a0b0c0d0e0f
--                10    0x4c0020  0x21        fb1eb741-0405-0607-0809-0a0b0c0d0e0f
--                10    0x4c0020  0x22        e3a9e279-0405-0607-0809-0a0b0c0d0e0f
```

This example shows the entries filtered by FCIDs. This example is filtered by a remote hypervisor N_Port FCID.

```
switch# show vmis database fcid 0x4c0020 vsan 10
Total 5 entries
--------------------------------------------------------------------------------
INTERFACE         VSAN  FCID      LOCAL VEID           GLOBAL VEID
--------------------------------------------------------------------------------
--                10    0x4c0020  0x1e        ba581b3d-0405-0607-0809-0a0b0c0d0e0f
--                10    0x4c0020  0x1f        abd77e50-0405-0607-0809-0a0b0c0d0e0f
--                10    0x4c0020  0x20        f241b12e-0405-0607-0809-0a0b0c0d0e0f
--                10    0x4c0020  0x21        fb1eb741-0405-0607-0809-0a0b0c0d0e0f
--                10    0x4c0020  0x22        e3a9e279-0405-0607-0809-0a0b0c0d0e0f
```

This example shows the VMIS entries filtered by Global VM ID and VSAN:

```
switch# show vmis database global-vmid e8e9161f-0405-0607-0809-0a0b0c0d0e0f vsan 1
Total 1 entries
--------------------------------------------------------------------------------
INTERFACE         VSAN  FCID      LOCAL VEID           GLOBAL VEID
--------------------------------------------------------------------------------
fc1/7             1     0xef000b  0x01        e8e9161f-0405-0607-0809-0a0b0c0d0e0f
```

This example shows the VEM IDs registered in a VSAN:

```
switch# show vmis database vem vsan 1
Total 2 entries
--------------------------------------------------------------------------------
INTERFACE         VSAN  FCID               VEM ID
--------------------------------------------------------------------------------
fc1/7             1     0xef000a  11223344-5566-7788-99aa-bbccddeeffaa
fc1/7             1     0xef000b  00010203-0405-0607-0809-0a0b0cef000b
```

This example shows VM entries that have migrated between VEMs:

The output shows two entries correspond to a VM before and after a VM has migrated between VEMs. The IDs associated with the VM before migration are not deleted immediately. These IDs will be deleted in the VMIS database after the switch IO timer expires. Until the IO timer expires, you will see two entries of the same VM in the VMIS database.

```
switch# show vmis database vmotion vsan 1
Total 2 entries
--------------------------------------------------------------------------------
INTERFACE         VSAN  FCID      LOCAL VEID           GLOBAL VEID
--------------------------------------------------------------------------------
fc1/7             1     0xef000b  0x1e        1b231602-0405-0607-0809-0a0b0c0d0e0f
fc1/7             1     0xef000a  0x32        1b231602-0405-0607-0809-0a0b0c0d0e0f
```

This example shows a Local VE ID range that is configured for each VSAN:

```
switch# show vmis range
VSAN      VEID Range
--------------------
   1      1-255
  10      1-255
  20      1-255
  30      1-255
```

This example shows statistics of a local switch's VMIS exchanges with locally attached hypervisor HBA driver clients (host side) and with other VMIS agents on other switches in the fabric (switch side) by VSAN:

```
switch# show vmis statistics
VSAN : 1
-----Host Side-----
qfpa/qfpa_rsp/qfpa_rjt : 1/1/0
uvem/uvem_rsp/uvem_rjt : 1/1/0
ggvid/ggvid_rsp/ggvid_rjt : 0/0/0
gfvid/gfvid_rsp/gfvid_rjt : 0/0/0
gvemid/gvemid_rsp/gvemid_rjt : 0/0/0
gvem/gvem_rsp/gvem_rjt : 0/0/0

-----Switch Side-----
```

```
gvemd_tx/gvemd_rsp_tx/gvemd_rjt_tx : 0/0/0
gvemd_rx/gvemd_rsp_rx/gvemd_rjt_rx : 0/0/0
uvemd_tx/uvemd_rsp_tx/uvemd_rjt_tx : 0/0/0
uvemd_rx/uvemd_rsp_rx/uvemd_rjt_rx : 0/0/0
```

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:

- Zoned end devices joining or leaving the fabric
- Name server registration change
- New zone enforcement
- An IP address change
- Any other similar event that affects the operation of the host

This section includes the following topics:

About RSCN Information

The RSCN service is provided by each MDS switch. End devices register with the local switch's Fabric Controller for the types of notifications they want to receive. When a registered event occurs, the Fabric Controller sends an RSCN to all locally attached end devices that have requested it. It also sends an SW-RSCN to all the remote switches in the fabric. The remote switches then send RSCNs to their locally attached nodes as needed.

An RSCN Request contains an event type and one or more FCIDs of affected nodes. The receiver should respond with either Accept or Reject.

After a node receives an RSCN, it is the responsibility of the node to query the FC Name Server to obtain updated information about the affected node. The RSCN itself does not contain the details of the event.

The RSCN feature configuration can be distributed by the CS service.

MDS switches support the following registration functions in the State Change Registration request:

- Fabric-detected registration
- Nx_Port-detected registration

If an SCR is sent to the MDS Fabric Controller requesting registration containing a request to register for one or more unsupported functions, the SCR will be rejected and none of the registration functions will be accepted. The following type of syslog will be logged showing the invalid registration function:

```
%RSCN-2-SCR_REJECT: %$VSAN 20%$ SCR rejected: Sending LS-RJT: Invalid registration function
 0xb from fcid 0xbb0601, Reason code-Logical error, Reason code expl-Request not supported
```

Displaying RSCN Information

Use the **show rscn** command to display RSCN information.

Displays RSCN Counter Information

```
switch(config)# show rscn statistics vsan 106

Statistics for VSAN: 106
------------------------
Number of SCR received         = 0
Number of SCR ACC sent         = 0
Number of SCR RJT sent         = 0
Number of RSCN received        = 0
Number of RSCN sent            = 0
Number of RSCN ACC received    = 0
Number of RSCN ACC sent        = 0
Number of RSCN RJT received    = 0
Number of RSCN RJT sent        = 0
Number of SW-RSCN received     = 0
Number of SW-RSCN sent         = 0
Number of SW-RSCN ACC received = 0
Number of SW-RSCN ACC sent     = 0
Number of SW-RSCN RJT received = 0
Number of SW-RSCN RJT sent     = 0
Number of CSWR received        = 3137
Number of CSWR sent            = 0
Number of CSWR ACC received    = 0
Number of CSWR ACC sent        = 3137
Number of CSWR RJT received    = 0
Number of CSWR RJT sent        = 0
Number of CSWR RJT not sent    = 0
```

Displaying Device Registration Information

```
switch# show rscn scr-table
Supported SCR functions:
  Fabric Detected registration
  Nport Detected registration

SCR table for VSAN: 1
FC-IDREGISTERED FOR
0x1b0300 fabric detected rscns
Total number of entries = 1
```

Note The SCR table is not configurable. It is dynamically updated with the current registrations from end devices. If an end device does not send an SCR request to the Fabric Controller then the **show rscn scr-table** command will not display any entries for that device.

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).



Note Some Nx ports may not understand multi-pid RSCN payloads. If not, disable the RSCN **multi-pid** option.

Configuring the multi-pid Option

To configure the **multi-pid** option, follow these steps:

Procedure

Step 1 switch# **config terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **rscn multi-pid vsan 105**

Sends RSCNs in a multi-pid format for VSAN 105.

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).

To suppress the transmission of these SW RSCNs over an ISL, follow these steps:

Procedure

Step 1 switch# **config terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **rscn suppress domain-swrscn vsan 105**

Suppresses transmission of domain format SW-RSCNs for VSAN 105.

Note You cannot suppress transmission of port address or area address format RSCNs.

Coalesced SW-RSCN

In order to improve the performance of the Fibre Channel protocols on the Cisco MDS 9000 switch, SW-RSCNs are delayed, collected and sent as a single coalesced SW-RSCN to all the switches in the fabric in a single Fibre Channel exchange.

Enabling Coalesced SW-RSCNs

Restrictions

- All the switches in the fabric should be running Cisco MDS 6.2(7) and above.
- This feature does not have interoperability with non-Cisco MDS switches.

To enable the coalesced SW-RSCNs, follow these step:

Procedure

Step 1 switch# **config terminal**

Enters configuration mode.

Step 2 switch(config)# rscn coalesce swrscn vsan 1

switch(config)#

Enables coalescing of Switch Registered State Change Notification (SWRSCN) in VSAN 1. The default delay is 500 milliseconds.

Step 3 switch(config)# rscn coalesce swrscn vsan 1 delay 800

switch(config)#

Enables coalescing of Switch Registered State Change Notification (SWRSCN) in VSAN 1. Delays the SW-RSCNs maximum by 800 milliseconds.



Note All the switches running 6.2(7) and above are capable of processing coalesced SW-RSCN by default, but they are capable of sending coalesced SW-RSCN only after enabling through CLI.

Disabling Coalesced SW-RSCNs

To disable the coalesced SW-RSCNs, follow these steps:

Procedure

Step 1 switch# **config terminal**

Enters configuration mode.

Step 2 switch(config)# no rscn coalesce swrscn vsan 1

switch(config)#

Disables coalescing of Switch Registered State Change Notification (SWRSCN) in VSAN 1.

Clearing RSCN Statistics

You can clear the counters and later view the counters for a different set of events. For example, you can keep track of how many RSCNs or SW-RSCNs are generated on a particular event (such as ONLINE or OFFLINE events). You can use these statistics to monitor responses for each event in the VSAN.

Use the **clear rscn statistics** command to clear the RSCN statistics for the specified VSAN.

```
switch# clear rscn statistics vsan 1
```

After clearing the RSCN statistics, you can view the cleared counters by issuing the **show rscn** command.

```
switch# show rscn statistics vsan 1
Statistics for VSAN: 1
-----------------------
Number of SCR received         = 0
Number of SCR ACC sent         = 0
Number of SCR RJT sent         = 0
Number of RSCN received        = 0
Number of RSCN sent            = 0
Number of RSCN ACC received    = 0
Number of RSCN ACC sent        = 0
Number of RSCN RJT received    = 0
Number of RSCN RJT sent        = 0
Number of SW-RSCN received     = 0
```

```
Number of SW-RSCN sent         = 0
Number of SW-RSCN ACC received = 0
Number of SW-RSCN ACC sent     = 0
Number of SW-RSCN RJT received = 0
Number of SW-RSCN RJT sent     = 0
Number of CSWR received        = 0
Number of CSWR sent            = 0
Number of CSWR ACC received    = 0
Number of CSWR ACC sent        = 0
Number of CSWR RJT received    = 0
Number of CSWR RJT sent        = 0
Number of CSWR RJT not sent    = 0
```

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.

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 RCSN 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.

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, on page 227.

Note Before performing a downgrade, make sure that you revert the RCSN timer value in your network to the default value. Failure to do so will disable the links across your VSANs and other devices.

To configure the RSCN timer, follow these steps:

Procedure

Step 1 switch# **config t**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **rscn distribute**

Enables RSCN timer configuration distribution.

Step 3 switch(config)# **rscn event-tov 300 vsan 10**

Sets the event time-out value in milliseconds for the selected VSAN. In this example, the event time-out value is set to 300 milliseconds for VSAN 12. The range is 0 to 2000 milliseconds. Setting a zero (0) value disables the timer.

Step 4 switch(config)# **no rscn event-tov 300 vsan 10**

Reverts to the default value (2000 milliseconds for Fibre Channel VSANs or 1000 milliseconds for FICON VSANs).

Step 5 switch(config)# **rscn commit vsan 10**

Commits the RSCN timer configuration to be distributed to the switches in VSAN 10.

Verifying the RSCN Timer Configuration

You verify the RSCN timer configuration using the **show rscn event-tov vsan** command.

```
switch# show rscn event-tov vsan 10
Event TOV : 1000 ms
```

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 N-ports 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 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.

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).

This section includes the following topics:

Enabling RSCN Timer Configuration Distribution

To enable RSCN timer configuration distribution, follow these steps:

Procedure

Step 1 switch# **config terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **rscn distribute**

Enables RSCN timer distribution.

Step 3 switch(config)# **no rscn distribute**

Disables (default) RSCN timer distribution.

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.

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.

To commit RSCN timer configuration changes, follow these steps:

Procedure

Step 1 switch# **config t**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **rscn commit vsan 10**

Commits the RSCN timer changes.

Discarding the RSCN Timer Configuration Changes

If you discard (terminate) the changes made to the pending database, the configuration database remains unaffected and the lock is released.

To discard RSCN timer configuration changes, follow these steps:

Procedure

Step 1 switch# **config t**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **rscn abort vsan 10**

Discards the RSCN timer changes and clears the pending configuration database.

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.

To use administrative privileges and release a locked DPVM session, use the **clear rscn session vsan** command in EXEC mode.

```
switch# clear rscn session vsan 10
```

Displaying RSCN Configuration Distribution Information

Use the **show cfs application name rscn** command to display the registration status for RSCN configuration distribution.

```
switch# show cfs application name rscn
 Enabled        : Yes
 Timeout        : 5s
 Merge Capable  : Yes
 Scope          : Logical
```

Use the **show rscn session status vsan** command to display session status information for RSCN configuration distribution.



Note A merge failure results when the RSCN timer values are different on the merging fabrics.

```
switch# show rscn session status vsan 1
Session Parameters for VSAN: 1
------------------------------
Last Action                : Commit
Last Action Result         : Success
Last Action Failure Reason : None
```

Use the **show rscn pending** command to display the set of configuration commands that would take effect when you commit the configuration.

Note The pending database includes both existing and modified configuration.

```
switch# show rscn pending
rscn event-tov 2000 ms vsan 1
rscn event-tov 2000 ms vsan 2
rscn event-tov 300 ms vsan 10
```

Use the **show rscn pending-diff** command to display the difference between pending and active configurations. The following example shows the time-out value for VSAN 10 was changed from 2000 milliseconds (default) to 300 milliseconds.

```
switch# show rscn pending-diff
- rscn event-tov 2000 ms vsan 10
+ rscn event-tov 300 ms vsan 10
```

Default Settings

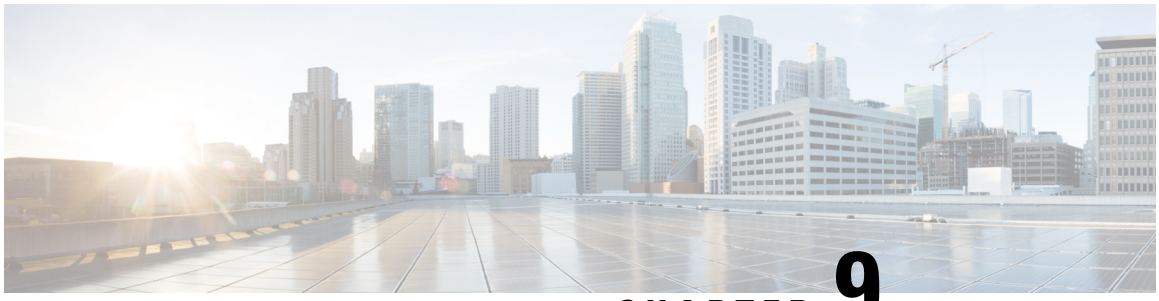
Table 23: Default RSCN Settings , on page 231 lists the default settings for RSCN.

Table 23: Default RSCN Settings

Parameters	Default
RSCN timer value	2000 milliseconds for Fibre Channel VSANs1000 milliseconds for FICON VSANs
RSCN timer configuration distribution	Disabled

Enabling Port Pacing

For detailed information, refer to the *Cisco MDS 9000 Family NX-OS System Management* .



CHAPTER 9

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:

- About SCSI LUN Discovery, on page 233
- Displaying SCSI LUN Information, on page 235

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

SCSI LUN discovery is done on demand.

Only Nx ports that are present in the name server database and that are registered as FC4 Type = SCSI_FCP are discovered.

Starting SCSI LUN Discovery

To start SCSI LUN discovery, follow one of these steps:

Procedure

Step 1 switch# **discover scsi-target local os all**

Example:

```
discovery started
```

Discovers local SCSI targets for all operating systems (OS). The operating system options are **aix**, **all**, **hpux**, **linux**, **solaris**, or **windows**

Step 2 switch# **discover scsi-target remote os aix**

Example:

```
discovery started
```

Discovers remote SCSI targets assigned to the AIX OS.

Step 3 switch# **discover scsi-target vsan 1 fcid 0x9c03d6**

Example:

```
discover scsi-target vsan 1 fcid 0x9c03d6
VSAN:    1 FCID: 0x9c03d6 PWWN: 00:00:00:00:00:00:00:00
PRLI RSP: 0x01 SPARM: 0x0012
SCSI TYPE: 0 NLUNS: 1
Vendor: Company 4 Model: ST318203FC    Rev: 0004
Other: 00:00:02:32:8b:00:50:0a
```

Discovers SCSI targets for the specified VSAN (1) and FC ID (0x9c03d6).

Step 4 switch# **discover scsi-target custom-list os linux**

Example:

```
discovery started
```

Discovers SCSI targets from the customized list assigned to the Linux OS.

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.

Use the **custom-list** option to initiate this discovery.

Initiating Customized Discovery

To initiate a customized discovery, follow one of these steps:

Procedure

Step 1 switch# **discover custom-list add vsan 1 domain 0X123456**

Adds the specified entry to the custom list.

Step 2 switch# **discover custom-list delete vsan 1 domain 0X123456**

Deletes the specified domain ID from the custom list.

Displaying SCSI LUN Information

Use the **show scsi-target** and **show fcns database** commands to display the results of the discovery. See Examples Displays the Discovered Targets, on page 235 to Displays Automatically Discovered Targets, on page 237.

Displays the Discovered Targets

```
switch# show scsi-target status
discovery completed
```

Note This command takes several minutes to complete, especially if the fabric is large or if several devices are slow to respond.

Displays the FCNS Database

```
switch# show fcns database
VSAN 1:
--------------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0xeb0000    N     21:01:00:e0:8b:2a:f6:54 (Qlogic)        scsi-fcp:init
0xeb0201    NL    10:00:00:00:c9:32:8d:76 (Emulex)        scsi-fcp:init
Total number of entries = 2
VSAN 7:
--------------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0xed0001    NL    21:00:00:04:cf:fb:42:f8 (Seagate)       scsi-fcp:target
Total number of entries = 1
VSAN 2002:
--------------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0xcafe00    N     20:03:00:05:30:00:2a:20 (Cisco)         FICON:CUP
Total number of entries = 1
```

Displays the Discovered Target Disks

```
switch# show scsi-target disk
--------------------------------------------------------------------------------
VSAN    FCID        PWWN                      VENDOR    MODEL             REV
--------------------------------------------------------------------------------
1       0x9c03d6    21:00:00:20:37:46:78:97   Company 4 ST318203FC        0004
1       0x9c03d9    21:00:00:20:37:5b:cf:b9   Company 4 ST318203FC        0004
1       0x9c03da    21:00:00:20:37:18:6f:90   Company 4 ST318203FC        0004
1       0x9c03dc    21:00:00:20:37:5a:5b:27   Company 4 ST318203FC        0004
1       0x9c03e0    21:00:00:20:37:36:0b:4d   Company 4 ST318203FC        0004
1       0x9c03e1    21:00:00:20:37:39:90:6a   Company 4 ST318203 CLAR18   3844
1       0x9c03e2    21:00:00:20:37:18:d2:45   Company 4 ST318203 CLAR18   3844
1       0x9c03e4    21:00:00:20:37:6b:d7:18   Company 4 ST318203 CLAR18   3844
1       0x9c03e8    21:00:00:20:37:38:a7:c1   Company 4 ST318203FC        0004
1       0x9c03ef    21:00:00:20:37:18:17:d2   Company 4 ST318203FC        0004
```

Displays the Discovered LUNs for All Operating Systems

```
switch# show scsi-target lun os all
ST336607FC from SEAGATE (Rev 0006)
FCID is 0xed0001 in VSAN 7, PWWN is 21:00:00:04:cf:fb:42:f8
--------------------------------------------------------------------------------
OS  LUN    Capacity Status  Serial Number    Device-Id
           (MB)
--------------------------------------------------------------------------------
WIN 0x0    36704    Online  3JA1B9QA00007338 C:1 A:0 T:3 20:00:00:04:cf:fb:42:f8
AIX 0x0    36704    Online  3JA1B9QA00007338 C:1 A:0 T:3 20:00:00:04:cf:fb:42:f8
SOL 0x0    36704    Online  3JA1B9QA00007338 C:1 A:0 T:3 20:00:00:04:cf:fb:42:f8
LIN 0x0    36704    Online  3JA1B9QA00007338 C:1 A:0 T:3 20:00:00:04:cf:fb:42:f8
HP  0x0    36704    Online  3JA1B9QA00007338 C:1 A:0 T:3 20:00:00:04:cf:fb:42:f8
```

Displays the Discovered LUNs for the Solaris OS

```
switch# show scsi-target lun os solaris
ST336607FC from SEAGATE (Rev 0006)
FCID is 0xed0001 in VSAN 7, PWWN is 21:00:00:04:cf:fb:42:f8
--------------------------------------------------------------------------------
OS  LUN    Capacity Status  Serial Number    Device-Id
           (MB)
--------------------------------------------------------------------------------
SOL 0x0    36704    Online  3JA1B9QA00007338 C:1 A:0 T:3 20:00:00:04:cf:fb:42:f8
```

The following command displays the port WWN that is assigned to each OS (Windows, AIX, Solaris, Linux, or HPUX)

Displays the pWWNs for each OS

```
switch# show scsi-target pwwn
------------------------------
OS      PWWN
------------------------------
WIN     24:91:00:05:30:00:2a:1e
```

```
AIX     24:92:00:05:30:00:2a:1e
SOL     24:93:00:05:30:00:2a:1e
LIN     24:94:00:05:30:00:2a:1e
HP      24:95:00:05:30:00:2a:1e
```

Displays Customized Discovered Targets

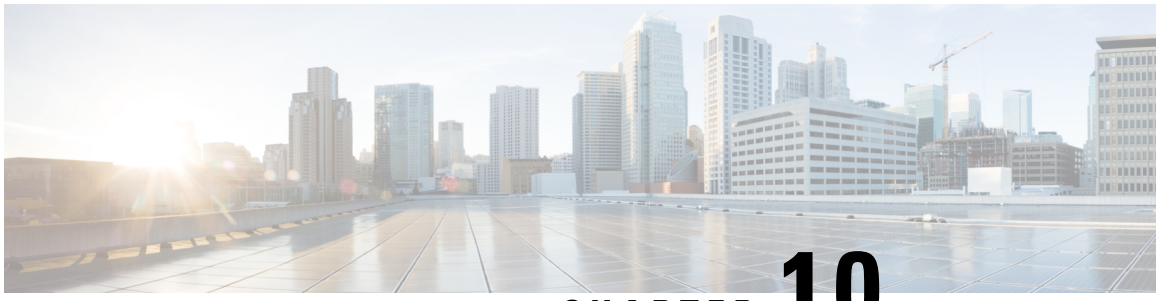
```
switch# show scsi-target custom-list
---------------
VSAN    DOMAIN
---------------
1       56
```

Use the **show scsi-target auto-poll** command to verify automatic discovery of SCSI targets that come online. The internal uuid number indicates that a CSM or an IPS module is in the chassis.

Displays Automatically Discovered Targets

```
switch(config)# show scsi-target auto-poll
name server polling is enabled
auto-polling is disabled, poll_start:0 poll_count:0 poll_type:0
USERS OF AUTO POLLING

--------------------
```



CHAPTER 10

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.

This chapter includes the following sections:

- About FICON, on page 239
- FICON Port Numbering, on page 245
- Configuring FICON, on page 252
- Configuring FICON Ports, on page 263
- FICON Configuration Files, on page 270
- Port Swapping, on page 274
- FICON Tape Acceleration, on page 274
- Configuring Zoning in a FICON VSAN, on page 277
- Moving a FICON VSAN to an Offline State, on page 278
- CUP In-Band Management, on page 278
- Displaying FICON Information, on page 279
- Default Settings, on page 284

About FICON

The Cisco MDS 9000 Family supports the Fibre Channel Protocol (FCP), FICON, iSCSI, NVMe, and FCIP capabilities within a single, high-availability platform (see Figure 23: Shared System Storage Network, on page 240).

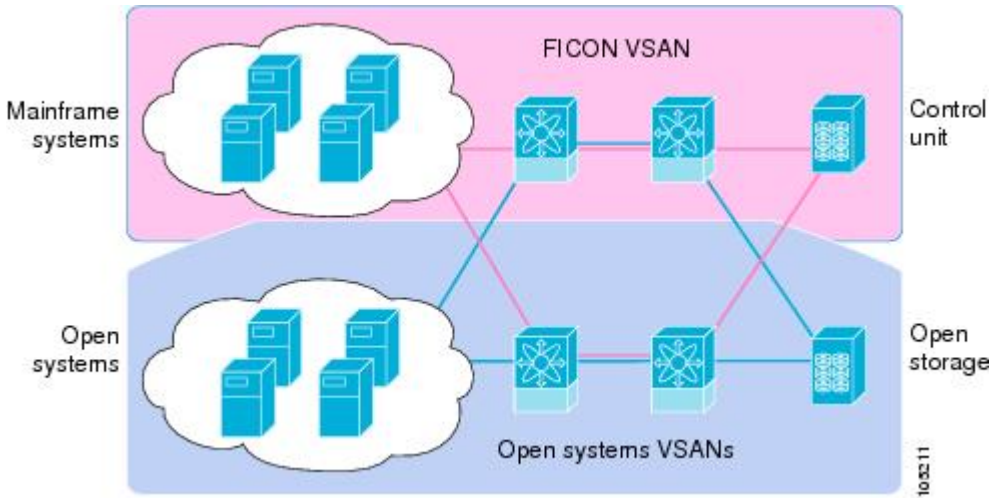
The FICON feature is supported only with the following platforms:

- Cisco MDS 9710 switches
- Cisco MDS 9706 switches
- Cisco MDS 9250i switches
- Cisco MDS 9220i switches
- Cisco MDS 9148V switches

FCP, NVMe, 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 (see the *Cisco MDS 9000 Series 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 23: Shared System Storage Network



This section includes the following topics:

FICON Requirements

The FICON feature has the following requirements:

- You can implement FICON features in the following switches:
  - Cisco MDS 9706 and MDS 9710 switches
  - Cisco MDS 9250i and MDS 9220i switches

Although in earlier releases the MAINFRAME_PKG license was required to configure FICON, beginning with NX-OS Release 9.4(1a), the FICON feature is a base feature of NX-OS and no special license is required.

MDS-Specific FICON Advantages

This section explains the additional FICON advantages in Cisco MDS switches and includes the following topics:

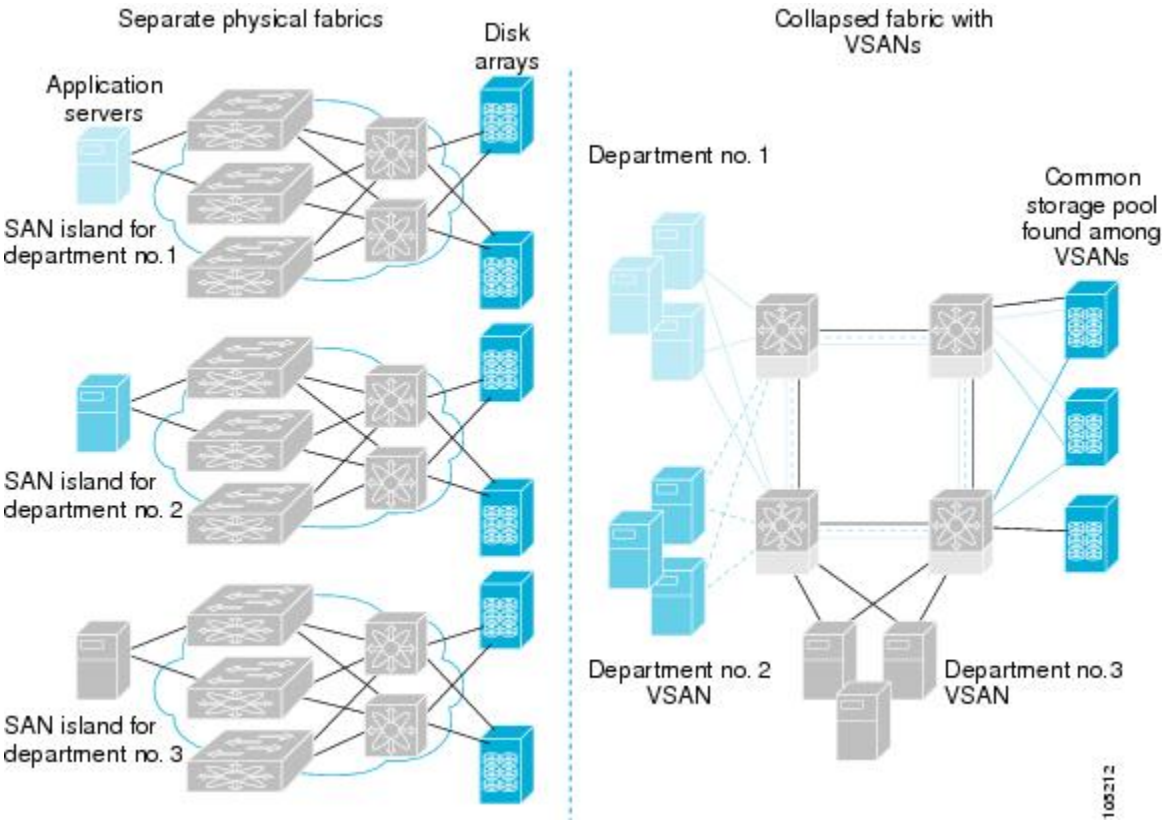
Fabric Optimization with VSANs

Generally, separate physical fabrics have a high level of switch management and have a higher implementation cost. The ports in each island may also be overprovisioned 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 overprovisioning 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 24: VSAN-Specific Fabric Optimization, on page 241).

Figure 24: 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 While you can configure VSANs in any Cisco MDS switch, you can only enable FICON in up to eight of these VSANs on switches that support the FICON feature.

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 an IBM Z Systems mainframe server. Each VSAN has its own set of fabric services (such as fabric server, name server, and zone server), FICON CUP, domain ID, Fabric Shortest Path First (FSPF) routing, operating mode, and security profile. FICON VSANs can span line cards and are dynamic in size. For example, one FICON VSAN with 8 ports can span 8 different line cards. FICON VSANs 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 VSAN is a

nondisruptive process. The maximum number of ports for a FICON VSAN is 254 per switch due to FICON addressing limitations.

FCIP Support

The multilayer architecture of the Cisco MDS 9000 Family enables a consistent feature set over protocol-agnostic switch fabric. Cisco MDS 9700 Series and 9200 Series switches transparently integrate FCP, NVMe, 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 can be extended over metro to global distances using ubiquitous IP infrastructure which simplifies business continuance strategies.

For more information, see the *Cisco MDS 9000 Series 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 Series 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 should 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. The default VSAN (VSAN 1) should never be used for production services.

Cisco MDS 9000-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 9700 Modular switches.

See the *Cisco MDS 9700 Series Multilayer Director Hardware Installation Guide*, the *Cisco MDS 9250i Multiservice Fabric Switch Hardware Installation Guide*, and the *Cisco MDS 9220i Fabric Switch Hardware Installation Guide.*

- High-availability FICON-enabled director — Cisco MDS 9700 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. The Cisco MDS 9710 supports up to 384 autosensing, 64/32/16/10/8/4/2-Gbps Fibre Channel ports for FCP, NVMe, and FICON connections as well as 1/10/25/40 Gbps IP Services ports for FCIP links. The Cisco MDS 9706 supports up to 192 autosensing, 64/32/16/10/8/4/2-Gbps Fibre Channel ports for FCP, NVMe, and FICON connections as well as 1/10/25/40 Gbps IP Services ports for FCIP links. See the *Cisco MDS 9000 Series High Availability Configuration Guide*.
- Infrastructure protection — Common software releases provide infrastructure protection across all Cisco MDS 9000 platforms. See the *Cisco MDS 9000 NX-OS Software Upgrade and Downgrade Guide*.
- VSAN technology — 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 Configuring and Managing VSANs, on page 9
- Port-level configurations — There are BB credits, beacon mode, and port security for each port. See the *Cisco MDS 9000 Series 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 the Configuring and Managing Zones, on page 39.
- Comprehensive security framework — 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. See the *Cisco MDS 9000 Series Security Configuration Guide* for information about RADIUS, TACACS+, FC-SP, and DHCHAP.
- Traffic encryption — IPSec is supported over FCIP. You can encrypt FICON, FCP, and NVMe traffic that is carried over FCIP. See the *Cisco MDS 9000 Series Comprehensive security framework 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, see 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 FICON in-band communications with the IBM Z Systems mainframe server. See the CUP In-Band Management, on page 278.
- Port address-based configurations — FICON port name attribute can be configured for ports in FICON VSANs. See the Configuring FICON Ports, on page 263.
- 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.

- Configuration files — Store and apply configuration files. See the FICON Configuration Files, on page 270.
- FICON and Open Systems Management Server features if installed. —See the VSANs for FICON and FCP Mixing, on page 242.
- Enhanced cascading support—See the CUP In-Band Management, on page 278.
- Date and time — Enable the IBM Z Systems Server to set the date and time for FICON VSANs on the switch. See the Allowing the Host to Control the Timestamp , on page 260.
- Configure SNMP trap recipients and community names — See the Configuring SNMP Control of FICON Parameters, on page 261.
- Call Home configurations — Configure the director name, location, description, and contact person. See the *Cisco MDS 9000 Series System Management Configuration Guide.*
- Configure preferred domain ID, FC ID persistence, and principal switch priority — For information about configuring domain parameters, see the *Cisco MDS 9000 Series System Management Configuration Guide*.
- Sophisticated SPAN diagnostics — 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, see the *Cisco MDS 9000 Series System Management Configuration Guide.*
- Configure R_A_TOV, E_D_TOV — See the “Fibre Channel Time-Out Values” section on page 11-1 .
- 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 Series System Management Configuration Guide*
- Port-level incident alerts—Display and clear port-level incident alerts. See the Clearing RLIR Information, on page 270.

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 each switch. See the *Cisco MDS 9000 Series Security Configuration Guide*).

The FICON topologies supported on the Cisco MDS 9000 Series switches are:

- **Single hop/traditional cascade** – This topology has two switches with a single hop (or set of ISLs between the switches. This support has been around since the introduction of FICON support in 2004.
- **Multi-hop cascade** – This topology allows for up to four (4) switches between the host channels and their associated control units. The ISLs between these switches can be fibre channel ISLs, port channels made of fibre channel ISLs, FCIP ISLs, or port channels made of FCIP ISLs. Multi-hop cascade was introduced in approximately 2017 and begins with the z13 System Z server forward.

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 or create a zoneset and associated zones for the VSAN. See Configuring zoning for FICON VSANs.
- Enable in-order delivery on the VSAN. See Configuring Fibre Channel Routing Services and Protocols, on page 167.
- Enable (and if required, configure) fabric binding on the VSAN. For more information about Fabric Binding, refer to the *Cisco MDS 9000 Series Security Configuration Guide*.
- Verify that conflicting FC IDs do not exist in the switch by configuring unique static domain IDs for each FICON VSAN on the switch or in the FICON fabric, if using enhanced FICON cascading. For information about configuring domain parameters, seethe *Cisco MDS 9000 Series System Management Configuration Guide*.
- Verify that the configured domain ID and requested domain ID match on the switch and these match what is configured for the switch in the HCD definitions on the IBM Z Systems Server. For information about configuring domain parameters, see the *Cisco MDS 9000 Series System Management Configuration Guide*.
- Add the CUP (area FE) to the zone, if you are using zoning. See the CUP In-Band Management, on page 278.

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
- User Assigned port numbers

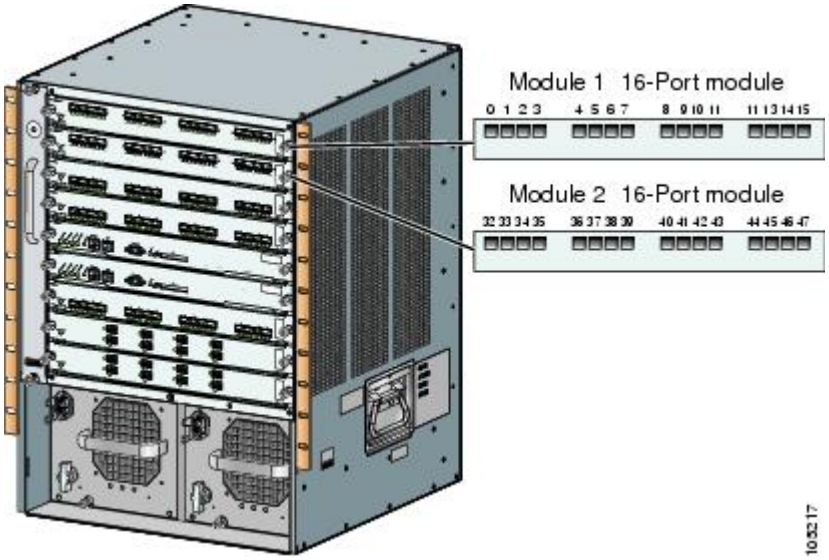
This section includes the following topics:

Note You must enable the FICON feature on the switch assigning FICON port numbers (see the About Enabling FICON on a VSAN, on page 252).

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 25: Default FICON Port Number in Numbering on the Cisco MDS 9000 Family Switch, on page 246).

Figure 25: Default FICON Port Number in Numbering on the Cisco MDS 9000 Family Switch



The default FICON port numbering is assigned based on the front panel location of the port and is specific to the slot in which the module resides. Forty-eight (48) port numbers are assigned to each slot on all Cisco MDS 9000 Family switches. 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 (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.

Note You can use the **ficon slot assign port-numbers** command to make use of excess ports mapped to a slot by default that are not addressable due to the module not having ports for them by manually assigning more port numbers to other slots. Before doing this, however, we recommend that you review the default port number assignments for Cisco MDS 9000 switches shown in Table 26: Default FICON Settings , on page 284 and Table 24: Default FICON Port Numbering in the Cisco MDS 9000 Family, on page 247, and that you read the following sections to gain a complete understanding of FICON port numbering: About the Reserved FICON Port Numbering Scheme, on page 248, FICON Port Numbering Guidelines, on page 249, and Assigning FICON Port Numbers to Slots, on page 249.

Note Only Fibre Channel, Port Channel, and FCIP ports are mapped to FICON port numbers. Other types of interfaces do not have a corresponding port number.

The following table lists the default port number assignment for the Cisco MDS 9000 Family of switches and directors.

Table 24: Default FICON Port Numbering in the Cisco MDS 9000 Family

Product	Slot Number	Port Number Allocation	To Port Channel/FCIP	Special (not assignable) Port Numbers
Cisco MDS 9250i Series	Slot 1	0 through 39	240 through 253	254 through 255
Cisco MDS 9220i Series	Slot 1	0 through 11	240 through 253	254 through 255
Cisco MDS 9710 Director	Slot 1	0 through 47	240 through 253	254 through 255
	Slot 2	48 through 95		
	Slot 3	96 through 143		
	Slot 4	144 through 191		
	Slot 5 - Supervisor	None		
	Slot 6 - Supervisor	None		
	Slot 7	192 through 239		
	Slot 8	None		
	Slot 9	None		
	Slot 10	None		
Cisco MDS 9706	Slot 1	0 through 47	240 through 253	254 through 255
	Slot 2	48 through 95		
	Slot 3 – Supervisor	None		
	Slot 4 – Supervisor	None		
	Slot 5	96 through 143		
	Slot 6	144 through 191		

Port Addresses

Following the deprecation of FICON Port Swap in NX-OS 9.4(1a), the port address is always the same as the port number.

Implemented and Unimplemented Port Addresses

An implemented port refers to any port address that is either assigned by default or is assignable using the **ficon slot x assign port-numbers** command. On the MDS 9000 Series switches, all port numbers from 0 to 253 are implemented giving customers flexibility for what values are used. The only unimplemented port is 255 while the value of 254 is always reserved for the CUP device.

About the Reserved FICON Port Numbering Scheme

A range of 254 port numbers are available for you to assign to all the ports on a switch. On the Cisco MDS 9710, you can have more than 254 physical ports on a switch and the excess ports do not have port numbers in the default numbering scheme. When you have more than 254 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 254 port numbers (0-253) and always has the FICON CUP device with port number 254 (0xFE).



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 9710 Director, ports 0 to 47 are considered uninstalled.
- The small form-factor pluggable (SFP) port is not present—For example, if a 48-port module is inserted in slot 2 in a Cisco MDS 9710 Director and no SFPs are inserted into the module, ports 48 to 95 are considered uninstalled.
- For slot 1, ports 0 to 47 have been assigned by default. 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 254 are considered implemented for any FICON-enabled VSAN. Therefore, VSAN 2 has port numbers 0 to 254 and one physical port, fc1/4. The corresponding physical ports 0 to 3, and 5 to 254 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 254).

Another scenario is if VSANs 2 through 5 are FICON-enabled, and trunking-enabled interface fc1/1 has VSANs 4 through 10, then port address 0 is uninstalled in VSAN 2 and 3.

- The port is part of a PortChannel—For example, if interface fc 1/1 is part of PortChanne1 5, port address 0 is uninstalled in all FICON VSANs. See Default Settings, on page 284.

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. As 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 from the pool of assigned logical port numbers.
- When a Fibre Channel ISL becomes a member of a PortChannel, the FICON Port number for this physical interface becomes uninstalled for this VSAN as it is now part of the FICON Port number assigned to the PortChannel.
- When a Fibre Channel ISL is removed from a PortChannel, the FICON port number for this physical interface becomes installed for this VSAN(s) that is it a part of.
- Each FCIP interface must be explicitly associated with a logical FICON port number.
- When an FCIP interface becomes a member of a PortChannel (which will have its own logical FICON port number), it also retains the logical FICON Port number associated with the FCIP interface itself.
- If logical port numbers are not assigned for PortChannels or FCIP interfaces, then these interfaces will not come up for the FICON VSAN(s).

See the About Port Numbers for FCIP and PortChannel, on page 250.

Assigning FICON Port Numbers to Slots

You can use the **show ficon port-number assign** and **show ficon first-available port-number** commands to determine which port numbers to use.

Caution When you assign, change, or release a port number, the port reloads.

To assign FICON port numbers to a slot, follow these steps:

Procedure

Step 1 switch# **config t**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **ficon slot 3 assign port-numbers 0-15, 48-63**

Reserves FICON port numbers 0 through 15 and 48 through 63 for up to 32 interfaces in slot 3.

If there are more interfaces in slot 3, they are not usable for FICON with this configuration.

Step 3 switch(config)# **ficon slot 3 assign port-numbers 0-15, 17-48**

Reserves FICON port numbers 0 through 15 for the first 16 interfaces and 17 through 48 for the next 32 interfaces in slot 3.

Step 4 switch(config)# **ficon slot 3 assign port-numbers 0-63**

Reserves FICON port numbers 0 through 63 for up to 64 interfaces in slot 3.

Step 5 switch(config)# **ficon slot 3 assign port-numbers 0-15, 56-63**

Changes the reserved FICON port numbers for up to 24 interfaces in slot 3.

Step 6 (Optional) switch(config)# **no ficon slot 3 assign port-numbers 0-15, 56-63**

Releases the FICON port numbers.

Displaying the FICON Port Number Assignments

Use the **show ficon port-numbers assign** command to display the port numbers assigned on the switch.

```
switch# show ficon port-numbers assign
ficon slot 1 assign port-numbers 0-47
ficon slot 2 assign port-numbers 48-95
ficon slot 3 assign port-numbers 96-143
ficon slot 4 assign port-numbers 144-191
ficon logical-port assign port-numbers 240-253
```

Use the **show ficon port-numbers assign slot** command to display the port numbers assigned to a specific slot.

```
switch# show ficon port-numbers assign slot 2
ficon slot 2 assign port-numbers 48-95
```

Use the **show ficon port-numbers assign** command to display the port numbers reserved for logical ports.

```
switch# show ficon port-numbers assign logical-port
ficon logical-port assign port-numbers 240-253
```

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, on page 263, and the Reserving FICON Port Numbers for FCIP and PortChannel Interfaces, on page 251, and the Binding Port Numbers to FCIP Interfaces , on page 264.

You can use the default port numbers if they are available (see Table 24: Default FICON Port Numbering in the Cisco MDS 9000 Family, on page 247) or if you reserve port numbers from the pool of port numbers that are not reserved for Fibre Channel interfaces (see the About the Reserved FICON Port Numbering Scheme, on page 248).

To find the first available port number to bind an FCIP or PortChannel interface, use the **show ficon first-available port-number** command (see Displays the Available Port Numbers, on page 280).

Tip The **show ficon vsan portaddress brief** command displays the port number to interface mapping. You can assign port numbers in the PortChannel/FCIP range that are not already assigned to a PortChannel or FCIP interface (see Displays Port Address Information in a Brief Format, on page 280).

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.

To reserve FICON port numbers for logical interfaces, follow these steps:

Procedure

Step 1 switch# **config t**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **ficon logical-port assign port-numbers 230-249**

Reserves port numbers 230 through 249 for FCIP and PortChannel interfaces.

Step 3 switch(config)# **ficon logical-port assign port-numbers 0xe6-0xf9**

Reserves port numbers 0xe6 through 0xf9 for FCIP and PortChannel interfaces.

Note You cannot change port numbers that are active. You must disable the interfaces using the **shutdown** command and unbind port numbers using the **no ficon portnumber** command. See the Configuring FICON Ports, on page 263.

Step 4 switch(config)# **no ficon logical-port assign port-numbers 230-249**

Releases the port numbers. Releasing the logical port numbers is particularly useful for switches that are not cascaded – thus allowing all 254 ports to be used for FICON channel and control unit connectivity.

Note You cannot release port numbers for interfaces that are active. You must disable the interfaces using the **shutdown** command and unbind port numbers using the **no ficon portnumber** command. See the Configuring FICON Ports, on page 263.

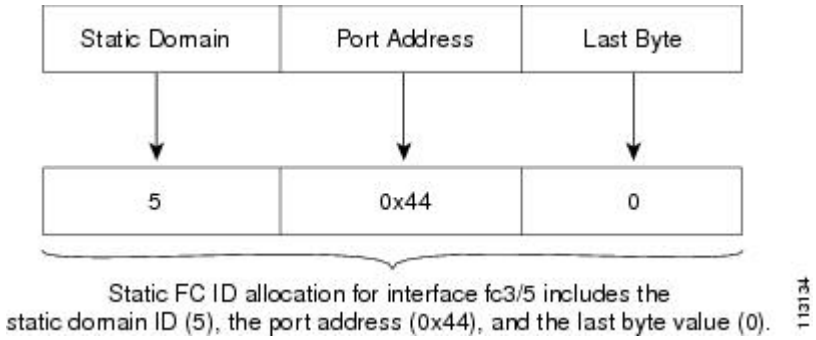
FC ID Allocation

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.

Note As the domain ID for FICON VSANs must be static, 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 default dynamic allocation scheme to use static FC IDs and vice versa (see Figure 26: Static FC ID Allocation for FICON, on page 252).

Figure 26: Static FC ID Allocation for FICON



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:

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:

- Use the automated **setup ficon** command.

  See the Setting Up a Basic FICON Configuration, on page 254.

- Manually address each prerequisite.

  See the About FICON, on page 239.

- Use Device Manager.

- 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). If all of the hosts and devices in this FICON VSAN support FICON Dynamic Routing (FDR), you can change the load balancing scheme to Source ID (SID) – Destination ID (DID) – Exchange ID (OXID).
  - The IPL configuration file is automatically created.

  See the About FICON Configuration Files, on page 271.



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 we recommend that only one instance of Device Manager monitor a FICON-enabled switch.

Enabling FICON on the Switch

By default, FICON is disabled in all switches in the Cisco MDS 9000 Family. You can enable FICON on the switch either explicitly or implicitly by enabling FICON on a VSAN. However, disabling FICON on all VSANs does not disable FICON on the switch. You must explicitly disable FICON.

To explicitly enable or disable FICON globally on the switch, follow these steps:

Procedure

Step 1 switch# **config t**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **feature ficon**

Enables FICON globally on the switch.

Step 3 switch(config)# **no feature ficon**

Disables FICON globally on the switch and removes all FICON configuration.

Setting Up a Basic FICON Configuration

This section steps you through the procedure to set up FICON on a specified VSAN in a Cisco MDS 9000 Family switch.

Note Press **Ctrl-C** at any prompt to skip the remaining configuration options and proceed with what is configured until that point.

Tip If you do not want to answer a previously configured question, or if you want to skip answers to any questions, press **Enter**. If a default answer is not available (for example, switch name), the switch uses what was previously configured and skips to the next question.

To enable and set up FICON, follow these steps:

Procedure

Step 1 Enter the **setup ficon** command at the EXEC command mode.

```
switch# setup ficon
                --- Ficon Configuration Dialog ---
This setup utility will guide you through basic Ficon Configuration
on the system.
Press Enter if you want to skip any dialog. Use ctrl-c at anytime
to skip all remaining dialogs.
```

Step 2 Enter **yes** (the default is **yes**) to enter the basic FICON configuration setup.

```
Would you like to enter the basic configuration dialog (yes/no) [yes]: yes
```

The FICON setup utility guides you through the basic configuration process. Press **Ctrl-C** at any prompt to end the configuration process.

Step 3 Enter the VSAN number for which FICON should be enabled.

```
Enter vsan [1-4093]:2
```

Step 4 Enter **yes** (the default is **yes**) to create a VSAN.

```
vsan 2 does not exist, create it? (yes/no) [yes]: yes
```

Step 5 Enter **yes** (the default is **yes**) to confirm your VSAN choice:

```
Enable ficon on this vsan? (yes/no) [yes]: yes
```

Note At this point, the software creates the VSAN if it does not already exist.

Step 6 Enter the domain ID number for the specified FICON VSAN.

```
Configure domain-id for this ficon vsan (1-239):2
```

Step 7 Enter **yes** (the default is **no**) to set up FICON in cascaded mode. If you enter **no**, skip to step 8 (see the CUP In-Band Management, on page 278).

```
Would you like to configure ficon in cascaded mode: (yes/no) [no]: yes
```

a) Assign the peer WWN for the attached FICON switch.

```
Configure peer wwn (hh:hh:hh:hh:hh:hh:hh:hh): 11:00:02:01:aa:bb:cc:00
```

b) Assign the peer domain ID for the attached FICON switch.

```
Configure peer domain (1-239) :4
```

c) Enter **yes** if you wish to configure additional peers (and repeat Steps 7a and 7b). Enter **no**, if you do wish to configure additional peers.

```
Would you like to configure additional peers: (yes/no) [no]: no
```

Step 8 Enter **yes** (the default is **yes**) to allow SNMP permission to modify existing port connectivity parameters (see the Configuring SNMP Control of FICON Parameters, on page 261).

```
Enable SNMP to modify port connectivity parameters? (yes/no) [yes]: yes
```

Step 9 Enter **no** (the default is **no**) to allow the host (mainframe) to modify the port connectivity parameters, if required (see the Allowing the Host to Change FICON Port Parameters, on page 259).

```
Disable Host from modifying port connectivity parameters? (yes/no) [no]: no
```

Step 10 Enter **yes** (the default is **yes**) to enable the **active equals saved** feature (see the Automatically Saving the Running Configuration, on page 262).

```
Enable active=saved? (yes/no) [yes]: yes
```

Step 11 Enter **yes** (the default is **yes**) if you wish to configure additional FICON VSANs.

```
Would you like to configure additional ficon vsans (yes/no) [yes]: no
```

Step 12 Review and edit the configuration that you have just entered.

Step 13 Enter no (the default is **no**) if you are satisfied with the configuration.

Note For documentation purposes, the following configurations shows three VSANs with different FICON settings. These settings provide a sample output for different FICON scenarios.

The following configuration will be applied:

```
fcdomain domain 2 static vsan 2
fcdomain restart disruptive vsan 2
fabric-binding database vsan 2
swwn 11:00:02:01:aa:bb:cc:00 domain 4
fabric-binding activate vsan 2
zone default-zone permit vsan 2
ficon vsan 2
no host port control
fcdomain domain 3 static vsan 3
fcdomain restart disruptive vsan 3
fabric-binding activate vsan 3 force
zone default-zone permit vsan 3
ficon vsan 3
no host port control
no active equals saved
vsan database
vsan 4
fcdomain domain 5 static vsan 4
fcdomain restart disruptive vsan 4
fabric-binding activate vsan 4 force
zone default-zone permit vsan 4
ficon vsan 4
no snmp port control
no active equals saved
Would you like to edit the configuration? (yes/no) [no]: no
```

Step 14 Enter yes (the default is **yes**) to use and save this configuration. The implemented commands are displayed. After FICON is enabled for the specified VSAN, you are returned to the EXEC mode switch prompt.

```
Use this configuration and apply it? (yes/no) [yes]: yes
`fcdomain domain 2 static vsan 2`
`fcdomain restart disruptive vsan 2`
`fabric-binding database vsan 2`
`swwn 11:00:02:01:aa:bb:cc:00 domain 4`
`fabric-binding activate vsan 2`
`zone default-zone permit vsan 2`
`ficon vsan 2`
`no host port control`
`fcdomain domain 3 static vsan 3`
`fcdomain restart disruptive vsan 3`
`fabric-binding activate vsan 3 force`
`zone default-zone permit vsan 3`
`ficon vsan 3`
`no host port control`
`no active equals saved`
```

Note If a new VSAN is created, two additional commands are displayed— **vsan database** and **vsan** *number*.

```
`vsan database`
`vsan 4`
`in-order-guarantee vsan 4`
`fcdomain domain 2 static vsan 4`
```

```
`fcdomain restart disruptive vsan 4`
`fabric-binding activate vsan 4 force`
`zone default-zone permit vsan 4`
`ficon vsan 4`
`no snmp port control`
Performing fast copy config...done. switch#
```

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, on page 262.

To manually enable FICON on a VSAN, follow these steps:

Procedure

Step 1 switch# **config t**

```
switch(config)#
```

Enters configuration mode.

Step 2 switch(config)# **vsan database**

```
switch(config-vsan-db)# vsan 5
switch(config-vsan-db)# show vsan usage
4 vsan configured
configured vsans:1-2,5,26
vsans available for configuration:3-4,6-25,27-4093
switch(config-vsan-db)# exit
```

Enables VSAN 5.

Step 3 switch(config)# **in-order-guarantee vsan 5**

Activates in-order delivery for VSAN 5.

See Configuring Fibre Channel Routing Services and Protocols, on page 167

Step 4 switch(config)# **fcdomain domain 20static vsan 5**

Configures the domain ID for VSAN 5.

For information about configuring domain parameters, refer to the *Cisco MDS 9000 Series System Management Configuration Guide*.

Step 5 switch(config)# **fabric-binding activate vsan 5 force**

Activates fabric binding on VSAN 5.

Refer to the *Cisco MDS 9000 Series Security Configuration Guide*.

Step 6 switch(config)# **zone default-zone permit vsan 5**

Sets the default zone to permit for VSAN 5.

Step 7 switch(config)# **ficon vsan 5**

switch(config-ficon)#

Enables FICON on VSAN 5.

Step 8 switch(config)# **no ficon vsan 6**

Disables the FICON feature on VSAN 6.

Step 9 switch(config-ficon)# **no host port control**

Prohibits mainframe users from moving the switch to an offline state.

See the Allowing the Host to Move the Switch Offline, on page 259.

Configuring the code-page Option

FICON strings are coded in Extended Binary-Coded Decimal Interchange Code (EBCDIC) format. Refer to the IBM System Z Server 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.

Tip This is an optional configuration. If you are not sure of the EBCDIC format to be used, we recommend retaining the **us-canada** (default) option.

To configure the **code-page** option in a VSAN, follow these steps:

Procedure

Step 1 switch# **config t**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **ficon vsan 2**

switch(config-ficon)#

Enables FICON on VSAN 2.

Step 3 switch(config-ficon)# **code-page italy**

Configures the **italy** EBCDIC format.

Step 4 switch(config-ficon)# **no code-page**

(Optional) Reverts to the factory default of using the **us-canada** EBCDIC format.

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.

To allow the host to move the switch to an offline state, follow these steps:

Procedure

Step 1 switch# **config terminal**
switch(config)#

Enters configuration mode.

Step 2 switch(config)# **ficon vsan 2**
switch(config-ficon)#

Enables FICON on VSAN 2.

Step 3 switch(config-ficon)# **no host control switch offline**

Prohibits mainframe users from moving the switch to an offline state.

Step 4 switch(config-ficon)# **host control switch offline**

Allows the host to move the switch to an offline state (default) and shuts down the ports.

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.

Use the **host port control** command to permit mainframe users to configure FICON parameters.

To allow the host (mainframe) to configure FICON parameters on the Cisco MDS switch, follow these steps:

Procedure

Step 1 switch# **config terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **ficon vsan 2**

switch(config-ficon)#

Enables FICON on VSAN 2.

Step 3 switch(config-ficon)# **no host port control**

Prohibits mainframe users from configuring FICON parameters on the Cisco MDS switch.

Step 4 switch(config-ficon)# **host port control**

Allows mainframe users to configure FICON parameters on the Cisco MDS switch (default).

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 IBM System Z Server 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.

The VSAN-clock current time is reported in the output of **show ficon vsan** *vsan-id*, **show ficon**, and **show accounting log** commands.

To configure host control of the timestamp, follow these steps:

Procedure

Step 1 switch# **config terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **ficon vsan 2**

switch(config-ficon)#

Enables FICON on VSAN 2.

Step 3 switch(config-ficon)# **no host set-timestamp**

Prohibits mainframe users from changing the VSAN-specific clock.

Step 4 switch(config-ficon)# **host set-timestamp**

Allows the host to set the clock on this switch (default).

Clearing the Time Stamp

Note You can clear time stamps only from the Cisco MDS switch—not the mainframe.

Use the **clear ficon vsan** *vsan-id* **timestamp** command in EXEC mode to clear the VSAN clock.

```
switch# clear ficon vsan 20 timestamp
```

Configuring SNMP Control of FICON Parameters

To configure SNMP control of FICON parameters, follow these steps:

Procedure

Step 1 switch# **config t**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **ficon vsan 2**

switch(config-ficon)#

Enables FICON on VSAN 2.

Step 3 switch(config-ficon)# **no snmp port control**

Prohibits SNMP users from configuring FICON parameters.

Step 4 switch(config-ficon)# **snmp port control**

Allows SNMP users to configure FICON parameters (default).

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.

Clearing FICON Device Allegiance

You can clear the current device allegiance by issuing the **clear ficon vsan** *vsan-id* **allegiance** command in EXEC mode.

```
switch# clear ficon vsan 2 allegiance
```

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 **active equals saved** option is automatically enabled on any FICON VSAN.

Table 25: Saving the Active FICON and Switch Configuration , on page 262 displays the results of the **Active = Saved** option **active equals saved** command and the implicit copy from the running configuration to the startup configuration (**copy running start**)**copy running-config startup-config** command in various scenarios.

When the Active=Saved option **active equals saved** command is enabled in any FICON-enabled VSAN in the fabric, then the following apply (see Number 1 and 2 in Table 25: Saving the Active FICON and Switch Configuration , on page 262):

- 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, on page 270).

If the Active=Saved option **active equals saved** command 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. Use the **copy running start** command explicitly (see number 3 in Table 25: Saving the Active FICON and Switch Configuration , on page 262).

Table 25: Saving the Active FICON and Switch Configuration

Number	FICON-enabled VSAN?	active equals saved Enabled?	Implicit copy 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.

Number	FICON-enabled VSAN?	active equals saved Enabled?	Implicit copy running start Issued?	Notes
4	No	Not applicable	—	—

Note If **active equals saved** is enabled, the Cisco NX-OS software ensures that you do not have to perform the **copy running startup** command for the FICON configuration as well. If your switch or fabric consists of multiple FICON-enabled VSANs, and one of these VSANs have **active equals saved** enabled, changes made to the non-FICON configuration results in all configurations being saved to the startup configuration.

To automatically save the running configuration, follow these steps:

Procedure

Step 1 switch# **config terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **ficon vsan 2**

switch(config-ficon)#

Enables FICON on VSAN 2.

Step 3 switch(config-ficon)# **active equals saved**

Enables the automatic save feature for all VSANs in the switch or fabric.

Step 4 switch(config-ficon)# **no active equals saved**

(Optional) Disables automatic save for this VSAN.

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 not installed, 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



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.

To bind a PortChannel with a FICON port number, follow these steps:

Procedure

Step 1 switch# **config terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **interface Port-channel 1**

switch(config-if)#

Enters the PortChannel interface configuration mode.

Step 3 switch(config-if)# **ficon portnumber 234**

Assigns the FICON port number to the selected PortChannel port.

Binding Port Numbers to FCIP Interfaces

You can bind (or associate) an FCIP interface with a FICON port number to bring up that interface.

To bind an FCIP interface with a FICON port number, follow these steps:

Procedure

Step 1 switch# **config t**

switch(config)#

Enters configuration mode.

Step 2 switch1(config)# **interface fcip 51**

switch1(config-if)#

Creates an FCIP interface (51).

Step 3 switch(config-if)# **ficon portnumber 208**

Assigns the FICON port number to the selected FCIP interface.

Configuring Port Blocking

FICON port blocking attribute has been deprecated from NX-OS 9.4(1a) forward. You can use the **shutdown**/ **'no' shutdown** commands to accomplish similar results. Port blocking was deprecated because the Z Systems Software that communicated with the switch over the FICON CUP interface to perform this function has long gone End of Support.

Port Prohibiting

FICON Port Prohibit attribute has been deprecated from NX-OS Release NX-OS 9.4(1a). If specific port to port protections are desired within FICON VSANs (above the protections implicitly given by the static routing nature of the System Z HCD configuration), zoning can be used.

Assigning a Port Address Name

To assign a port address name, follow these steps:

Procedure

Step 1 switch# **config t**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **ficon vsan 2**

switch(config-ficon)#

Enables FICON on VSAN 2.

Step 3 switch(config-ficon)# **portaddress 7**

switch(config-ficon-portaddr)#

Selects port address 7 for further configuration.

Step 4 switch(config-ficon-portaddr)# **name SampleName**

Assigns a name to the port address.

Note The port address name is restricted to 24 alphanumeric characters.

Step 5 switch(config-ficon-portaddr)# **no name SampleName**

Deletes a previously configured port address name.

About RLIR

The Registered Link Incident Report (RLIR) application provides a method for a switch port to send an 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 a 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.

The RLIR data is written to persistent storage when you enter the **copy running-config startup-config** command.

The RLIR data is written to persistent storage when you **copy** the running configuration to the startup configuration.

Specifying an RLIR Preferred Host

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.”

To specify the RLIR preferred host for a VSAN, follow these steps:

Procedure

Step 1 switch# **config terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **rlir preferred-cond fcid 0x772c00 vsan 5**

Specifies FC ID 0x772c00 as the RLIR preferred host in VSAN 5. (FC ID 0x772c00 is used here as an example.)

Step 3 (Optional) switch(config)# **no rlir preferred-cond fcid 0x772c00 vsan 5**

Removes FC ID 0x772c00 as the RLIR preferred host for VSAN 5.

Example

To display the RLIR preferred host configuration, use the **show rlir erl** command.

```
switch# show rlir erl
Established Registration List for VSAN: 5
-----------------------------------------------
FC-ID LIRR FORMAT REGISTERED FOR
-----------------------------------------------
0x772c00 0x18 conditional receive(*)
0x779600 0x18 conditional receive
0x779700 0x18 conditional receive
0x779800 0x18 conditional receive
Total number of entries = 4
(*) - Denotes the preferred host
```

Displaying RLIR Information

The **show rlir statistics** command displays the complete statistics of LIRR, RLIR, and DRLIR frames. It lists the number of frames received, sent, and rejected. Specify the VSAN ID to obtain VSAN statistics for a specific VSAN. If you do not specify the VSAN ID, then the statistics are shown for all active VSANs (see Examples Displays RLIR Statistics for All VSANs, on page 267 and Displays RLIR Statistics for a Specified VSAN, on page 268).

Displays RLIR Statistics for All VSANs

```
switch# show rlir statistics
Statistics for VSAN: 1
-----------------------
Number of LIRR received      = 0
Number of LIRR ACC sent      = 0
Number of LIRR RJT sent      = 0
Number of RLIR sent          = 0
Number of RLIR ACC received  = 0
Number of RLIR RJT received  = 0
Number of DRLIR received     = 0
Number of DRLIR ACC sent     = 0
Number of DRLIR RJT sent     = 0
Number of DRLIR sent         = 0
Number of DRLIR ACC received = 0
Number of DRLIR RJT received = 0
Statistics for VSAN: 100
-----------------------
Number of LIRR received      = 26
```

```
Number of LIRR ACC sent      = 26
Number of LIRR RJT sent      = 0
Number of RLIR sent          = 815
Number of RLIR ACC received  = 815
Number of RLIR RJT received  = 0
Number of DRLIR received     = 417
Number of DRLIR ACC sent     = 417
Number of DRLIR RJT sent     = 0
Number of DRLIR sent         = 914
Number of DRLIR ACC received = 828
Number of DRLIR RJT received = 0
```

Displays RLIR Statistics for a Specified VSAN

```
switch# show rlir statistics vsan 4
Statistics for VSAN: 4
------------------------
Number of LIRR received      = 0
Number of LIRR ACC sent      = 0
Number of LIRR RJT sent      = 0
Number of RLIR sent          = 0
Number of RLIR ACC received  = 0
Number of RLIR RJT received  = 0
Number of DRLIR received     = 0
Number of DRLIR ACC sent     = 0
Number of DRLIR RJT sent     = 0
Number of DRLIR sent         = 0
Number of DRLIR ACC received = 0
Number of DRLIR RJT received = 0
```

The **show rlir erl** command shows the list of Nx ports that are registered to receive the RLIRs with the switch. If the VSAN ID is not specified, the details are shown for all active VSANs (see Examples Displays All ERLs, on page 268 and Displays ERLs for the Specified VSAN, on page 269).

Displays All ERLs

```
switch# show rlir erl
Established Registration List for VSAN: 2
-----------------------------------------
FC-ID       LIRR FORMAT    REGISTERED FOR
-----------------------------------------
0x0b0200    0x18           always receive
Total number of entries = 1
Established Registration List for VSAN: 100
-----------------------------------------
FC-ID       LIRR FORMAT    REGISTERED FOR
-----------------------------------------
0x0b0500    0x18           conditional receive
0x0b0600    0x18           conditional receive
Total number of entries = 2
```

In Displays All ERLs, on page 268, if the Registered For column states that an FC ID is conditional receive, the source port is registered as a valid recipient of subsequent RLIRs. This source port is selected as an RLIR recipient only if no other ERL recipient is selected.

In Displays All ERLs, on page 268, if the Registered For column states that an FC ID is always receive, the source port is registered as a valid recipient of subsequent RLIRs. This source port is always selected as an LIR recipient.

Note If an always receive RLIR is not registered for any N port or if the delivery of an RLIR fails for one of those ports, then the RLIR is sent to a port registered to conditional receive RLIRs.

Displays ERLs for the Specified VSAN

```
switch# show rlir erl vsan 100
Established Registration List for VSAN: 100
-----------------------------------------------
FC-ID       LIRR FORMAT    REGISTERED FOR
-----------------------------------------------
0x0b0500    0x18           conditional receive
0x0b0600    0x18           conditional receive
Total number of entries = 2
```

Displays the LIR History

```
switch# show rlir history
Link incident history
-------------------------------------------------------------------------------------------------------------------

 Host Time Stamp         Switch Time Stamp       VSAN   Domain   Port    Intf        Link
Incident       Loc/Rem
-------------------------------------------------------------------------------------------------------------------

 Sep 20 12:42:44 2006    Sep 20 12:42:44 2006    ****    ****    0x0b    fc1/12      Loss
of sig/sync     LOC
 Reported Successfully to: [0x640001] [0x640201]
 Sep 20 12:42:48 2006    Sep 20 12:42:48 2006    ****    ****    0x0b    fc1/12      Loss
of sig/sync     LOC
 Reported Successfully to: [0x640001] [0x640201]
 *** ** **:**:** ****    Sep 20 12:42:51 2006    1001    230     0x12    ****        Loss
of sig/sync     REM
 Reported Successfully to: [0x640001] [0x640201]
 Sep 20 12:42:55 2006    Sep 20 12:42:55 2006    ****    ****    0x0b    fc1/12      Loss
of sig/sync     LOC
 Reported Successfully to: None [No Registrations]
 *** ** **:**:** ****    Sep 20 12:45:56 2006    1001    230     0x12    ****        Loss
of sig/sync     REM
 Reported Successfully to: None [No Registrations]
 *** ** **:**:** ****    Sep 20 12:45:56 2006    1001    230     0x12    ****        Loss
of sig/sync     REM
 Reported Successfully to: None [No Registrations]
 Sep 20 12:52:45 2006    Sep 20 12:52:45 2006    ****    ****    0x0b    fc1/12      Loss
of sig/sync     LOC
 Reported Successfully to: None [No Registrations]
**** - Info not required/unavailable
```

Displays Recent LIRs for a Specified Interface

```
switch# show rlir recent interface fc1/1-4
Recent link incident records
-------------------------------------------------------------------------------
Host Time Stamp          Switch Time Stamp        Port Intf   Link Incident
-------------------------------------------------------------------------------
Thu Dec 4 05:02:29 2003  Wed Dec 3 21:02:56 2003   2   fc1/2  Implicit Incident
Thu Dec 4 05:02:54 2003  Wed Dec 3 21:03:21 2003   4   fc1/4  Implicit Incident
```

Displays Recent LIRs for a Specified Port Number

```
switch# show rlir recent portnumber 1-4
Recent link incident records
-------------------------------------------------------------------------------
Host Time Stamp          Switch Time Stamp        Port Intf    Link Incident
-------------------------------------------------------------------------------
Thu Dec 4 05:02:29 2003  Wed Dec 3 21:02:56 2003   2   fc1/2   Implicit Incident
Thu Dec 4 05:02:54 2003  Wed Dec 3 21:03:21 2003   4   fc1/4   Implicit Incident
```

Clearing RLIR Information

Use the **clear rlir statistics** command to clear all existing statistics for a specified VSAN.

```
switch#
clear rlir statistics vsan 2
```

Use the **clear rlir history** command to clear the RLIR history where all link incident records are logged for all interfaces.

```
switch# clear rlir history
```

Use the **clear rlir recent interface** command to clear the most recent RLIR information for a specified interface.

```
switch# clear rlir recent interface fc 1/2
```

Use the **clear rlir recent portnumber** command to clear the most recent RLIR information for a specified port number.

```
switch# clear rlir recent portnumber 16
```

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 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 2 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:

- 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.

Refer to the *Cisco MDS 9000 Series Fundamentals Configuration Guide* for details on the normal configuration files used by Cisco MDS switches.

This section includes the following topics:

About FICON Configuration Files

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.

Applying the Saved Configuration Files to the Running Configuration

You can apply the configuration from the saved files to the running configuration using the **ficon vsan** *number* **apply file** *filename* command.

```
switch# ficon vsan 2 apply file SampleFile
```

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.

To edit the contents of a specified FICON configuration file, follow these steps:

Procedure

Step 1 switch# **config t**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **ficon vsan 2**

switch(config-ficon)#

Enables FICON on VSAN 2.

Step 3 switch(config-ficon)# **file IplFile1**

switch(config-ficon-file)#

Accesses the FICON configuration file called IplFile1 for VSAN 2. If this file does not exist, it is created.



Note All FICON file names are restricted to eight alphanumeric characters.

Step 4 switch(config-ficon)# **no file IplFileA**

(Optional) Deletes a previously created FICON configuration file.

Step 5 switch(config-ficon-file)# **portaddress 3**

switch(config-ficon-file-portaddr)#

Enters the submode for port address 3 to edit the contents of the configuration file named IplFile1.



Note The running configuration is not applied to the current configuration. The configuration is only applied when the **ficon vsan***number***aply file***filename* command is issued

Step 6 switch(config-ficon-file-portaddr)#**name P3**

Edits the content of the configuration file named IplFile1 by assigning the name P3 to port address 3. If the name did not exist, it is created. If it existed, it is overwritten.

Displaying FICON Configuration Files

Use the **show ficon vsan** *vsan-id* **file all** command to display the contents of all FICON configuration files.

```
switch# show ficon vsan 69 file all
File IPL      is locked
FICON configuration file IPL3     in vsan 69
Description:
  Port address 0(0)
   Port name is
   Port is not blocked
   Prohibited port addresses are 255(0xff)

  Port address 1(0x1)
   Port name is
   Port is not blocked
   Prohibited port addresses are 255(0xff)

  Port address 2(0x2)
   Port name is
   Port is not blocked
   Prohibited port addresses are 255(0xff)
```

Use the **show ficon vsan** *vsan-id* **file name** command to display the contents of a specific FICON configuration file.

```
switch# show ficon vsan 69 file name IPL3
FICON configuration file IPL3     in vsan 69
 Description:
 Port address 0(0)
  Port name is
  Port is not blocked
  Prohibited port addresses are 255(0xff)

 Port address 1(0x1)
  Port name is
  Port is not blocked
  Prohibited port addresses are 255(0xff)

 Port address 2(0x2)
  Port name is
  Port is not blocked
  Prohibited port addresses are 255(0xff)
```

Use the **show ficon vsan** *vsan-id* **filename** *filename* **portaddress** command to display the FICON configuration file information for a specific FICON port.

```
switch# show ficon vsan 69 file name IPL3 portaddress 2
FICON configuration file IPL3     in vsan 69
Description:
 Port address 2(0x2)
  Port name is
  Port is not blocked
  Prohibited port addresses are 255(0xff)
```

Copying FICON Configuration Files

Use the **ficon vsan** *vsan-id* **copy file** *existing-file-name save-as-file-name* command in EXEC mode to copy an existing FICON configuration file.

```
switch#
ficon vsan 69 copy file IPL IPL3
```

You can see the list of existing configuration files by issuing the **show ficon vsan** *vsan-id* command.

```
switch# show

ficon vsan 69
Ficon information for VSAN 69
  Ficon is online
  VSAN is active
  Host port control is Enabled
  Host offline control is Enabled
  User alert mode is Disabled
  SNMP port control is Enabled
  Host set director timestamp is Enabled
  Active=Saved is Enabled
  Loadbalance is srcid-dstid
  Number of implemented ports are 254
  Key Counter is 11
  FCID last byte is 0

  Date/Time is set by host to Fri Jan 26 21:28:56.122170 2024
  Device Allegiance not locked
  Codepage is us-canada
  Saved configuration files
    IPL
    IPL3
```

Port Swapping

The FICON Port Swapping ability has been deprecated from NX-OS Release 9.4(1a). If there is a need to move the FICON port address from one physical port to another, the **ficon slot x assign port-numbers** command can be used. This will move both the port number and port address.

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 Z System and native tape drives (both IBM and Oracle)
- The link between Z System and Virtual Tape Systems (both IBM and Oracle)
- The back-end link between the VSM (Oracle VSM) and tape drive (Oracle)

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

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, on page 276.

The following images show the supported configurations.

Note For information about FCIP tape acceleration, refer to the *Cisco MDS 9000 Series IP Services Configuration Guide*.

Configuring FICON Tape Acceleration

FICON tape acceleration has the following configuration considerations:

- 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.

To configure FICON tape acceleration, follow these steps:

Procedure

Step 1 switch# **config t**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **interface fcip 2**

switch(config-if)#

Specifies an FCIP interface and enters interface configuration submode.

Step 3 switch(config-if)# **ficon-tape-accelerator vsan 100**

```
This configuration change will disrupt all traffic on the FCIP interface in all
VSANs.  Do you wish to continue? [no] y
```

Enables FICON tape acceleration over an FCIP interface.

Step 4 switch(config-if)# **no ficon-tape-accelerator vsan 100**

```
This configuration change will disrupt all traffic on the FCIP interface in all
VSANs.  Do you wish to continue? [no] y
```

Disables (default) FICON tape acceleration over an FCIP interface.

What to do next

Use the **show running-config** command to verify the FICON tape acceleration over FCIP configuration.

```
switch# show running-config | begin "interface fcip"
interface fcip2
  ficon-tape-accelerator vsan 100
  no shutdown
...
```

Configuring FICON Tape Read Acceleration

All the configuration guidelines and restrictions applicable for FICON tape acceleration are also applicable for FICON tape read acceleration. FICON tape acceleration and FICON tape read acceleration can coexist.

To configure FICON tape read acceleration, follow these steps:

Procedure

Step 1 switch# **config t**
switch(config)#

Enters configuration mode.

Step 2 switch(config)# **interface fcip 2**
switch(config-if)#

Specifies an FCIP interface and enters interface configuration submode.

Step 3 switch(config-if)# **ficon-tape-read-accelerator**

```
This configuration change will disrupt all traffic on the FCIP interface in all
VSANs.  Do you wish to continue? [no]
```

Enables FICON tape read acceleration over an FCIP interface.

Step 4 switch(config-if)# **no ficon-tape-read-accelerator**

```
This configuration change will disrupt all traffic on the FCIP interface in all
VSANs.  Do you wish to continue? [no]
```

Disables (default) FICON tape read acceleration over an FCIP interface.

Configuring Zoning in a FICON VSAN

FICON Environments have an implicit method for controlling which FICON ports are allowed to talk to each other – the IOCDS on the System Z Server. This is why in many cases, actually coding zoning is not needed in FICON environments and by default FICON VSANs are configured to use a setting where all ports in the VSAN can talk each other without configuring zones. This setting is the configuration command **zonedefault-zone permit vsan xx.**

It is supported, however, to configure more traditional zoning for FICON VSANs. When this is done, it is disruptive and should likely be done at initial setup. If this is done later, it needs to be carefully planned so as not to create an outage of the environment.

First, the default behavior of zoning needs to be changed from permit. This is done by the following command.

```
switch# config terminal
switch(config)#no zonedefault-zone permit vsan 20
```

Next a Zoneset is defined within the FICON VSAN and then zones within this zoneset. The exact configuration of these zones is discussed elsewhere but below are a few recommendations / guidelines for how to configure zoning for FICON environments.

- Do not configure one large zone with all of the FICON ports in it. This is NOT the same as running with default zone permit.
- The most recommended zoning configuration for FICON is one that closely mirrors the IOCDS and that can be maintained in-sync with IOCDS changes as they occur. Creating zones for each CHPID or even small group of CHPIDs and then adding to this all of the Control Units that are mapped to these CHPIDs in the IOCDS. Once this is in place, any change that is made to the IOCDS can be mirrored to the zoning database in the associated FICON VSAN and everything then stays working and in-sync.
- Remember to configure the FICON VSAN CUP device in the VSAN with the CHPIDs that access it per the IOCDS. The WWN and FCID for the CUP(s) in the FICON VSAN by using the **show fcns database vsan xx**command.

```
switch# show fcns database vsan 69
VSAN 69:
--------------------------------------------------------------------------
FCID        TYPE  PWNN                    (VENDOR)        FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0x103000    N     c0:50:76:c7:6c:88:13:81 (IBM)           fcsb2-ch-cu fcsb2-cu-ch
0x103100    N     c0:50:76:c7:6c:88:13:91 (IBM)           fcsb2-ch-cu fcsb2-cu-ch
0x103200    N     c0:50:76:c7:6c:88:16:01 (IBM)           fcsb2-ch-cu fcsb2-cu-ch
0x103300    N     c0:50:76:c7:6c:88:16:11 (IBM)           fcsb2-ch-cu fcsb2-cu-ch
..

0x104200    N     c0:50:76:c7:6c:80:2e:c1 (IBM)           FICON:CU
0x104300    N     c0:50:76:c7:6c:80:2e:d1 (IBM)           FICON:CU
0x104400    N     c0:50:76:c7:6c:80:2f:01 (IBM)           FICON:CU
..

0x10aa00    N     c0:50:76:c9:10:80:15:41 (IBM)           FICON:CU
0x10ab00    N     c0:50:76:c9:10:80:15:51 (IBM)           FICON:CU
0x10fe00    N     21:01:00:2a:6a:a4:37:02 (Cisco)         FICON:CUP
```

Zoning can be done by WWN or by FCID (or any of the other supported methods) but usually WWN will be the best so that if CHPIDs need to move around, a change to the zoning database is not needed.

No zoning for the CUP device in the switch is needed for FICON VSANs that use default permit for zoning.

For more detailed information on the process of creating and maintaining zonesets and zones, refer to seethe *Cisco MDS 9000 Series Fabric Configuration Guide.*

Moving a FICON VSAN to an Offline State

Issue the **ficon vsan** *vsan-id* **offline** command in EXEC mode to log out all ports in the VSAN that need to be suspended.

Issue the EXEC-level **ficon vsan** *vsan-id* **online** command in EXEC mode to remove the offline condition and to allow ports to log on again.

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, on page 259).

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.

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 as well as send in asysnchronous information to the switch for things like port performance and device topology discovery.

Displaying Control Unit Information

Displays Control Unit Information, on page 278 displays configured control device information.

Displays Control Unit Information

```
switch# show ficon control-device sb3
Control Unit Image:0x80b9c2c
VSAN:20 CU:0x20fe00 CUI:0 CUD:0 CURLP:(nil)
ASYNC LP:(nil) MODE:1 STATE:1 CQ LEN:0 MAX:0
PRIMARY LP:   VSAN:0 CH:0x0 CHI:0 CU:0x0 CUI:0
ALTERNATE LP: VSAN:0 CH:0x0 CHI:0 CU:0x0 CUI:0
Logical Path:0x80b9fb4
VSAN:20 CH:0x200600 CHI:15 CU:0x20fe00 CUI:0 STATE:1 FLAGS:0x1
LINK: OH:0x0 OC:0x0 IH:0x0 IC:0x0
DEV:  OH:0x0 OC:0x0 IH:0x0 IC:0x0
SENSE: 00 00 00 00 00 00 00 46
       30 20 00 00 00 00 00 00
       00 00 00 00 00 00 00 00
```

```
       00 00 00 00 00 00 00 00
IUI:0x0 DHF:0x0 CCW:0x0 TOKEN:0x0 PCCW:0x0 FCCW:0x0 PTOKEN:0x0 FTOKEN:0x0
CMD:0x0 CCW_FLAGS:0x0 CCW_COUNT:0 CMD_FLAGS:0x0 PRIO:0x0 DATA_COUNT:0
STATUS:0x0 FLAGS:0x0 PARAM:0x0 QTP:0x0 DTP:0x0
CQ LEN:0 MAX:0 DESTATUS:0x0
```

Displaying FICON Information

This section includes the following topics:

Receiving FICON Alerts

In Displays Configured FICON Information, on page 279, the user alert mode is Enabled output confirms that you will receive an alert to indicate any changes in the FICON configuration.

Displays Configured FICON Information

```
switch# show ficon
  Ficon information for VSAN 69
  Ficon is online
  VSAN is active
  Host port control is Enabled
  Host offline control is Enabled
  User alert mode is Enabled
  SNMP port control is Enabled
  Host set director timestamp is Enabled
  Active=Saved is Disabled
  Number of implemented ports are 254
  Key Counter is 73723
  FCID last byte is 0(0)
  Serial number is 04.002A6A64AF85
  Date/Time is set by host to Fri Jan 26 21:38:10.991999 2024
  Device allegiance is locked by Host
  Codepage is us-canada
  Saved configuration files
    IPL
```

Displaying FICON Port Address Information

Examples Displays Port Address Information, on page 279 to Displays Port Address Counter Information, on page 280 display FICON Port Address information.

Displays Port Address Information

```
switch# show ficon vsan 69 portaddress
Port Address 4(0x4) is up in vsan 69
 Port number is 4(0x4), Interface is fc1/5
 Port name is
 Admin port mode is auto
 Port mode is F, FCID is 0xc00400
 Peer is type 008561 model T01 manufactured by IBM
 Serial num is 00000007AFB8, FICON tag is 0x80A8

Port Address 5(0x5) is up in vsan 69
```

```
 Port number is 5(0x5), Interface is fc1/6
 Port name is
 Admin port mode is auto
 Port mode is F, FCID is 0xc00500
 Peer is type 008561 model T01 manufactured by IBM
 Serial num is 00000007AFB8, FICON tag is 0x80A9

Port Address 6(0x6) is up in vsan 69
 Port number is 6(0x6), Interface is fc1/7
 Port name is
 Admin port mode is auto
 Port mode is F, FCID is 0xc00600
 Peer is type 008561 model T01 manufactured by IBM
 Serial num is 00000007AFB8, FICON tag is 0x80AA
```

Displays the Available Port Numbers

```
switch# show ficon first-available port-number
Port number 129(0x81) is available
```

In Displays Port Address Information in a Brief Format, on page 280, the interface column is populated with the corresponding interface if the port number is installed. If the port number is uninstalled, this space remains blank and indicates an unbound port number. For example, 56 is an unbound port number in Displays Port Address Information in a Brief Format, on page 280.

Displays Port Address Information in a Brief Format

```
switch# show ficon vsan 69 portaddress 50-55 brief
--------------------------------------------------------------------------------
Port    Port    Interface          Admin     Status           Oper   FCID
Address Number                     Blocked                    Mode
--------------------------------------------------------------------------------
0x32    0x32    fc2/3              off       up               F      0xc03200
0x33    0x33    fc2/4              off       up               F      0xc03300
0x34    0x34    fc2/5              off       up               F      0xc03400
0x35    0x35    fc2/6              off       up               F      0xc03500
0x36    0x36    fc2/7              off       up               F      0xc03600
0x37    0x37    fc2/8              off       up               F      0xc03700
```

Displays Port Address Counter Information, on page 280 displays the counters in FICON version format 1 (32-bit format)

Displays Port Address Counter Information

```
switch# show ficon vsan 69 portaddress 8 counters
Port Address 8(0x8) is up in vsan 20
    Port number is 8(0x8), Interface is fc1/8
    Version presented 1, Counter size 32b
    242811 frames input, 9912794 words
      484 class-2 frames, 242302 class-3 frames
      0 link control frames, 0 multicast frames
      0 disparity errors inside frames
      0 disparity errors outside frames
      0 frames too big, 0 frames too small
      0 crc errors, 0 eof errors
      0 invalid ordered sets
      0 frames discarded c3
```

```
  0 address id errors
116620 frames output, 10609188 words
  0 frame pacing time
  0 link failures
  0 loss of sync
  0 loss of signal
  0 primitive seq prot errors
  0 invalid transmission words
  1 lrr input, 0 ols input, 5 ols output
  0 error summary
```

Displaying the Configured FICON State

If FICON is enabled on a VSAN, you can display the port address information for that VSAN (see Displays the Specified Port Address When FICON Is Enabled, on page 281).

Displays the Specified Port Address When FICON Is Enabled

```
switch# show ficon vsan 69 portaddress 5
Port Address 5(0x5) is up in vsan 69
Port number is 5(0x5), Interface is fc1/6
 Port name is
 Admin port mode is auto
 Port mode is F, FCID is 0xc00500
 Peer is type 008561 model T01 manufactured by IBM
 Serial num is 00000007AFB8, FICON tag is 0x80A9
```

Displaying Buffer Information

In Displays the History Buffer for the Specified VSAN, on page 281, 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.

Displays the History Buffer for the Specified VSAN

```
switch# show ficon vsan 20 director-history
Director History Buffer for vsan 20
-------------------------------------------
Key Counter           Ports Address
                      Changed
-------------------------------------------
74556                 43
74557                 44
74558                 45
74559                 46
74560                 47
74561                 48
```

```
74562                   49
74563                   50
74564                   51
74565                   52
74566                   53
74567                   54
74568                   55
74569                   56
74570                   57
74571                   58
74572                   59
74573                   60
74574                   61
74575                   62
74576                   63
74577                   64
74578
74579
74580                   1-3,5,10,12,14-16,34-40,43-45,47-54,56-57,59-64
74581                   3,5
74582                   64
74583
74584                   1-3,10,12,14-16,34-40,43-45,47-54,56-57,59-64
74585                   1
74586                   2
74587                   3
```

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.

Displaying FICON Information in the Running Configuration

Displays the Running Configuration Information, on page 282 displays the FICON-related information in the running configuration.

Displays the Running Configuration Information

```
switch# show running-config
Building Configuration ...
in-order-guarantee
vsan database
  vsan 11 name "FICON11" loadbalancing src-dst-id
  vsan 75 name "FICON75" loadbalancing src-dst-id
fcdomain domain 11 static vsan 11
fcdomain domain 119 static vsan 75
fcdroplatency network 100 vsan 11
fcdroplatency network 500 vsan 75
feature fabric-binding
```

```
fabric-binding database vsan 11
  swwn 20:00:00:0d:ec:01:20:c0 domain 10
fabric-binding database vsan 75
  swwn 20:00:00:0d:ec:00:d6:40 domain 117
fabric-binding activate vsan 11
fabric-binding activate vsan 75
ficon vsan 75
interface port-channel 1
  ficon portnumber 0x80
  switchport mode E
snmp-server user mblair network-admin auth md5 0x688fa3a2e51ba5538211606e59ac292
7 priv 0x688fa3a2e51ba5538211606e59ac2927 localizedkey
snmp-server user wwilson network-admin auth md5 0x688fa3a2e51ba5538211606e59ac29
27 priv 0x688fa3a2e51ba5538211606e59ac2927 localizedkey
snmp-server host 171.71.187.101 traps version 2c public udp-port 1163
snmp-server host 172.18.2.247 traps version 2c public udp-port 2162
vsan database
  vsan 75 interface fc1/1
...
interface mgmt0
  ip address 172.18.47.39 255.255.255.128
  switchport speed 100
  switchport duplex full
no system health
ficon vsan 75
  file IPL
```

Displaying FICON Information in the Startup Configuration

Displays the Startup Configuration, on page 283 displays the FICON-related information in the startup configuration.

Displays the Startup Configuration

```
switch# show startup-config
...
ficon vsan 2
file IPL
```

Displays the Startup Configuration Status, on page 283 displays the switch response to an implicitly-issued copy running start command. In this case, only a binary configuration is saved until you explicitly issue the **copy running start** command again (seeTable 25: Saving the Active FICON and Switch Configuration , on page 262)

Displays the Startup Configuration Status

```
switch# show startup-config
No ASCII config available since configuration was last saved internally
on account of 'active=saved' mode.
Please perform an explicit 'copy running startup` to get ASCII configuration
```

Displaying FICON-Related Log Information

Displays Logging Levels for the FICON Feature, on page 284 and Displays FICON-Related Log File Contents, on page 284 display the logging information for FICON-related configurations.

Displays Logging Levels for the FICON Feature

```
switch# show logging level ficon
Facility        Default Severity        Current Session Severity
--------        ----------------        ------------------------
ficon                   2                       2
0(emergencies)          1(alerts)       2(critical)
3(errors)               4(warnings)     5(notifications)
6(information)          7(debugging)
```

Displays FICON-Related Log File Contents

```
switch# show logging logfile
...
2024 Jan 29 12:45:40.784113 Challenger-9710 %PORT-5-IF_UP: %$VSAN 69: 2024 Mon Jan 29
16:45:40.122353%$ Interface fc2/17 is up in mode F
2024 Jan 29 12:45:40.857400 Challenger-9710 %PORT-5-IF_UP: %$VSAN 69: 2024 Mon Jan 29
16:45:40.122371%$ Interface fc2/14 is up in mode F
2024 Jan 29 12:45:40.866201 Challenger-9710 %PORT-5-IF_UP: %$VSAN 69: 2024 Mon Jan 29
16:45:40.122373%$ Interface fc1/6 is up in mode F
2024 Jan 29 12:45:40.882935 Challenger-9710 %PORT-5-IF_UP: %$VSAN 69: 2024 Mon Jan 29
16:45:40.122378%$ Interface fc1/22 is up in mode F
2024 Jan 29 12:45:40.942220 Challenger-9710 %PORT-5-IF_UP: %$VSAN 69: 2024 Mon Jan 29
16:45:40.122392%$ Interface fc2/16 is up in mode F
2024 Jan 29 12:45:40.943643 Challenger-9710 %PORT-5-IF_UP: %$VSAN 69: 2024 Mon Jan 29
16:45:40.122392%$ Interface fc2/19 is up in mode F
```

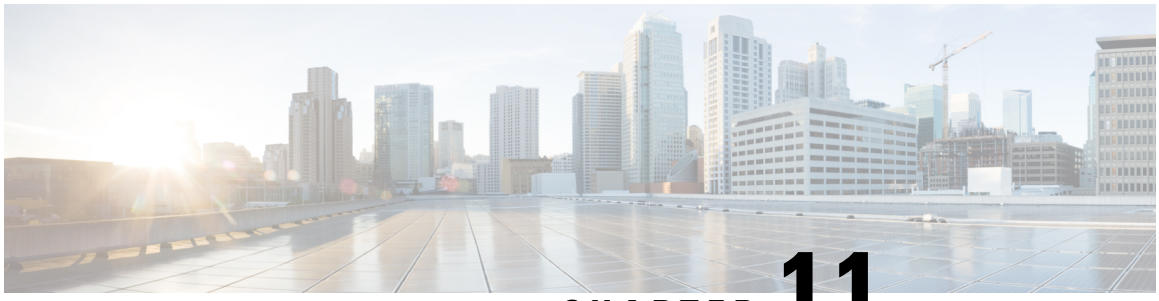
Default Settings

Table 26: Default FICON Settings , on page 284 lists the default settings for FICON features.

Table 26: Default FICON Settings

Parameters	Default
FICON feature	Disabled.
Port numbers	Always the 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.

Parameters	Default
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.



CHAPTER 11

Advanced Features and Concepts

This chapter describes the advanced features provided in switches in the Cisco MDS 9000 Family. It includes the following sections:

- Common Information Model, on page 287
- Fibre Channel Time-Out Values, on page 287
- Organizationally Unique Identifiers, on page 292
- World Wide Names, on page 293
- FC ID Allocation for HBAs, on page 295
- Default Settings, on page 298

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.

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/

Note The CIM Functionality and SMI-S is now supported with Cisco Prime Data Center Network Manager (DCNM). Please refer to “Cisco Prime DCNM Installation Guide” and “SMI-S and Web Services Programming Guide, Cisco DCNM for SAN.

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 4,000 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.

This section includes the following topics:

Timer Configuration 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 Fibre Channel timers across all VSANs, follow these steps:

Procedure

Step 1 switch# **config terminal**

switch(config)

Enters configuration mode.

Step 2 switch(config)# **fctimer R_A_TOV 6000**

Configures the R_A_TOV value for all VSANs to be 6000 msec. This type of configuration is not permitted unless all VSANs are suspended.

Timer Configuration 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*.

To configure per-VSAN Fiber Channel timers, follow these steps:

Procedure

Step 1 switch# **config terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **fctimer D_S_TOV 6000 vsan 2**

```
Warning: The vsan will be temporarily suspended when updating the timer value This configuration
would impact whole fabric. Do you want to continue? (y/n) y

Since this configuration is not propagated to other switches, please configure the same value in all
 the switches
```

Configures the D_S_TOV value to be 6000 msec for VSAN 2. Suspends the VSAN temporarily. You have the option to end this command, if required.

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.

Enabling fctimer Distribution

To enable or disable fctimer fabric distribution, follow these steps:

Procedure

Step 1 switch# **config terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **fctimer distribute**

Enables fctimer configuration distribution to all switches in the fabric. Acquires a fabric lock and stores all future configuration changes in the pending database.

Step 3 switch(config)# **no fctimer distribute**

Disables (default) fctimer configuration distribution to all switches in the fabric.

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.

To commit the fctimer configuration changes, follow these steps:

Procedure

Step 1 switch# **config terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **fctimer commit**

Distributes the fctimer configuration changes to all switches in the fabric and releases the lock. Overwrites the effective database with the changes made to the pending database.

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.

To discard the fctimer configuration changes, follow these steps:

Procedure

Step 1 switch# **config terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **fctimer abort**

Discards the fctimer configuration changes in the pending database and releases the fabric lock.

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.

To use administrative privileges and release a locked fctimer session, use the **clear fctimer session** command.

```
switch# clear fctimer session
```

Database Merge Guidelines

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

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.

Note The number of pending fctimer configuration operations cannot be more than 15. At that point, you must commit or terminate the pending configurations before performing any more operations.

Displaying Configured fctimer Values

Use the **show fctimer** command to display the configured fctimer values (see the following examples).

Displays Configured Global TOVs

```
switch# show fctimer

F_S_TOV   D_S_TOV   E_D_TOV   R_A_TOV
----------------------------------------
5000 ms   5000 ms   2000 ms   10000 ms
```

Note The F_S_TOV constant, though not configured, is displayed in the output of the **show fctimer** command.

Displays Configured TOVs for a Specified VSAN

```
switch# show fctimer vsan 10

vsan no.  F_S_TOV   D_S_TOV   E_D_TOV   R_A_TOV
-------------------------------------------------
10        5000 ms   5000 ms   3000 ms   10000 ms
```

Organizationally Unique Identifiers

Organizationally Unique Identifiers (OUIs) are unique 24 bit numbers that identify an organization globally. OUIs are extended by the organisation they are assigned to, to create 48 bit or 60 bit Extended Unique Identifiers (EUIs). Cisco obtains OUIs from IEEE and uses them to construct EUIs. These are assigned and burnt in to each system. A system may have one or more EUIs assigned to it. The EUIs are used in various forms such as MAC addresses, WWNs, SNMP identifiers, and so on.

Cisco MDS NX-OS software has an OUI database based on which certain software functionalities are made available. If a new Cisco device with an unrecognized OUI is added to a fabric, there is a possibility that some of these functionalities might be affected. To avoid this issue, the ability to manually add OUIs to the OUI database using the CLI is available.

Guidelines and Limitations

- ISSU—After an upgrade, there may be instances of duplicate OUIs in the default (built-in) and static (user defined) lists. In such a scenario, we recommend that you compare static OUIs with those in the default list and delete the duplicate static OUIs.

- ISSD—Delete all the configured or static OUIs before performing a downgrade to a release that does not support the **wwn oui** *oui-id* command.

For more information on deleting OUIs, see the Adding and Deleting OUIs, on page 293 section.

Adding and Deleting OUIs

To add an OUI to the OUI database, enter the **wwn oui** *oui-id* command in global configuration mode. To delete an OUI from the OUI database, enter the **no wwn oui** *oui-id* command in global configuration mode.

For detailed information about the **wwn oui** command, see the *Cisco MDS 9000 Family Command Reference* .

Configuration Examples for Adding and Deleting OUIs

Example: Adding and Deleting OUIs

```
switch# configure terminal
switch(config)# wwn oui 0x10001c
switch(config)# no wwn oui 0x10001c
switch(config)# end
```

Example: Displaying OUIs

```
switch# show wwn oui
OUI        Vendor               Default/Static
-----------------------------------------------------
0x0000fc   Cisco                Static
0x00000c   Cisco                Default
0x000196   Cisco                Default
0x000197   Cisco                Default
0x0001c7   Cisco                Default
0x0001c9   Cisco                Default
```

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 27: Standardized NAA WWN Formats , on page 293).

Table 27: Standardized NAA WWN Formats

NAA Address	NAA Type	WWW 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

NAA Address	NAA Type	WWW Format	
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.

This section includes the following topics:

Displaying WWN Information

Use the **show wwn** commands to display the status of the WWN configuration. See the following examples:

Displays the Status of All WWNs

```
switch# show wwn status
         Type 1 WWNs: Configured:     64 Available:     48 (75%) Resvd.: 16
    Types 2 & 5 WWNs: Configured: 524288 Available: 450560 (85%) Resvd.: 73728
NKAU & NKCR WWN Blks: Configured:   1760 Available:   1760 (100%)
        Alarm Status:      Type1:   NONE Types 2&5:   NONE
```

Displays Specified Block ID Information

```
switch# show wwn status block-id 51

WWNs in this block: 21:00:ac:16:5e:52:00:03 to 21:ff:ac:16:5e:52:00:03
Num. of WWNs:: Configured: 256 Allocated:    0 Available: 256
Block Allocation Status: FREE
```

Displays the WWN for a Specific Switch

```
switch# show wwn switch
Switch WWN is 20:00:ac:16:5e:52:00:00
```

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.

Configuring a Secondary MAC Address

To allocate secondary MAC addresses, follow these steps:

Procedure

Step 1 switch# **config terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **wwn secondary-mac 00:99:55:77:55:55 range 64**

```
This command CANNOT be undone.

Please enter the BASE MAC ADDRESS again: 00:99:55:77:55:55

Please enter the mac address RANGE again: 64

From now on WWN allocation would be based on new MACs.

Are you sure? (yes/no) no

You entered: no. Secondary MAC NOT programmed
```

Configures the secondary MAC address. This command cannot be undone.

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, on page 295).

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 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.

This section includes the following topics:

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.



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.

Use the **fcinterop FCID allocation auto** command to change the FC ID allocation and the **show running-config** command to view the currently allocated mode.

- When you issue a **write erase**, the list inherits the default list of company IDs shipped with a relevant release.

To allocate company IDs, follow these steps:

Procedure

Step 1 switch# **config terminal**
switch(config)#

Enters configuration mode.

Step 2 switch(config)# **fcid-allocation area company-id 0x003223**

Adds a new company ID to the default list.

Step 3 switch(config)# **no fcid-allocation area company-id 0x00E069**

Deletes a company ID from the default list.

Step 4 switch(config)# **fcid-allocation area company-id 0x003223**

Adds a new company ID to the default list.

Verifying the Company ID Configuration

You can view the configured company IDs by issuing the **show fcid-allocation area** command (see Displays the List of Default and Configured Company IDs, on page 297). Default entries are listed first and the user-added entries are listed next. Entries are listed even if they were part of the default list and you later removed them.

Displays the List of Default and Configured Company IDs

```
switch# show fcid-allocation area
FCID area allocation company id info:
 00:50:2E <-------------------- Default entry
 00:50:8B
 00:60:B0
 00:A0:B8
 00:E0:69
 00:30:AE + <------------------ User-added entry
 00:32:23 +
 00:E0:8B * <------------- Explicitly deleted entry (from the original default list)
Total company ids: 7
+ - Additional user configured company ids.
* - Explicitly deleted company ids from default list.
```

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.

You can also view or obtain the company IDs in a specific WWN by issuing the **show fcid-allocation company-id-from-wwn** command (see Displays the Company ID for the Specified WWN, on page 297). Some WWN formats do not support company IDs. In these cases, you many need to configure the FC ID persistent entry.

Displays the Company ID for the Specified WWN

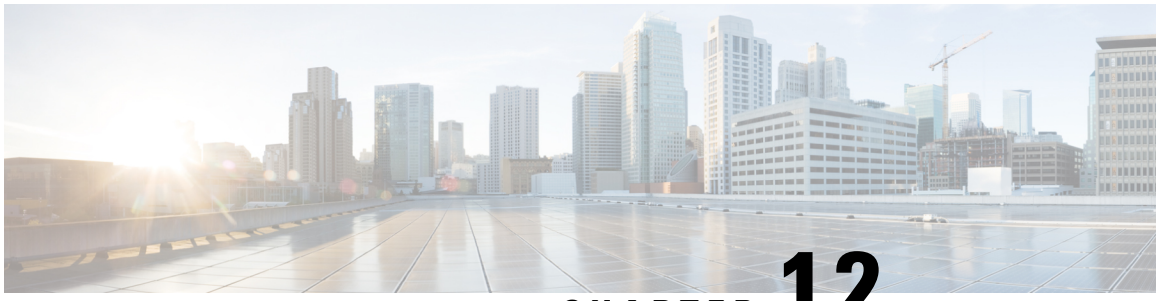
```
switch# show fcid-allocation company-id-from-wwn 20:00:00:05:30:00:21:60
Extracted Company ID: 0x000530
```

Default Settings

Table 28: Default Settings for Advanced Features , on page 298 lists the default settings for the features included in this chapter.

Table 28: 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



CHAPTER 12

Configuring Fibre Channel Common Transport Management Security

This chapter describes the Fibre Channel Common Transport (FC-CT) Management Security feature for Cisco MDS 9000 Series switches.

- About Fibre Channel Common Transport , on page 299
- Configuration Guidelines, on page 299
- Configuring the Fibre Channel Common Transport Query, on page 300
- Verifying Fibre Channel Common Transport Management Security, on page 300
- Default Settings, on page 301

About Fibre Channel Common Transport

With the FC-CT management security feature, you can configure the network in such a manner that only a storage administrator or a network administrator can send queries to a switch and access information such as devices that are logged in devices in the fabric, switches in the fabric, how they are connected, how many ports each switch has and where each port is connected, configured zone information and privilege to add or delete zone and zone sets, and host bus adapter (HBA) details of all the hosts connected in the fabric.

Note In Cisco MDS NX-OS Release 6.2(9), the FC management feature is disabled by default. To enable FC management feature, use the fc-management enable command.

You can configure which pWWNs can send FC-CT management query and modify request to the management server. When any of the modules, such as a zone server, unzoned Fibre Channel name server (FCNS), or Fabric Configuration Server (FCS) receives an FC-CT management query, they perform a read operation on the FC-management database. If device is found in FC-management database, a reply is sent according to the permissions granted. If the device is not found in the FC-management database, each module sends a reject. If FC-management is disabled, each module processes each management query.

Configuration Guidelines

The FC-management security feature has the following configuration guidelines:

- When the FC-management security feature is enabled on a Cisco MDS switch, all management queries to the server are rejected unless the port world-wide name (pWWN) of the device that is sending management queries is added to FC-management database.
- When you enable FC Management, FC-CT management server queries from N_Port Virtualization (NPV) switches to N_Port Identifier Virtualization (NPIV) switches are rejected. We recommend that you add the switch world-wide name (sWWN) of the NPV switch to the FC management database of the NPIV switch after enabling the FC-management security feature.

Configuring the Fibre Channel Common Transport Query

To configure the FC-CT management security, follow these steps:

Procedure

Step 1 switch# **config terminal**

Enters configuration mode.

Step 2 switch(config)# fc-management enable

Enables the FC-CT management security.

Step 3 switch(config)# **fc-management database vsan 1**

Configures the FC-CT management Security database.

Step 4 switch(config-fc-mgmt)# pwwn 1:1:1:1:1:1:1:1 feature all operation both

Adds the pWWN to the FC management database. You also can use these optional keywords when configuring the pwwn command:

- fcs— Enables or disables FC-CT query for fabric conf-server.
- fdmi—Enables or disables FC-CT query for FDMI.
- unzoned-ns—Enables or disables FC-CT query for unzoned name-server.
- zone—Enables or disables FC-CT query for zone-server.

Step 5 switch# **show fc-managment database**

Displays the configured FC-CT management information.

Verifying Fibre Channel Common Transport Management Security

The **show fc-management database** command displays the configured FC-CT management security feature information, see example Displays the Contents of the Fibre Channel Common Transport Query, on page 301.

Displays the Contents of the Fibre Channel Common Transport Query

```
switch# show fc-management database
------------------------------------------------------------
VSAN PWWN FC-CT Permissions per FC services
------------------------------------------------------------
1 01:01:01:01:01:01:01:01 Zone(RW), Unzoned-NS(RW), FCS(RW), FDMI(RW)
1 02:02:02:02:02:02:02:02 Zone(R), Unzoned-NS(R), FCS(R), FDMI(R)
1 03:03:03:03:03:03:03:03 Zone(W), Unzoned-NS(W), FCS(W), FDMI(W)
------------------------------------------------------------
Total 3 entries
switch#
```

To verify the if the FC-management security feature is enabled or not, use the show fc-management status command:

```
switch# show fc-management status
Mgmt Security Disabled
switch#
```

Default Settings

Table 29: Default FC Management Settings , on page 301 lists the default settings for the FC management security feature in a Cisco MDS 9000 Family switch.

Table 29: Default FC Management Settings

Parameters	Default
FC-management	Disabled