Cisco Nexus 9000 Series NX-OS SAN Switching Configuration Guide, Release 10.1(x)

First Published: 2021-02-16

Last Modified: 2024-01-12

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 REFERENCED IN THIS DOCUMENTATION ARE SUBJECT TO CHANGE WITHOUT NOTICE. EXCEPT AS MAY OTHERWISE BE AGREED BY CISCO IN WRITING, ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS DOCUMENTATION ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED.

The Cisco End User License Agreement and any supplemental license terms govern your use of any Cisco software, including this product documentation, and are located at: http://www.cisco.com/go/softwareterms.Cisco product warranty information is available at http://www.cisco.com/go/warranty. US Federal Communications Commission Notices are found here http://www.cisco.com/c/en/us/products/us-fcc-notice.html.

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 products and features described herein as in development or available at a future date remain in varying stages of development and will be offered on a when-and if-available basis. Any such product or feature roadmaps are subject to change at the sole discretion of Cisco and Cisco will have no liability for delay in the delivery or failure to deliver any products or feature roadmap items that may be set forth in this document.

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.

The documentation set for this product strives to use bias-free language. For the purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on RFP documentation, or language that is used by a referenced third-party product.

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: www.cisco.com go trademarks. 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)

© 2021–2024 Cisco Systems, Inc. All rights reserved.

CONTENTS

Trademarks ?

PREFACE

Preface xiii

Audience xiii

Document Conventions xiii

Related Documentation for Cisco Nexus 9000 Series Switches xiv

Documentation Feedback xiv

Communications, Services, and Additional Information xiv

CHAPTER 1

New and Changed Information 1

New and Changed Information 1

CHAPTER 2

Hardware Support for SAN Switching 3

Hardware Support for SAN Switching 3

Supported Platforms 3

CHAPTER 3

Overview 5

Licensing Requirements 5

SAN Switching Overview 5

SAN Switching General Guidelines and Limitations 7

CHAPTER 4

Enabling FC/FCoE Switch Mode 9

Guidelines and Limitations for FC Switch Mode 9

Enabling FC/FCoE 9

Disabling FC/FCoE 10

Disabling LAN Traffic on an FCoE Link 11

Configuring the FC-Map 12
Configuring the Fabric Priority 12
Setting the Advertisement Interval 13

CHAPTER 5 **Configuring FCoE 15**
FCoE Topologies 15
Directly Connected CNA Topology 15
Remotely Connected CNA Topology 17
FCoE Best Practices 18
Directly Connected CNA Best Practice 18
Remotely Connected CNA Best Practice 19
Guidelines and Limitations 20
Configuring FC/FCoE 21
Perform TCAM Carving 21
Configuring LLDP 22
Configuring Default QoS 22
Configuring User Defined QoS 22
Configuring Traffic Shaping 24
FCoE with vPC Configuration Examples 25
Cisco Nexus 9000 Series Switch vPC Configuration Example 26
Cisco Nexus 9000 Series Switch FCoE Configuration Example 30

CHAPTER 6 **Configuring Fibre Channel Interfaces 37**
Information About Fibre Channel Interfaces 37
Virtual Fibre Channel Interfaces 37
VF Port 38
Interface Modes 38
E Port 38
F Port 38
TE Port 38
TF Port 39
Auto Mode 39
Interface States 39
Administrative States 39

Operational States 40
Reason Codes 40
Buffer-to-Buffer Credits 42
Licensing Requirements for Fibre Channel 43
Enabling the Fibre Channel Port License 43
Configuring QoS for no-drop Support 44
Physical Fibre Channel Interfaces 47
Long-Distance ISLs 47
Configuring Fibre Channel Interfaces 48
Configuring a Fibre Channel Interface 48
Configuring a Range of Fibre Channel Interfaces 48
Setting the Interface Administrative State 49
Configuring Interface Modes 49
Configuring the Interface Description 50
Configuring Unified Ports 51
Configuring Port Speeds 52
Configuring Trunk Mode 53
Note 54
Autosensing 54
Understanding Bit Error Thresholds 55
Configuring Global Attributes for Fibre Channel Interfaces 56
Configuring Switch Port Attribute Default Values 56
Information About N Port Identifier Virtualization 56
Enabling N Port Identifier Virtualization 57
Example Port Channel Configurations 58
Verifying Fibre Channel Interfaces 58
Verifying SFP Transmitter Types 58
Verifying Interface Information 58
Verifying BB_Credit Information 60
Default Fibre Channel Interface Settings 61

CHAPTER 7 **Configuring and Managing VSANs 63**
Configuring and Managing VSANs 63
Information About VSANs 63

VSAN Topologies 63
VSAN Advantages 66
VSANs Versus Zones 66
Guidelines and Limitations for VSANs 67
About VSAN Creation 68
Creating VSANs Statically 68
Port VSAN Membership 69
Assigning Static Port VSAN Membership 69
Default VSANs 70
Isolated VSANs 70
Displaying Isolated VSAN Membership 71
Operational State of a VSAN 71
Static VSAN Deletion 71
Deleting Static VSANs 72
Interop Mode 73
Displaying the Static VSAN Configuration 73
Default Settings for VSANs 73

CHAPTER 8 **Configuring SAN Port Channels** 75
Configuring SAN Port Channels 75
Information About SAN Port Channels 75
Understanding Port Channels and VSAN Trunking 76
Understanding Load Balancing 77
Configuring SAN Port Channels 79
SAN Port Channel Configuration Guidelines 80
Creating a SAN Port Channel 81
About Port Channel Modes 82
About SAN Port Channel Deletion 84
Interfaces in a SAN Port Channel 84
About Interface Addition to a SAN Port Channel 84
Adding an Interface to a SAN Port Channel 86
Forcing an Interface Addition 86
About Interface Deletion from a SAN Port Channel 87
Deleting an Interface from a SAN Port Channel 87

SAN Port Channel Protocol 87
About Manually Configured Channel Groups 88
Example Port Channel Configurations 88
Verifying SAN Port Channel Configuration 89
Default Settings for SAN Port Channels 90

CHAPTER 9 **Configuring Fibre Channel Domain Parameters 91**
Information About Domain Parameters 91
Fibre Channel Domains 91
Domain Restarts 92
Restarting a Domain 93
Switch Priority 93
Configuring Switch Priority 93
Configuring Fabric Names 94
Incoming RCFs 95
Rejecting Incoming RCFs 95
Autoreconfiguring Merged Fabrics 95
Enabling Autoreconfiguration 96
Domain IDs 96
Domain IDs - Guidelines 96
Configuring Static or Preferred Domain IDs 98
Allowed Domain ID Lists 99
Configuring Allowed Domain ID Lists 99
CFS Distribution of Allowed Domain ID Lists 100
Enabling Distribution 100
Locking the Fabric 101
Committing Changes 101
Discarding Changes 101
Clearing a Fabric Lock 102
Displaying CFS Distribution Status 102
Displaying Pending Changes 102
Displaying Session Status 103
Contiguous Domain ID Assignments 103
Enabling Contiguous Domain ID Assignments 103

FC IDs 104
Persistent FC IDs 104
Enabling the Persistent FC ID Feature 105
Persistent FC ID Configuration Guidelines 105
Configuring Persistent FC IDs 105
Unique Area FC IDs for HBAs 106
Configuring Unique Area FC IDs for an HBA 107
Persistent FC ID Selective Purging 108
Purging Persistent FC IDs 108
Verifying the fcdomain Configuration 109
Default Settings for Fibre Channel Domains 110

CHAPTER 10 **Configuring FCoE VLANs and Virtual Interfaces 113**
Information About Virtual Interfaces 113
Guidelines and Limitations for FCoE VLANs and Virtual Interfaces 113
Configuring Virtual Interfaces 115
Mapping a VSAN to a VLAN 115
Creating a Virtual Fibre Channel Interface 116
Associating a Virtual Fibre Channel Interface to a VSAN 117
Creating an Implicit Virtual Fibre Channel Port Channel Interface 118
Configuring virtual Fibre Channel – Port Channel Interface 119
Verifying the Virtual Interface 120
Mapping VSANs to VLANs Example Configuration 123

CHAPTER 11 **Managing FLOGI, Name Server, and RSCN Databases 127**
Managing FLOGI, Name Server and RSCN Databases 127
Fabric Login 127
Name Server Proxy 128
About Registering Name Server Proxies 128
Registering Name Server Proxies 128
Rejecting Duplicate pWWNs 129
Name Server Database Entries 129
Displaying Name Server Database Entries 129
FDMI 132

Displaying FDMI 132
RSCN 134
Displaying RSCN Information 134
Multi-pid Option 134
Configuring the multi-pid Option 135
Suppressing Domain Format SW-RSCNs 136
Clearing RSCN Statistics 136
Configuring the RSCN Timer 136
Verifying the RSCN Timer Configuration 138
RSCN Timer Configuration Distribution 138
Default Settings for RSCN 141

CHAPTER 12 **Distributing Device Alias Services** 143
Distributing Device Alias Services 143
Information About Device Aliases 143
Device Alias Features 143
Device Alias Requirements 144
Device Alias Databases 144
Creating Device Aliases 144
Device Alias Modes 145
Device Alias Mode Guidelines and Limitations for Device Alias Services 146
Configuring Device Alias Modes 146
Device Alias Distribution 147
Locking the Fabric 148
Committing Changes 148
Discarding Changes 149
Overriding the Fabric Lock 150
Disabling and Enabling Device Alias Distribution 150
Device Alias Database Merge Guidelines 151
Verifying the Device Alias Configuration 151
Default Settings for Device Alias Services 152

CHAPTER 13 **Configuring and Managing Zones** 153
Information About Zones 153

Information About Zoning 153
Zoning Features 153
Zoning Example 155
Zone Implementation 155
Active and Full Zone Sets 156
Configuring a Zone 158
Configuration Examples 158
Zone Sets 159
Activating a Zone Set 160
Default Zone 161
Configuring the Default Zone Access Permission 161
FC Alias Creation 162
Creating FC Aliases 162
Creating Zone Sets and Adding Member Zones 164
Zone Enforcement 165
Zone Set Distribution 165
Enabling Full Zone Set Distribution 165
Enabling a One-Time Distribution 166
Recovering from Link Isolation 167
Importing and Exporting Zone Sets 167
Zone Set Duplication 168
Copying Zone Sets 168
Renaming Zones, Zone Sets, and Aliases 168
Cloning Zones, Zone Sets and FC Aliases 169
Clearing the Zone Server Database 170
Verifying the Zone Configuration 170
Enhanced Zoning 171
Enhanced Zoning 171
Changing from Basic Zoning to Enhanced Zoning 172
Changing from Enhanced Zoning to Basic Zoning 172
Enabling Enhanced Zoning 173
Modifying the Zone Database 173
Releasing Zone Database Locks 174
Verifying Enhanced Zone Information 175

Merging the Database 175
Configuring Zone Merge Control Policies 176
Default Zone Policies 176
Configuring System Default Zoning Settings 177
About Smart Zoning 178
Smart Zoning Member Configuration 179
Enabling Smart Zoning on a VSAN 179
Setting Default Value for Smart Zoning 179
Converting Zones Automatically to Smart Zoning 180
Configuring Device Types for Zone Members 181
Removing Smart Zoning Configuration 181
Disabling Smart Zoning at Zone Level in the Basic Zoning Mode 182
Disabling Smart Zoning at Zone Level for a VSAN in the Enhanced Zoning Mode 182
Default Settings for Zones 183

CHAPTER 14 **Advanced Fibre Channel Features** 185
Advanced Fibre Channel Features and Concepts 185
Fibre Channel Timeout Values 185
Timer Configuration Across All VSANs 185
Timer Configuration Per-VSAN 186
fctimer Distribution 187
Enabling or Disabling fctimer Distribution 187
Committing fctimer Changes 188
Discarding fctimer Changes 188
Overriding the Fabric Lock 189
Fabric Database Merge Guidelines 189
Verifying Configured fctimer Values 190
World Wide Names 190
Verifying the WWN Configuration 190
Link Initialization WWN Usage 191
Configuring a Secondary MAC Address 191
FC ID Allocation for HBAs 192
Default Company ID List 192
Verifying the Company ID Configuration 193

Switch Interoperability 194
About Interop Mode 194
Configuring Interop Mode 3 195
Default Settings for Advanced Fibre Channel Features 196

Preface

This preface includes the following sections:

- Audience, on page xiii
- Document Conventions, on page xiii
- Related Documentation for Cisco Nexus 9000 Series Switches, on page xiv
- Documentation Feedback, on page xiv
- Communications, Services, and Additional Information, on page xiv

Audience

This publication is for network administrators who install, configure, and maintain Cisco Nexus switches.

Document Conventions

Command descriptions use the following conventions:

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

Convention	Description
`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 includes the quotation marks.

Examples use the following conventions:

Convention	Description
`screen font`	Terminal sessions and information the switch displays are in screen font.
`boldface screen font`	Information that 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.

Related Documentation for Cisco Nexus 9000 Series Switches

The entire Cisco Nexus 9000 Series switch documentation set is available at the following URL:

http://www.cisco.com/en/US/products/ps13386/tsd_products_support_series_home.html

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to nexus9k-docfeedback@cisco.com. We appreciate your feedback.

Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at Cisco Profile Manager.
- To get the business impact 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.

Preface

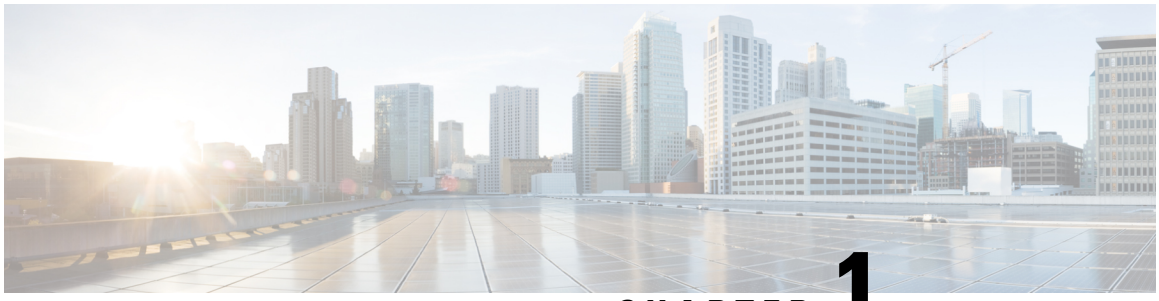

CHAPTER 1

New and Changed Information

This chapter contains the following sections:

- New and Changed Information, on page 1

New and Changed Information

This chapter provides release-specific information for each new and changed feature in the *Cisco Nexus 9000 Series NX-OS SAN Switching Configuration Guide, Release 10.1(x)*.

Table 1: New and Changed Features for Release 10.1(x)

Feature	Description	Changed in Release	Where Documented
FC FCoE Switch mode Support	Added FC FCoE switch mode support for Cisco Nexus N9K-C93360YC-FX2 platform switch.	10.1(1)	Enabling FC/FCoE Switch Mode, on page 9

New and Changed Information

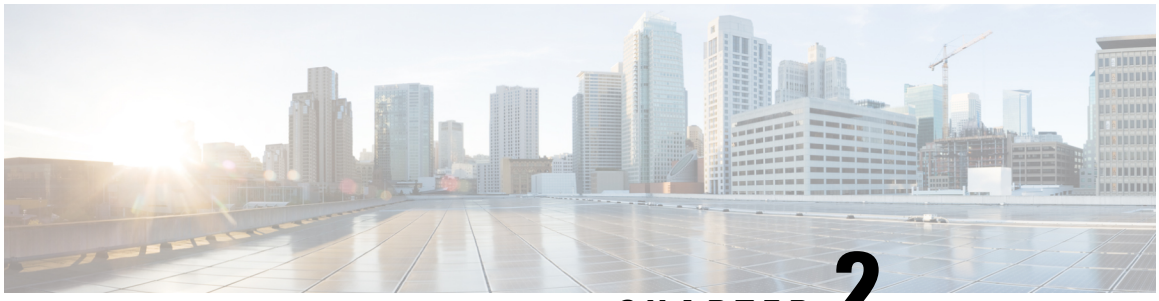

CHAPTER 2

Hardware Support for SAN Switching

- Hardware Support for SAN Switching, on page 3
- Supported Platforms, on page 3

Hardware Support for SAN Switching

The following table lists the Cisco Nexus 9000 Series hardware that supports SAN switching.

Table 2: Cisco Nexus 9300 Series Switches: Supported Hardware

Model (PID)	FC E Port	FCoE E Port	FC Edge Port	FCoE Edge Port	FEX Support
N9K-C93180YC-FX	Yes	No	Yes	Yes	No
N9K-C93360YC-FX2	Yes	No	Yes	Yes	No

The following FC SFPs are supported:

- DS-SFP-FC8G-SW
- DS-SFP-FC16G-SW
- DS-SFP-FC32G-SW
- DS-SFP-FC32G LW is supported only for long distance ISLs (supported on N9K-C93180YC-FX)

Supported Platforms

Starting with Cisco NX-OS release 7.0(3)I7(1), use the Nexus Switch Platform Support Matrix to know from which Cisco NX-OS releases various Cisco Nexus 9000 and 3000 switches support a selected feature.

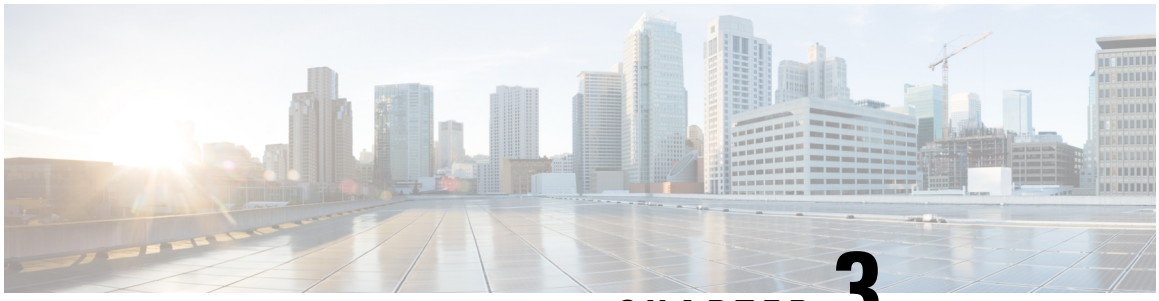

CHAPTER 3

Overview

This chapter contains the following sections:

- Licensing Requirements, on page 5
- SAN Switching Overview, on page 5
- SAN Switching General Guidelines and Limitations, on page 7

Licensing Requirements

For a complete explanation of Cisco NX-OS licensing recommendations and how to obtain and apply licenses, see the *Cisco NX-OS Licensing Guide* and the *Cisco NX-OS Licensing Options Guide*.

SAN Switching Overview

This chapter provides an overview of SAN switching for Cisco Nexus 9000 devices. This chapter includes the following sections:

Domain Parameters

The Fibre Channel domain (fcdomain) feature performs principal switch selection, domain ID distribution, FC ID allocation, and fabric reconfiguration functions as described in the FC-SW-2 standards. The domains are configured per VSAN . If you do not configure a domain ID, the local switch uses a random ID.

N port virtualizer (NPV) is a complementary feature that reduces the number of Fibre Channel domain IDs in core-edge SANs. Cisco Nexus 9000 series fabric switches operating in the NPV mode do not join a fabric; they only pass traffic between core switch links and end devices, which eliminates the domain IDs for these switches. NPIV is used by edge switches in the NPV mode to log in to multiple end devices that share a link to the core switch.

VSAN Trunking

Trunking, also known as VSAN trunking, enables interconnect ports to transmit and receive frames in more than one VSAN over the same physical link. Trunking is supported on E ports and F ports.

Virtual SANs

Virtual SANs (VSANs) partition a single physical SAN into multiple VSANs. VSANs allow the 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.

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 reconfiguration and error conditions within an individual VSAN. The strict traffic segregation provided by VSANs can ensure that the control and data traffic of a specified VSAN are confined within the VSAN's own domain, which increases SAN security. VSANs can reduce costs by facilitating consolidation of isolated SAN islands into a common infrastructure without compromising availability.

You can create administrator roles that are limited in scope to certain VSANs. For example, you can set up a network administrator role to allow configuration of all platform-specific capabilities and other roles 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.

Zoning

Zoning provides access control for devices within a SAN. The 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)

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 the hardware, and none of them cause performance degradation.

Device Alias Services

The software supports Device Alias Services (device alias) fabric wide. Device alias distribution allows you to move host bus adapters (HBAs) between VSANs without manually reentering alias names.

Fibre Channel Routing

Fabric Shortest Path First (FSPF) is the protocol used by Fibre Channel fabrics. FSPF 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 if a failure occurs on 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.

Advanced Fibre Channel Features

You can configure Fibre Channel protocol-related timer values for distributed services, error detection, and resource allocation.

You must uniquely associate the WWN to a single switch. The principal switch selection and the allocation of domain IDs rely on the WWN. .

Fibre Channel standards require that you allocate a unique FC ID to an N port that is attached to an F port in any switch.

Fabric Configuration Servers

The Fabric Configuration Server (FCS) provides discovery of topology attributes and maintains a repository of configuration information of fabric elements. A management application is usually connected to the FCS on the switch through an N port. Multiple VSANs constitute a fabric, where one instance of the FCS is present per VSAN.

SAN Switching General Guidelines and Limitations

The following are the general following guidelines and limitations of SAN switching:

- SAN switching is supported only on Cisco Nexus C93180YC-FX and C93360YC-FX2 switches.
- Beginning with 9.3(5), smart-zoning is supported.
- VE-port or virtual expansion port (ISL) is not supported.
- Dynamic Port VSAN Membership (*DPVM*) not supported.
- Fabric Extender (*FEX*) with switch mode is not supported
- IP over Fibre Channel (*IPFC*) function is not supported.
- Inter VSAN Routing(*IVR*) is not supported
- XML and DME of CLIs are not supported.
- 100G FCoE is not supported.
- OBFL (show logging onboard) feature support is limited to the error statistics.

Note For more information on OBFL, see: *Cisco Nexus 9000 Series NX-OS Troubleshooting Guide, Release 9.3(x)*

- Nexus 9000 only supports the IDLE fill pattern on 8 Gbps Fibre Channel interfaces. For Nexus 9000 FC interface to operate at 8 Gbps, peer device must be configured to use a matching IDLE fill pattern. Most server and target FC interfaces do not support this and thus cannot connect to Nexus 9000 at 8 Gbps. To interoperate with other Fibre Channel switches at 8 Gbps ensure the peer switch FC interface also uses a matching IDLE fill pattern. For Cisco MDS switches, configure using the **switchport fill-pattern** interface configuration command. To connect to a peer Nexus 9000 at 8 Gbps, use no fill pattern configuration, as both devices use matching IDLE fill patterns by default.
- Beginning with Cisco NX-OS Release 9.3(5), 25G FCoE is supported.
- Beginning with Cisco NX-OS Release 9.3(5), enhanced zoning is supported.
- Beginning with Cisco NX-OS Release 9.3(5), enhanced device alias is supported.
- Beginning with Cisco NX-OS Release 9.3(5), FDMI is supported.
- Beginning with Cisco NX-OS Release 9.3(5), FCoE over vPC is supported.
- Beginning with Cisco NX-OS Release 9.3(5), ISSU is supported in switch mode.

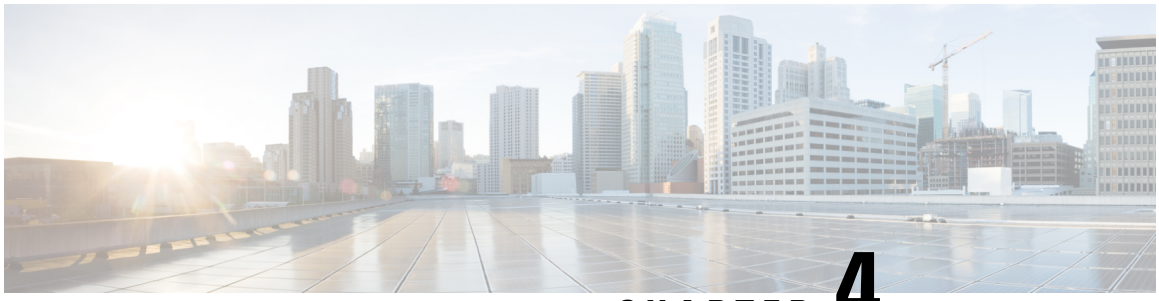

CHAPTER 4

Enabling FC/FCoE Switch Mode

This chapter contains the following sections:

To enable FC/FCoE switch mode on Cisco Nexus 9000 series switches, you must configure **feature-set fcoe**.

Note For more information about enabling NPV mode on Cisco Nexus 9000 series switches, see the relevant version of *Cisco Nexus 9000 Series NX-OS FC-NPV and FCoE-NPV Configuration Guide* on cisco.com.

- Guidelines and Limitations for FC Switch Mode, on page 9
- Enabling FC/FCoE , on page 9
- Disabling FC/FCoE , on page 10
- Disabling LAN Traffic on an FCoE Link, on page 11
- Configuring the FC-Map , on page 12
- Configuring the Fabric Priority, on page 12
- Setting the Advertisement Interval, on page 13

Guidelines and Limitations for FC Switch Mode

- Beginning with Release 10.1(1), FC switch mode is supported on Cisco Nexus 93360YC-FX2.

Enabling FC/FCoE

You can enable FC/FCoE on the switch; however, enabling FCoE on VLAN 1 is not supported.

Note Alternatively, you can use the **FC set up script** included in the **Cisco NX-OS Setup Utility** to enable FC/FCoE. For more information, see the relevant version of *Cisco Nexus 9000 Series NX-OS Fundamentals Configuration Guide* on cisco.com.

Note All the Fibre Channel features of the Cisco Nexus device are packaged in the FC Plugin. When you enable FC/FCoE, the switch software checks for the SAN_ENTERPRISE_PKG license. If it finds the license, the software loads the plugin. The package FC_PORT_ACTIVATION_PKG is required to enable FC port license.

After the FC Plugin is loaded, the following occurs:

- All Fibre Channel and FCoE-related CLI are available

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **install feature-set fcoe**
3. switch(config)# **feature-set fcoe**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **install feature-set fcoe**	Installs feature set FCoE.
Step 3	switch(config)# **feature-set fcoe**	Enables the FC/FCoE capability.

Example

This example shows how to enable FC/FCoE on the switch:

```
switch# configure terminal
switch(config)# install feature-set fcoe
switch(config)# feature-set fcoe
```

Disabling FC/FCoE

After you disable FC/FCoE, all FC/FCoE commands are removed from the CLI and the FC/FCoE configuration is deleted.

Note The command **no feature-set fcoe** is not allowed if there are FC ports on the switch. If there are FC ports on the switch, you must convert them to Ethernet ports before issuing this command. On Cisco Nexus C93180YC-FX and C93360YC-FX2 switches, you must reload the switch after disabling feature-set fcoe.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **no feature-set fcoe**

3. switch(config)# **no install feature-set fcoe**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **no feature-set fcoe**	Disables the FC/FCoE capability.
Step 3	switch(config)# **no install feature-set fcoe**	Uninstalls feature set FCoE.

Example

This example shows how to disable FC/FCoE on the switch:

```
switch# configure terminal
switch(config)# no feature-set fcoe
switch(config)# no install feature-set fcoe
```

Disabling LAN Traffic on an FCoE Link

You can disable LAN traffic on an FCoE link.

DCBX allows the switch to send a LAN Logical Link Status (LLS) message to a directly connected CNA. Enter the **shutdown lan** command to send an LLS-Down message to the CNA. This command causes all VLANs on the interface that are not enabled for FCoE to be brought down. If a VLAN on the interface is enabled for FCoE, it continues to carry SAN traffic without any interruption.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface ethernet** *slot*/*port*
3. switch(config-if)# **shutdown lan**
4. (Optional) switch(config-if)# **no shutdown lan**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface ethernet** *slot*/*port*	Specifies an interface to configure, and enters interface configuration mode.
Step 3	switch(config-if)# **shutdown lan**	Shuts down Ethernet traffic on the interface. If the interface is part of an FCoE VLAN, the shutdown has no impact on the FCoE traffic.
Step 4	(Optional) switch(config-if)# **no shutdown lan**	Reenables Ethernet traffic on the interface.

Configuring the FC-Map

Note We recommend using the "Mapping a VSAN to a VLAN " method for preserving fabric isolation and leaving the FC-MAP default.

You can prevent data corruption due to cross-fabric talk by configuring an FC-Map that identifies the Fibre Channel fabric for this Cisco Nexus device. When the FC-Map is configured, the switch discards the MAC addresses that are not part of the current fabric.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **fcoe fcmap** *fabric-map*
3. (Optional) switch(config)# **no fcoe fcmap** *fabric-map*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **fcoe fcmap** *fabric-map*	Configures the global FC-Map. The default value is 0E.FC.00. The range is from 0E.FC.00 to 0E.FC.FF.
Step 3	(Optional) switch(config)# **no fcoe fcmap** *fabric-map*	Resets the global FC-Map to the default value of 0E.FC.00.

Example

This example shows how to configure the global FC-Map:

```
switch# configure terminal
switch(config)# fcoe fcmap 0x0efc2a
```

Configuring the Fabric Priority

The Cisco Nexus device advertises its priority. The priority is used by the CNAs in the fabric to determine the best switch to connect to.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **fcoe fcf-priority** *fabric-priority*
3. (Optional) switch(config)# **no fcoe fcf-priority** *fabric-priority*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **fcoe fcf-priority** *fabric-priority*	Configures the global fabric priority. The default value is 128. The range is from 0 (higher) to 255 (lower).
Step 3	(Optional) switch(config)# **no fcoe fcf-priority** *fabric-priority*	Resets the global fabric priority to the default value of 128.

Example

This example shows how to configure the global fabric priority:

```
switch# configure terminal
switch(config)# fcoe fcf-priority 42
```

Setting the Advertisement Interval

You can configure the interval for Fibre Channel fabric advertisement on the switch.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **fcoe fka-adv-period** *interval*
3. (Optional) switch(config)# **no fcoe fka-adv-period** *interval*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **fcoe fka-adv-period** *interval*	Configures the advertisement interval for the fabric. The default value is 8 seconds. The range is from 4 to 60 seconds.
Step 3	(Optional) switch(config)# **no fcoe fka-adv-period** *interval*	Resets the advertisement interval for the fabric to its default value of 8 seconds.

Example

This example shows how to configure the advertisement interval for the fabric:

```
switch# configure terminal
switch(config)# fcoe fka-adv-period 42
```

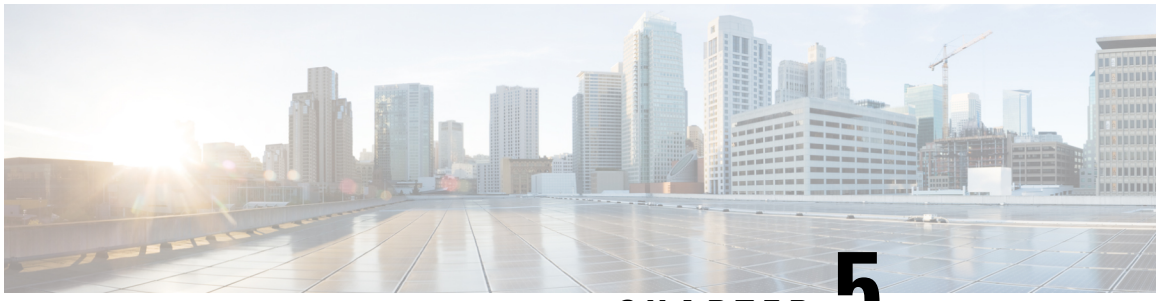

CHAPTER 5

Configuring FCoE

This chapter contains the following sections:

- FCoE Topologies, on page 15
- FCoE Best Practices, on page 18
- Guidelines and Limitations, on page 20
- Configuring FC/FCoE, on page 21

FCoE Topologies

Directly Connected CNA Topology

The Cisco Nexus device can be deployed as a Fibre Channel Forwarder (FCF) as shown in the following figure.

Figure 1: Directly Connected Fibre Channel Forwarder

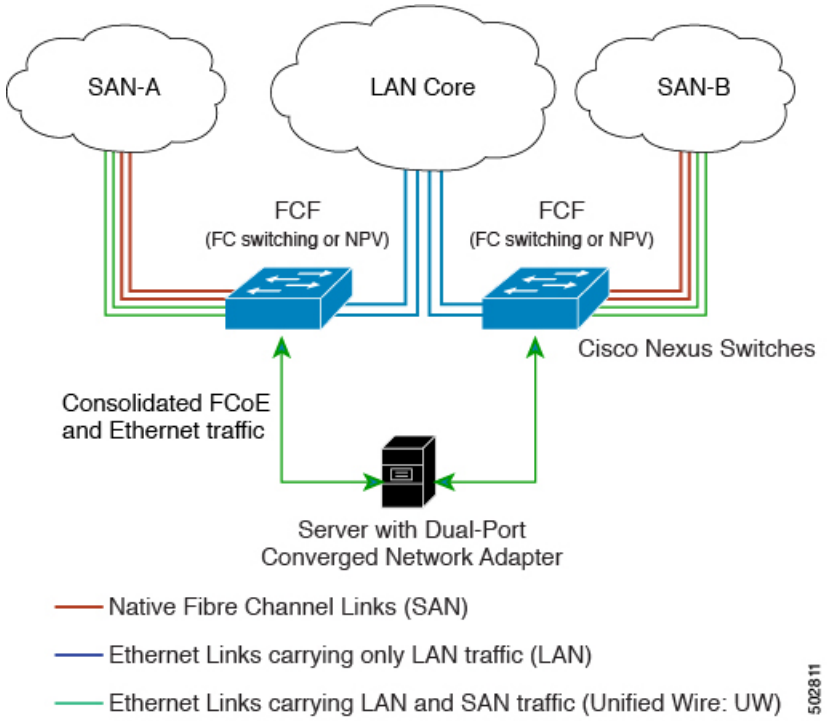

The following rules are used to process FIP frames to avoid the FCF being used as a transit between an FCoE node (ENode) and another FCF. These rules also prevent login sessions between ENodes and FCFs in different fabrics.

- FIP solicitation and login frames received from the CNAs are processed by the FCF and are not forwarded.
- If an FCF receives solicitations and advertisements from other FCFs over an interface, the following occurs:
  - The frames are ignored and discarded if the FC-MAP value in the frame matches the value of the FCF (the FCF is in the same fabric).
  - The interface is placed in the "FCoE Isolated" state if the FC-MAP value in the FIP frame does not match that of the FCF (the FCF is in a different fabric).

CNAs cannot discover or log in to FCFs that are reachable only through a transit Cisco Nexus FCF. The Cisco Nexus device cannot perform the FCoE transit function between a CNA and another FCF due to hardware limitations.

Because the Cisco Nexus FCF cannot perform the transit FCoE function, you must design your network topology so that the active Spanning Tree Protocol (STP) path of FCoE VLANs is always over the directly connected links between the CNA and the FCF. Make sure that you configure the FCoE VLAN on the directly connected links only.

Remotely Connected CNA Topology

The Cisco Nexus device can be deployed as a Fibre Channel Forwarder (FCF) for remotely connected CNAs, but not as a FIP snooping bridge, as shown in the following figure.

Figure 2: Remotely Connected Fibre Channel Forwarder

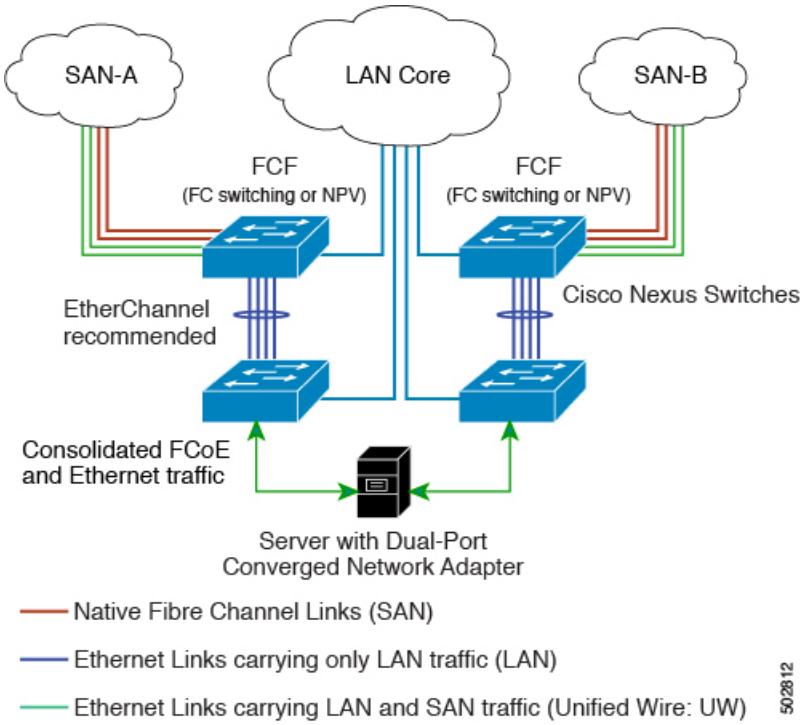

The following rules are used to process FIP frames to avoid the FCF being used as a transit between an ENode and another FCF. These rules also prevent login sessions between ENodes and FCFs in different fabrics.

- FIP solicitation and login frames received from the CNAs are processed by the FCF and are not forwarded.
- If an FCF receives solicitations and advertisements from other FCFs over an interface, the following occurs:
  - The frames are ignored and discarded if the FC-MAP value in the frame matches the value of the FCF (the FCF is in the same fabric).
  - The interface is placed in the "FCoE Isolated" state if the FC-MAP value in the FIP frame does not match that of the FCF (the FCF is in a different fabric).

Because the Cisco Nexus FCF cannot perform the transit FCoE function, you must design your network topology so that the active STP path of FCoE VLANs is always over the directly connected links between the CNA and the FCF. Make sure that you configure the FCoE VLAN on the directly connected links only.

FCoE Best Practices

Directly Connected CNA Best Practice

The following figure shows a best practices topology for an access network that is using directly connected CNAs with Cisco Nexus devices.

Figure 3: Directly Connected CNA

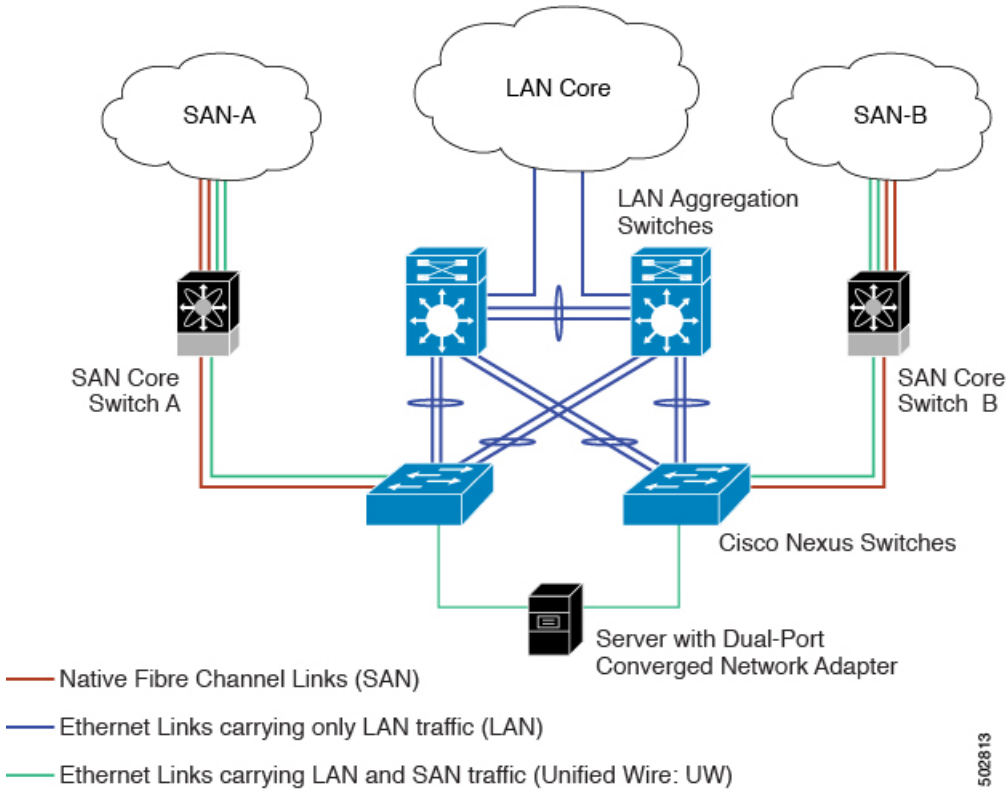

Follow these configuration best practices for the deployment topology in the preceding figure:

1. You must configure a unique dedicated VLAN at every converged access switch to carry traffic for each Virtual Fabric (VSAN) in the SAN (for example, VLAN 1002 for VSAN 1, VLAN 1003 for VSAN 2, and so on). If you enable Multiple Spanning Tree (MST), you must use a separate MST instance for FCoE VLANs.
2. You must configure the unified fabric (UF) links as trunk ports. Do not configure the FCoE VLAN as a native VLAN. You must configure all FCoE VLANs as members of the UF links to allow extensions for VF_Port trunking and VSAN management for the virtual Fibre Channel interfaces.

Note A unified wire carries both Ethernet and FCoE traffic.

3. You must configure the UF links as spanning-tree edge ports.

4. You must not configure the FCoE VLANs as members of Ethernet links that are not designated to carry FCoE traffic because you want to ensure that the scope of the STP for the FCoE VLANs is limited to UF links only.
5. If the converged access switches (in the same SAN fabric or in another) need to be connected to each other over Ethernet links for a LAN alternate path, you must explicitly configure such links to exclude all FCoE VLANs from membership. This action ensures that the scope of the STP for the FCoE VLANs is limited to UF links only.
6. You must use separate FCoE VLANs for FCoE in SAN-A and SAN-B.

Remotely Connected CNA Best Practice

The following figure shows a best practices topology for an access network using remotely connected CNAs with Cisco Nexus devices.

Figure 4: Remotely Connected CNAs

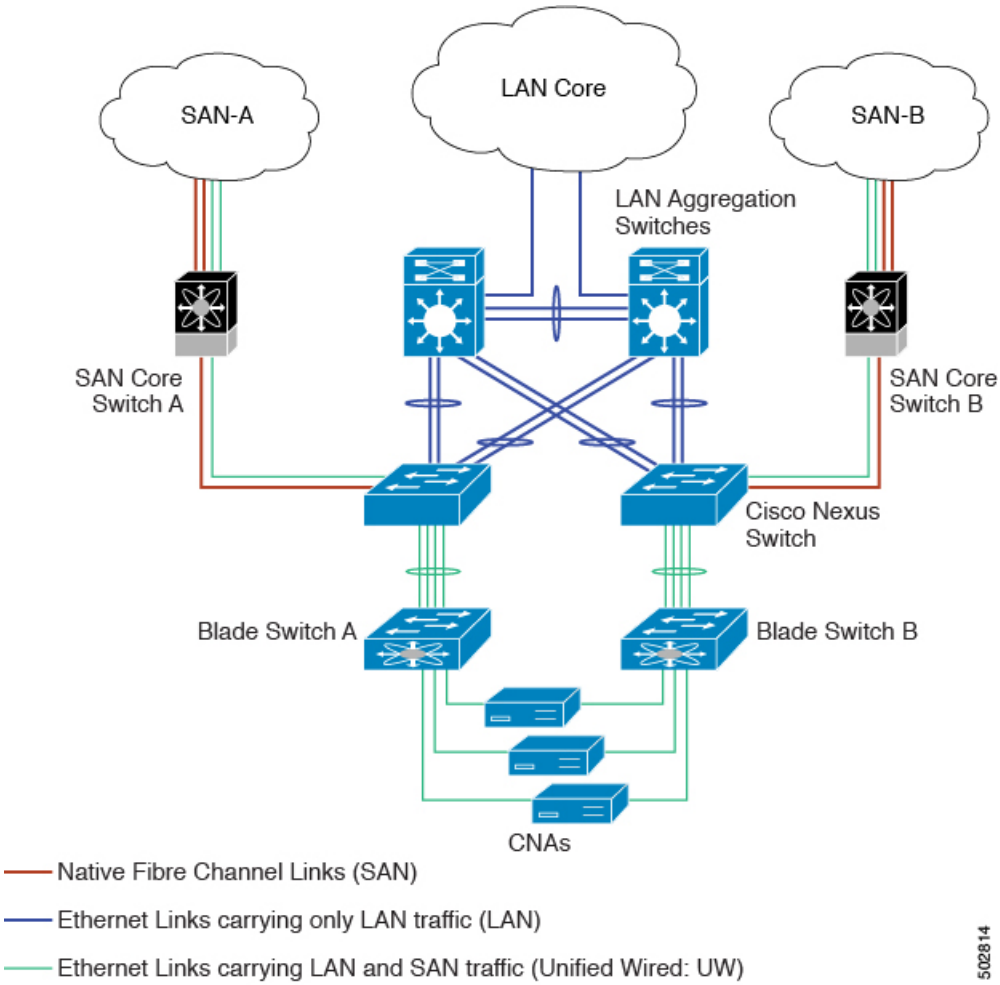

Follow these configuration best practices for the deployment topology in the preceding figure:

1. You must configure a unique dedicated VLAN at every converged access switch to carry traffic for each Virtual Fabric (VSAN) in the SAN (for example, VLAN 1002 for VSAN 1, VLAN 1003 for VSAN 2, and so on). If you enable MST, you must use a separate MST instance for FCoE VLANs.
2. You must configure the unified fabric (UF) links as trunk ports. Do not configure the FCoE VLAN as a native VLAN. You must configure all FCoE VLANs as members of the UF links to allow extensions for VF_Port trunking and VSAN management for the virtual Fibre Channel interfaces.

Note A unified fabric link carries both Ethernet and FCoE traffic.

3. You must configure the CNAs and the blade switches as spanning-tree edge ports.
4. A blade switch must connect to exactly one Cisco Nexus device converged access switch, preferably over an EtherChannel, to avoid disruption due to STP reconvergence on events such as provisioning new links or blade switches.
5. You must configure the Cisco Nexus device converged access switch with a better STP priority than the blade switches that are connected to it. This requirement allows you to create an island of FCoE VLANs where the converged access switch is the spanning-tree root and all the blade switches connected to it become downstream nodes.
6. Do not configure the FCoE VLANs as members of Ethernet links that are not designated to carry FCoE traffic because you want to ensure that the scope of the STP for the FCoE VLANs is limited to UF links only.
7. If the converged access switches and/or the blade switches need to be connected to each over Ethernet links for the purposes of LAN alternate pathing, you must explicitly configure such links to exclude all FCoE VLANs from membership. This action ensures the scope of the STP for FCoE VLANs is limited to UF links only.
8. You must use separate FCoE VLANs for FCoE in SAN-A and SAN-B.

Guidelines and Limitations

FC/FCoE has the following guidelines and limitations:

- Enabling FCoE on VLAN 1 is not supported.
- Enabling FCoE requires enabling the LLDP feature using **feature lldp**, as LLDP is not enabled by default.
- FCOE is not supported with Copper SFPs.
- Direct connect FCoE (that is, a direct connect to CNAs through a bind interface) is not supported on a port channel of a Cisco Nexus device interface if it is configured to have more than one interface. Direct connect FCoE is supported on port channels with a single link to allow for FCoE from a CNA connected through a vPC with one 10/40 GB link to each upstream switch.
- Ethernet interfaces used for vFC must have the QOS policy configured manually regardless of default or custom policy defined globally.

Note For a description of the default quality of service (QoS) policies for FC/FCoE, see the Quality of Service guide for your device. For the Nexus software release that you are using. The available versions of this document can be found at the following URL: https://www.cisco.com/c/en/us/support/switches/nexus-9000-series-switches/products-installation-and-configuration-guides-list.html

Configuring FC/FCoE

Perform TCAM Carving

This section explains how to perform TCAM carving.

SUMMARY STEPS

1. Perform TCAM carving.
2. Use the command **show hardware access-list tcam region** to view the configured TCAM region size.
3. Save the configuration and use the command **reload** to reload the switch.

DETAILED STEPS

Step 1 Perform TCAM carving.

Example:

```
Switch(config)# hardware access-list tcam region ing-racl 1536
Switch(config)# hardware access-list tcam region ing-ifacl 256
Switch(config)# hardware access-list tcam region ing-redirect 256
```

Step 2 Use the command **show hardware access-list tcam region** to view the configured TCAM region size.

Example:

```
Switch(config)# show hardware access-list tcam region
Switch(config)#
```

Step 3 Save the configuration and use the command **reload** to reload the switch.

Example:

```
Switch(config)# reload
Switch(config)#
```

What to do next

You must reload the switch after carving TCAM

Configuring LLDP

This section explains how to configure LLDP.

SUMMARY STEPS

1. **configure terminal**
2. [**no**]**feature lldp**

DETAILED STEPS

Step 1 **configure terminal**

Enters global configuration mode.

Step 2 [**no**]**feature lldp**

Enables or disables LLDP on the device. LLDP is disabled by default.

Configuring Default QoS

There are four types of FCoE default policies: network QoS, output queuing, input queuing, and QoS. You can enable the FCoE default policies by enabling the FCoE feature using the **feature-set fcoe command** command. The default QoS ingress policy, **default-fcoe-in-policy**, is implicitly attached to all FC and SAN-port-channel interfaces to enable FC to FCoE traffic; this can be verified by using **show interface** {*fc slot/port* | *san-port-channel* <*no*>} **all** command. The default QoS policy uses CoS3 and Q1 for all FC and FCoE traffic.

Configuring User Defined QoS

To use a different queue or CoS value for FCoE traffic, create user-defined policies. The user-defined QoS ingress policy has to be created and attached explicitly to both FC and FCoE interfaces to enable traffic to use a different queue or CoS. User-defined QoS policies must be created and activated for system-wide QoS.

Note The Ethernet or port-channel interface must be configured with MTU 9216 (or the maximum available MTU size) for supporting FCoE.

The following example demonstrates how to configure and activate user-defined QoS policies that use CoS3 and Q2 for all FC and FCoE traffic.

- Creating a user-defined network QOS policy:

```
switch(config)# policy-map type network-qos fcoe_nq
switch(config-pmap-nqos)# class type network-qos c-nq1
switch(config-pmap-nqos-c)# mtu 1500
switch(config-pmap-nqos-c)# class type network-qos c-nq2
switch(config-pmap-nqos-c)# mtu 9216
```

```
switch(config-pmap-nqos-c)# pause pfc-cos 3
switch(config-pmap-nqos-c)# class type network-qos c-nq3
switch(config-pmap-nqos-c)# mtu 1500
switch(config-pmap-nqos-c)# class type network-qos c-nq-default
switch(config-pmap-nqos-c)# mtu 1500
switch(config-pmap-nqos-c)# exit
switch(config-pmap-nqos)# exit
switch(config)#
```

- Creating a user-defined input queuing policy:

```
switch(config)# policy-map type queuing fcoe-in-policy
switch(config-pmap-que)# class type queuing c-in-q2
switch(config-pmap-c-que)# bandwidth percent 50
switch(config-pmap-c-que)# class type queuing c-in-q-default
switch(config-pmap-c-que)# bandwidth percent 50
switch(config-pmap-c-que)# exit
switch(config-pmap-que)# exit
switch(config)
```

- Creating a user-defined output queuing policy:

```
switch(config)# policy-map type queuing fcoe-out-policy
switch(config-pmap-que)# class type queuing c-out-q3
switch(config-pmap-c-que)# priority level 1
switch(config-pmap-c-que)# class type queuing c-out-q-default
switch(config-pmap-c-que)# bandwidth remaining percent 50
switch(config-pmap-c-que)# class type queuing c-out-q1
switch(config-pmap-c-que)# bandwidth remaining percent 0
switch(config-pmap-c-que)# class type queuing c-out-q2
switch(config-pmap-c-que)# bandwidth remaining percent 50
switch(config-pmap-c-que)# exit
switch(config-pmap-que)# exit
switch(config)#
```

- Creating a user-defined QoS input policy:

```
switch(config)# class-map type qos match-any fcoe
switch(config-cmap-qos)# match protocol fcoe
switch(config-cmap-qos)# match cos 3
switch(config-cmap-qos)# exit
switch(config)#
switch(config)# policy-map type qos fcoe_qos_policy
switch(config-pmap-qos)# class fcoe
switch(config-pmap-c-qos)# set cos 3
switch(config-pmap-c-qos)# set qos-group 2
switch(config-pmap-c-qos)# exit
switch(config-pmap-qos)# exit
switch(config)#
```

- Activating a user-defined system QoS policy:

```
switch(config)# system qos
switch(config-sys-qos)# service-policy type queuing input fcoe-in-policy
switch(config-sys-qos)# service-policy type queuing output fcoe-out-policy
switch(config-sys-qos)# service-policy type network-qos fcoe_nq
switch(config-sys-qos)# exit
switch(config)#
```

- Applying a QoS input policy to an FC or FCoE interface:

```
switch# conf
```

```
switch(config)# interface {fc <slot>/<port> | ethernet <slot>/<port> | san-port-channel
 <no> | port-channel <no>}
switch(config-if)# service-policy type qos input fcoe_qos_policy
```

- Removing a QoS input policy from an FC or FCoE interface:

```
switch# conf
switch(config)# interface {fc <slot>/<port> | ethernet <slot>/<port> | san-port-channel
 <no> | port-channel <no>}
switch(config-if)# no service-policy type qos input fcoe_qos_policy
```

- Verifying a QoS input policy applied to an FC or FCoE interface:

```
switch# show running-config interface {fc <slot>/<port> | interface <slot>/<port> |
san-port-channel <no> | port-channel <no>} all
```

Note

- When a user-defined QoS policy is used, the same QoS input policy must be applied to all FC and FCoE interfaces in the switch.
- Do not configure **match protocol fcoe** under more than one QoS class map, as FCoE traffic is supported only on a single CoS.

Configuring Traffic Shaping

Traffic shaping is used to control access to available bandwidth and to regulate the flow of traffic in order to avoid congestion that can occur when the sent traffic exceeds the access speed. Because traffic shaping limits the rate of transmission of data, you may use this command only when necessary.

The following example demonstrates how to configure traffic shaper:

- The following command displays the default system level settings for all FC interfaces:

```
switch(config)# show running-config all | i i rate
hardware qos fc rate-shaper
switch(config)#
```

- The following example shows how to configure rate shaper. This command is applied on all FC interfaces:

Note Rarely, you may see input discards on any of the 4G, 8G, 16G, or 32G interfaces. Use the command *hardware qos fc rate-shaper [low]*, to configure the rate shape. Because this is a system level configuration, it will apply to all the FC ports and will reduce the rates for all FC ports. The default option of the command *hardware qos fc rate-shaper* is applicable to all FC interfaces.

```
switch(config)# hardware qos fc rate-shaper low
switch(config)#
switch(config)#end
```

FCoE with vPC Configuration Examples

The Cisco Nexus N9K-93180YC-FX and N9K-C93360YC-FX2 devices support vPCs. The vPCscan be configured to increase bandwidth and provide increased load-balancing to the Ethernet fabric. The following are example configurations to explain how to configure FCoE when using vPCs on the Cisco Nexus 9000 Series switches:

Figure 5: FCoE Traffic Flow with Host vPC

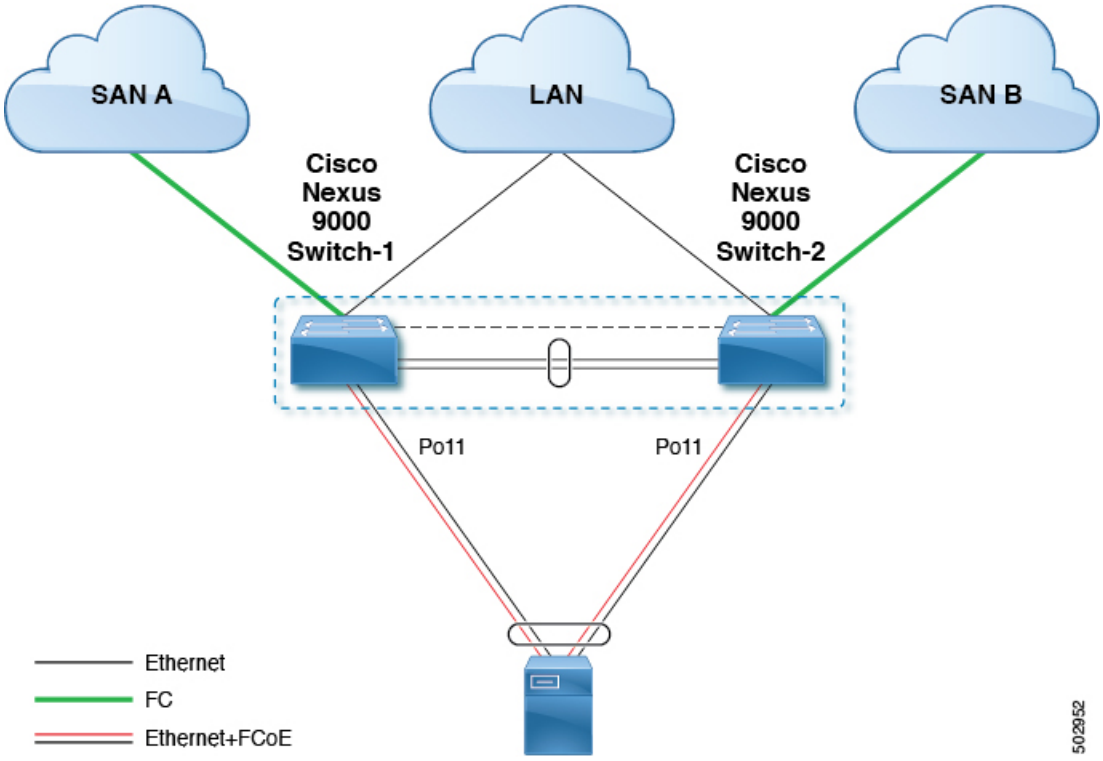

Note FCoE VLANs should not be trunked across vPC peer-links.

Note Only FC uplinks are supported on Cisco Nexus N9K switches (switchmode) that connects to core switches.

The configuration example includes the following parameters:

```
switchname: tme-switch-1
switchname: tme-switch-2
mgmt ip: 172.25.182.66
mgmt ip: 172.25.182.67
```

The configuration example includes the following hardware:

- Emulex CNA or CISCO CNA
- 2 Cisco Nexus 9000 switches running Cisco NX-OS Release 10.1(1) or later releases.

The configuration example includes the following considerations and requirements:

- Generation 2 CNAs that support DCBX are required.
- Single host CNA port channel connection to a separate switch. FCoE interfaces will not be brought up if the port channel on a single switch contains more than one member port in a port channel or vPC.
- Cisco NX-OS Release 10.1(1).

Cisco Nexus 9000 Series Switch vPC Configuration Example

This example presumes that the basic configuration has been completed on the switch (for example, IP Address (mgmt0), switchname, and password for the administrator).

Note The configuration must be done on both peer switches in the vPC topology.

SUMMARY STEPS

1. **feature vpc**
2. **vPC domain**
3. **vpc peer-link**
4. **show vpc peer-keepalive**
5. **int po**
6. **vpc**
7. **show vpc statistics**

DETAILED STEPS

	Command or Action	Purpose

Step 1	**feature vpc**
	Example:
	`tme-switch-1# conf t`
	`Enter configuration commands, one per line. End with CNTL/Z.`
	`tme-switch-1(config)# feature vpc`
	`tme-switch-1(config)#`
	
	`tme-switch-2# conf t`
	`Enter configuration commands, one per line. End with CNTL/Z.`
	`tme-switch-2(config)# feature vpc`
	`tme-switch-2(config)#`
		Enable the vPC feature on both peer switches.

Step 2	**vPC domain**
	Example:
	`tme-switch-1(config)# vpc domain 2`
	`tme-switch-1(config-vpc-domain)# peer-keepalive destination 192.165.200.229`
		Configure the vPC domain and peer-keep alive destinations.
		Note In this set up, switch tme-switch-1 has the mgmt IP address of 192.165.200.229 and switch tme-switch-2 has the mgmt IP address of 192.165.200.230.

	Command or Action	Purpose
	`tme-switch-2(config)# vpc domain 2` `tme-switch-2(config-vpc-domain)# peer-keepalive destination 192.165.200.230`	

Step 3	**vpc peer-link**
	Example:
	`tme-switch-1(config)# int port-channel 1`
	`tme-switch-1(config-if)# vpc peer-link`
	Note The spanning tree port type is changed to network port type on vPC peer-link. This will enable STP Bridge Assurance on vPC peer-link provided that the STP Bridge Assurance (which is enabled by default) is not disabled.
	`tme-switch-2(config)# int port-channel 1`
	`tme-switch-2(config-if)# vpc peer-link`
		Configure the port channel interface that will be used as the vPC peer-link.

Step 4	**show vpc peer-keepalive**
	Example:
	`tme-switch-1(config)# show vpc peer-keepalive`
	`vPC keep-alive status : peer is alive`
	`--Destination : 172.25.182.167`
	`--Send status : Success`
	`--Receive status : Success`
	`--Last update from peer : (0) seconds, (975) msec`
	`tme-switch-1(config)#`
	`tme-switch-2(config)# show vpc peer-keepalive`
	`--PC keep-alive status : peer is alive`
	`--Destination : 172.25.182.166`
	`--Send status : Success`
	`--Receive status : Success`
	`--Last update from peer : (0) seconds, (10336) msec`
	`tme-switch-2(config)#`
		Verify that the peer-keepalive can be reached.

Step 5	**int po**
	Example:
	`tme-switch-1(config-if-range)# int po 1`
	`tme-switch-1(config-if)# switchport mode trunk`
	`tme-switch-1(config-if)# no shut`
	`tme-switch-1(config-if)# exit`
	`tme-switch-1(config)# int eth 1/39-40`
	`tme-switch-1(config-if-range)# switchport mode trunk`
	`tme-switch-1(config-if-range)# channel-group 1`
	`tme-switch-1(config-if-range)# no shut`
	`tme-switch-1(config-if-range)#`
	`tme-switch-2(config-if-range)# int po 1`
		Add member ports to the vpc-peer link port channel and bring up the port channel interface.

	Command or Action	Purpose
	tme-switch-2(config-if)# switchport mode trunk tme-switch-2(config-if)# no shut tme-switch-2(config-if)# exit tme-switch-2(config)# int eth 1/39-40 tme-switch-2(config-if-range)# switchport mode trunk tme-switch-2(config-if-range)# channel-group 1 tme-switch-2(config-if-range)# no shut tme-switch-2(config-if-range)# tme-switch-1(config-if-range)# show int po1 port-channel 1 is up Hardware: Port-Channel, address: 000d.ecde.a92f (bia 000d.ecde.a92f) MTU 1500 bytes, BW 20000000 Kbit, DLY 10 usec, reliability 255/255, txload 1/255, rxload 1/255 Encapsulation ARPA Port mode is trunk full-duplex, 10 Gb/s Beacon is turned off Input flow-control is off, output flow-control is off Switchport monitor is off Members in this channel: Eth1/39, Eth1/40 Last clearing of "show interface" counters never 1 minute input rate 1848 bits/sec, 0 packets/sec 1 minute output rate 3488 bits/sec, 3 packets/sec tme-switch-1(config-if-range)# tme-switch-2(config-if-range)# show int po1 port-channel1 is up Hardware: Port-Channel, address: 000d.ecdf.5fae (bia 000d.ecdf.5fae) MTU 1500 bytes, BW 20000000 Kbit, DLY 10 usec, reliability 255/255, txload 1/255, rxload 1/255 Encapsulation ARPA Port mode is trunk full-duplex, 10 Gb/s Beacon is turned off Input flow-control is off, output flow-control is off Switchport monitor is off Members in this channel: Eth1/39, Eth1/40 Last clearing of "show interface" counters never minute input rate 1848 bits/sec, 0 packets/sec minute output rate 3488 bits/sec, 3 packets/sec tme-switch-2(config-if-range)#	
Step 6	**vpc** **Example:** tme-switch-1(config)# int po 11 tme-switch-1(config-if)# vpc 11 tme-switch-1(config-if)# switchport mode trunk tme-switch-1(config-if)# no shut tme-switch-1(config-if)# int eth 1/1 tme-switch-1(config-if)# switchport mode trunk tme-switch-1(config-if)# channel-group 11 tme-switch-1(config-if)# spanning-tree port type edge trunk tme-switch-1(config-if)#	Create the vPC and add member interfaces. **Note** To run FCoE over a vPC topology, the port channel can only have a sinlge member interface. **Note** The vPC number configured under the port channel interface must match on both Nexus 9000 switches. The port channel interface number does not have to match on both switches.

	Command or Action	Purpose

	Warning Edge port type (portfast) should only be enabled on ports connected to a single host. Connecting some devices such as hubs, concentrators, switches, or bridges to this interface when edge port type (portfast) is enabled, can cause temporary bridging loops. Caution should be used in this type of configuration.
	
	`tme-switch-2(config)# int po 11`
	`tme-switch-2(config-if)# vpc 11`
	`tme-switch-2(config-if)# switchport mode trunk`
	`tme-switch-2(config-if)# no shut`
	`tme-switch-2(config-if)# int eth 1/1`
	`tme-switch-2(config-if)# switchport mode trunk`
	`tme-switch-2(config-if)# channel-group 11`
	`tme-switch-2(config-if)# spanning-tree port type edge trunk`
	
	Warning Edge port type (portfast) should only be enabled on ports connected to a single host. Connecting some devices such as hubs, concentrators, switches, or bridges to this interface when edge port type (portfast) is enabled, can cause temporary bridging loops. Caution should be used in this type of configuration.	

Step 7	**show vpc statistics**
	
	Example:
	
	`tme-switch-1(config-if)# show vpc statistics vpc 11`
	`port-channel11 is up`
	`vPC Status: Up, vPC number: 11`
	`Hardware: Port-Channel, address: 000d.ecde.a908 (bia 000d.ecde.a908)`
	`MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec,`
	`reliability 255/255, txload 1/255, rxload 1/255`
	`Encapsulation ARPA`
	`Port mode is trunk`
	`full-duplex, 10 Gb/s`
	`Beacon is turned off`
	`Input flow-control is off, output flow-control is off`
	`Switchport monitor is off`
	`Members in this channel: Eth1/1`
	`Last clearing of "show interface" counters never`
	`minute input rate 4968 bits/sec, 8 packets/sec`
	`minute output rate 792 bits/sec, 1 packets/sec`
	`tme-switch-1(config-if)#`
	
	`tme-switch-2(config-if)# show vpc statistics vpc 11`
	`port-channel11 is up`
	`vPC Status: Up, vPC number: 11`
	`Hardware: Port-Channel, address: 000d.ecdf.5fae`
		Verify that the vPC interfaces are up and operational.

	Command or Action	Purpose
	(bia 000d.ecdf.5fae) MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec, reliability 255/255, txload 1/255, rxload 1/255 Encapsulation ARPA Port mode is trunk full-duplex, 10 Gb/s Beacon is turned off Input flow-control is off, output flow-control is off Switchport monitor is off Members in this channel: Eth1/1 Last clearing of "show interface" counters never minute input rate 4968 bits/sec, 8 packets/sec minute output rate 792 bits/sec, 1 packets/sec tme-switch-1(config-if)#	

Cisco Nexus 9000 Series Switch FCoE Configuration Example

After setting up vPC between the two Nexus 9000 switches, you can configure the FCoE topology. This procedure presumes that basic configuration has been executed on the Nexus 9000 switch that will provide IP Address (mgmt0), switch name, password for admin, etc. and that the vPC configuration has been completed as outlined in the previous section. The following steps will walk through the basic FCoE configuration necessary to set up an FCoE topology in conjunction with the vPC topology.

SUMMARY STEPS

1. **install feature-set fcoe**
2. **feature-set fcoe**
3. **vsan database**
4. **interface port-channel**
5. **int vfc**
6. **show int brief**
7. **show flogi database**
8. **show vpc statistics**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**install feature-set fcoe**	Install FCoE feature.
Step 2	**feature-set fcoe** **Example:** tme-switch-1(config)# feature-set fcoe Please configure the following for fcoe to be fully functional: - hardware access-list tcam region ing-racl TCAM size - hardware access-list tcam region ing-ifacl TCAM size - hardware access-list tcam region ing-redirect TCAM size	Enable FCoE on the Cisco Nexus 9000 switch. **Note** This can take a few moments to complete. You must ensure to complete the TCAM carving before doing this step. After completing the TCAM carving, you must reload the switch.

	Command or Action	Purpose
	tme-switch-1(config)# tme-switch-2(config)# feature-set fcoe Please configure the following for fcoe to be fully functional: - hardware access-list tcam region ing-racl TCAM size - hardware access-list tcam region ing-ifacl TCAM size - hardware access-list tcam region ing-redirect TCAM size tme-switch-2(config)#	
Step 3	**vsan database** **Example:** tme-switch-1(config)# vsan database tme-switch-1(config-vsan-db)# vsan 100 tme-switch-1(config-vsan-db)# exit tme-switch-1(config)# vlan 100 tme-switch-1(config-vlan)# fcoe vsan 100 tme-switch-1(config-vlan)# show vlan fcoe VLAN VSAN Status -------- -------- -------- 100 100 Operational tme-switch-1(config-vlan)# tme-switch-2(config)# vsan database tme-switch-2(config-vsan-db)# vsan 101 tme-switch-2(config-vsan-db)# exit tme-switch-2(config)# vlan 101 tme-switch-2(config-vlan)# fcoe vsan 101 tme-switch-2(config-vlan)# show vlan fcoe VLAN VSAN Status -------- -------- -------- 101 101 Operational tme-switch-2(config)#	Create a VSAN and map it to a VLAN that has been designated to carry FCoE traffic. **Note** VLAN and VSAN numbers are not required to be the same.
Step 4	**interface port-channel** **Example:** tme-switch-1(config)# interface port-channel 11 tme-switch-1(config-if)# switchport trunk allowed vlan 1, 100 tme-switch-1(config-if)# mtu 9216 tme-switch-1(config-if)# service-policy type qos input default-fcoe-in-policy tme-switch-1(config-if)# show int trunk —————————————————————————— Port Native Status Port —————————————————————————— Eth1/1 1 trnk-bndl Po11 Eth1/39 1 trnk-bndl Po1 Eth1/40 1 trnk-bndl Po1 Po1 1 trunking -- Po11 1 trunking -- —————————————————————————— Port Vlans Allowed on Trunk ——————————————————————————	Configure the VLANs that are allowed to transverse the vPC links.

	Command or Action	Purpose

```
Eth1/1 1,100
Eth1/39 1-3967,4048-4093
Eth1/40 1-3967,4048-4093
Po1 1-3967,4048-4093
Po11 1,100

--------------------------------------------------------------------------------
Port Vlans Err-disabled on Trunk
--------------------------------------------------------------------------------
Eth1/1 none
Eth1/39 100
Eth1/40 100
Po1 100
Po11 none

--------------------------------------------------------------------------------
Port STP Forwarding
--------------------------------------------------------------------------------
Eth1/1 none
Eth1/39 none
Eth1/40 none
Po1 1
Po11 1,100
tme-switch-1(config-if)#

tme-switch-2(config)# int po 11
tme-switch-2(config-if)# switchport trunk allowed
 vlan 1, 101
tme-switch-1(config-if)# mtu 9216
tme-switch-1(config-if)# service-policy type qos
input default-fcoe-in-policy
tme-switch-2(config-if)# show int trunk
--------------------------------------------------------------------------------
Port Native Status Port
--------------------------------------------------------------------------------
Eth1/1 1 trnk-bndl Po11
Eth1/39 1 trnk-bndl Po1
Eth1/40 1 trnk-bndl Po1
Po1 1 trunking --
Po11 1 trunking --

--------------------------------------------------------------------------------
Port Vlans Allowed on Trunk
--------------------------------------------------------------------------------
Eth1/1 1,101
Eth1/39 1-3967,4048-4093
Eth1/40 1-3967,4048-4093
Po1 1-3967,4048-4093
Po11 1,101

--------------------------------------------------------------------------------
Port Vlans Err-disabled on Trunk
--------------------------------------------------------------------------------
Eth1/1 none
Eth1/39 101
Eth1/40 101
Po1 101
Po11 none

--------------------------------------------------------------------------------
Port STP Forwarding
--------------------------------------------------------------------------------
```

	Command or Action	Purpose
	Eth1/1 none Eth1/39 none Eth1/40 none Po1 1 Po11 1,101 tme-switch-2(config-if)#	
Step 5	**int vfc** **Example:** tme-switch-1(config)# int vfc 1 tme-switch-1(config-if)# bind interface po11 tme-switch-1(config-if)# no shut tme-switch-1(config-if)# tme-switch-2(config)# int vfc 1 tme-switch-2(config-if)# bind interface po11 tme-switch-2(config-if)# no shut tme-switch-2(config-if)# tme-switch-1(config)# vsan database tme-switch-1(config-vsan-db)# vsan 100 interface vfc 1 tme-switch-1(config)# show vsan membership vsan 1 interfaces: fc2/1 fc2/2 fc2/3 fc2/4 fc2/5 fc2/6 fc2/7 fc2/8 vsan 100 interfaces: vfc1 vsan 4079(evfp_isolated_vsan) interfaces: vsan 4094(isolated_vsan) interfaces: tme-switch-1(config)# tme-switch-2(config)# vsan database tme-switch-2(config-vsan-db)# vsan 101 interface vfc 1 tme-switch-2(config)# show vsan membership vsan 1 interfaces: fc2/1 fc2/2 fc2/3 fc2/4 fc2/5 fc2/6 fc2/7 fc2/8 vsan 101 interfaces: vfc1 vsan 4079(evfp_isolated_vsan) interfaces: vsan 4094(isolated_vsan) interfaces: tme-switch-2(config)#	Create a virtual Fibre Channel interface (vfc) and add it to the VSAN that was created in the previous step.
Step 6	**show int brief** **Example:** tme-switch-1(config-if)# show int brief -- Ethernet VLAN Type Mode Status Reason Speed --	Verify that the vfc is up and operational:

	Command or Action	Purpose
	Eth1/1 1 eth trunk up none 10G(D) Eth1/2 1 eth access up none 10G(D) Eth1/38 1 eth access down SFP not inserted 10G(D) Eth1/39 1 eth trunk up none 10G(D) Eth1/40 1 eth trunk up none 10G(D) —————————————————————————— Port-channel VLAN Type Mode Status Reason Speed —————————————————————————— Po1 1 eth trunk up none a-10G(D) none Po11 1 eth trunk up none a-10G(D) none —————————————————————————— Port VRF Status IP Address Speed MTU —————————————————————————— mgmt0 -- up 172.25.182.166 1000 1500 —————————————————————————— Interface Vsan Admin Admin Status SFP Oper Oper Port —————————————————————————— vfc1 100 F on up -- F auto -- tme-switch-1(config-if)# tme-switch-2(config-if)# show int brief —————————————————————————— Ethernet VLAN Type Mode Status Reason Speed Port —————————————————————————— Eth1/1 1 eth trunk up none 10G(D) 11 Eth1/2 1 eth access up none 10G(D) -- Eth1/38 1 eth access down SFP not inserted 10G(D) -- Eth1/39 1 eth trunk up none 10G(D) 1 Eth1/40 1 eth trunk up none 10G(D) 1 —————————————————————————— Port-channel VLAN Type Mode Status Reason Speed Protocol —————————————————————————— Po1 1 eth trunk up none a-10G(D) none Po11 1 eth trunk up none a-10G(D) none —————————————————————————— Port VRF Status IP Address Speed MTU —————————————————————————— mgmt0 -- up 172.25.182.167 1000 1500 —————————————————————————— Interface Vsan Admin Admin Status SFP Oper Oper —————————————————————————— vfc1 101 F on up -- F auto -- tme-switch-2(config-if)#	
Step 7	**show flogi database** **Example:** tme-switch-1# show flogi database —————————————————————————— INTERFACE VSAN FCID PORT NAME NODE NAME —————————————————————————— vfc1 100 0x540000 21:00:00:c0:dd:11:2a:01	Verify that the virtual Fibre Channel interface has logged into the fabric.

	Command or Action	Purpose
	20:00:00:c0:dd:11:2a:01 Total number of flogi = 1. tme-switch-2# show flogi database -- INTERFACE VSAN FCID PORT NAME NODE NAME -- vfc1 101 0x540000 21:00:00:c0:dd:11:2a:01 20:00:00:c0:dd:11:2a:01 Total number of flogi = 1.	
Step 8	**show vpc statistics** **Example:** tme-switch-1(config-if)# show vpc statistics vpc 11 port-channel11 is up vPC Status: Up, vPC number: 11 Hardware: Port-Channel, address: 000d.ecde.a908 (bia 000d.ecde.a908) MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec, reliability 255/255, txload 1/255, rxload 1/255 Encapsulation ARPA Port mode is trunk full-duplex, 10 Gb/s Beacon is turned off Input flow-control is off, output flow-control is off Switchport monitor is off Members in this channel: Eth1/1 Last clearing of "show interface" counters never 1 minute input rate 4968 bits/sec, 8 packets/sec 1 minute output rate 792 bits/sec, 1 packets/sec tme-switch-2(config-if)# show vpc statistics vpc 11 port-channel11 is up vPC Status: Up, vPC number: 11 Hardware: Port-Channel, address: 000d.ecdf.5fae (bia 000d.ecdf.5fae) MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec, reliability 255/255, txload 1/255, rxload 1/255 Encapsulation ARPA Port mode is trunk full-duplex, 10 Gb/s Beacon is turned off Input flow-control is off, output flow-control is off Switchport monitor is off Members in this channel: Eth1/1 Last clearing of "show interface" counters never 1 minute input rate 4968 bits/sec, 8 packets/sec 1 minute output rate 792 bits/sec, 1 packets/sec	Verify that the vPC is up and operational.

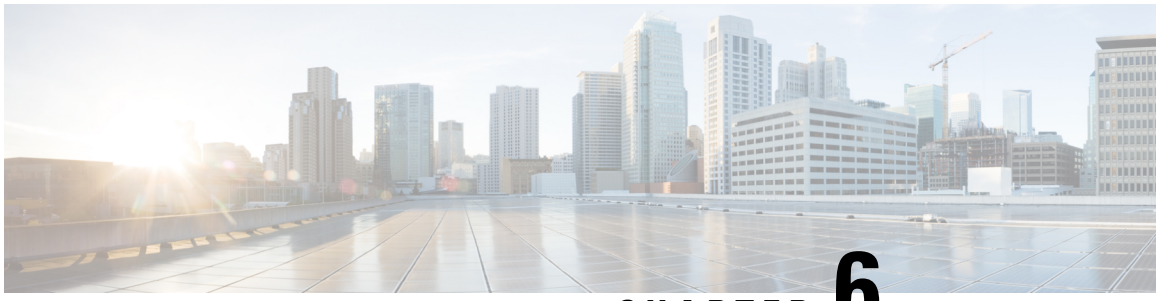

CHAPTER 6

Configuring Fibre Channel Interfaces

This chapter contains the following sections:

- Information About Fibre Channel Interfaces, on page 37
- Configuring Fibre Channel Interfaces, on page 48
- Configuring Global Attributes for Fibre Channel Interfaces, on page 56
- Verifying Fibre Channel Interfaces, on page 58
- Default Fibre Channel Interface Settings, on page 61

Information About Fibre Channel Interfaces

Virtual Fibre Channel Interfaces

Fibre Channel over Ethernet (FCoE) encapsulation allows a physical Ethernet cable to simultaneously carry Fibre Channel and Ethernet traffic. In Cisco Nexus devices, an FCoE-capable physical Ethernet interface can carry traffic for one virtual Fibre Channel (vFC) interface.

Like any interface in Cisco NX-OS, vFC interfaces are manipulable objects with properties such as configuration and state. Native Fibre Channel and vFC interfaces are configured using the same CLI commands.

Note vFC interfaces support only F mode and operate in trunk mode only.

The following capabilities are not supported for virtual Fibre Channel interfaces:

- SAN port channels.
- The SPAN destination cannot be a vFC interface.
- Buffer-to-buffer credits.
- Exchange link parameters (ELP), or Fabric Shortest Path First (FSPF) protocol.
- Configuration of physical attributes (speed, rate, mode, transmitter information, MTU size).
- Port tracking.

VF Port

vFC interfaces always operate in trunk mode; vFC interfaces do not operate in any other mode. You can configure allowed VSANs on a vFC by using the **switchport trunk allowed vsan** command under the vfc interface (which is similar to FC TF and TE ports). For vFC interfaces that are connected to hosts, port VSAN is the only VSAN that supports logins (FLOGI). We recommend that you restrict the allowed VSANs for such vFC interfaces to the port VSAN by using the **switchport trunk allowed vsan** command in the interface mode to configure a VF port.

Includes support for 160 vFC interfaces.

The vFC VSAN assignment and the global VLAN-to-VSAN mapping table enables the Cisco Nexus device to choose the appropriate VLAN for a VF port.

Interface Modes

Each physical Fibre Channel interface in a switch may operate in one of several port modes: E mode, TE mode, F mode, and TF mode. A physical Fibre Channel interface can be configured as an E port or an F port, an F port, or an SD port. Interfaces may also be configured in Auto mode; the port type is determined during interface initialization.

Fibre Channel interfaces may operate in E mode or an F mode.

Virtual Fibre Channel interfaces can only be configured in F mode.

Interfaces are automatically assigned VSAN 1 by default.

Each interface has an associated administrative configuration and an operational status:

- The administrative configuration does not change unless you modify it. This configuration has various attributes that you can configure in administrative mode.
- The operational status represents the current status of a specified attribute such as the interface speed. This status cannot be changed and is read-only. Some values may not be valid when the interface is down (for example, the operational speed).

E Port

In expansion port (E port) mode, an interface functions as a fabric expansion port. This port may be connected to another E port to create an Inter-Switch Link (ISL) between two switches. E ports carry frames between switches for configuration and fabric management. They serve as a conduit between switches for frames destined to remote N ports. E ports support class 3 and class F service.

An E port connected to another switch may also be configured to form a SAN port channel.

F Port

In fabric port (F port) mode, an interface functions as a fabric port. This port may be connected to a peripheral device (host or disk) operating as a node port (N port). An F port can be attached to only one N port. F ports support class 3 service.

TE Port

In trunking E port (TE port) mode, an interface functions as a trunking expansion port. It may be connected to another TE port to create an extended ISL (EISL) between two switches. TE ports connect to another Cisco

Nexus device or a Cisco MDS 9000 Family switch. They expand the functionality of E ports to support the following:

- VSAN trunking
- Fibre Channel trace (fctrace) feature

In TE port mode, all frames are transmitted in EISL frame format, which contains VSAN information. Interconnected switches use the VSAN ID to multiplex traffic from one or more VSANs across the same physical link. This feature is referred to as VSAN trunking in the Cisco Nexus device. TE ports support class 3 and class F service.

TF Port

When the switch is operating in NPV mode, the interfaces that connect the switch to the core network switch are configured as NP ports. NP ports operate like N ports that function as proxies for multiple physical N ports.

In trunking F port (TF port) mode, an interface functions as a trunking expansion port. It may be connected to another trunked N port (TN port) or trunked NP port (TNP port) to create a link between a core switch and an NPV switch or an HBA to carry tagged frames. TF ports expand the functionality of F ports to support VSAN trunking.

In TF port mode, all frames are transmitted in an EISL frame format, which contains VSAN information. Interconnected switches use the VSAN ID to multiplex traffic from one or more VSANs across the same physical link. This feature is referred to as VSAN trunking in Cisco Nexus devices. TF ports support class 3 and class F service.

Auto Mode

Interfaces configured in auto mode can operate in one of the following modes: E, F, TE, and TF. The port mode is determined during interface initialization. For example, if the interface is connected to a node (host or disk), it operates in F port mode. If the interface is attached to a third-party switch, it operates in E port mode. If the interface is attached to another switch in the Cisco Nexus device or Cisco MDS 9000 Family, it may become operational in TE port mode.

Interface States

The interface state depends on the administrative configuration of the interface and the dynamic state of the physical link.

Administrative States

The administrative state refers to the administrative configuration of the interface. The table below describes the administrative states.

Table 3: Administrative States

Administrative State	Description
Up	Interface is enabled.

Administrative State	Description
Down	Interface is disabled. If you administratively disable an interface by shutting down that interface, the physical link layer state change is ignored.

Operational States

The operational state indicates the current operational state of the interface. The table below describes the operational states.

Table 4: Operational States

Operational State	Description
Up	Interface is transmitting or receiving traffic as desired. To be in this state, an interface must be administratively up, the interface link layer state must be up, and the interface initialization must be completed.
Down	Interface cannot transmit or receive (data) traffic.
Trunking	Interface is operational in TE or TF mode.

Reason Codes

Reason codes are dependent on the operational state of the interface. The following table describes the reason codes for operational states.

Table 5: Reason Codes for Interface States

Administrative Configuration	Operational Status	Reason Code
Up	Up	None.
Down	Down	Administratively down. If you administratively configure an interface as down, you disable the interface. No traffic is received or transmitted.
Up	Down	See the table below.

If the administrative state is up and the operational state is down, the reason code differs based on the nonoperational reason code. The table below describes the reason codes for nonoperational states.

Note Only some of the reason codes are listed in the table.

Table 6: Reason Codes for Nonoperational States

Reason Code (long version)	Description	Applicable Modes
Link failure or not connected	The physical layer link is not operational.	All
SFP not present	The small form-factor pluggable (SFP) hardware is not plugged in.	All
Initializing	The physical layer link is operational and the protocol initialization is in progress.	All
Reconfigure fabric in progress	The fabric is currently being reconfigured.	
Offline	The switch software waits for the specified R_A_TOV time before retrying initialization.	
Inactive	The interface VSAN is deleted or is in a suspended state. To make the interface operational, assign that port to a configured and active VSAN.	
Hardware failure	A hardware failure is detected.	
Error disabled	Error conditions require administrative attention. Interfaces may be error-disabled for various reasons. For example: • Configuration failure. • Incompatible buffer-to-buffer credit configuration. To make the interface operational, you must first fix the error conditions causing this state and then administratively shut down andor enable the interface.	
Isolation because limit of active port channels is exceeded.	The interface is isolated because the switch is already configured with the maximum number of active SAN port channels.	
Isolation due to ELP failure	The port negotiation failed.	Only E ports and TE ports
Isolation due to ESC failure	The port negotiation failed.	
Isolation due to domain overlap	The Fibre Channel domains (fcdomain) overlap.	
Isolation due to domain ID assignment failure	The assigned domain ID is not valid.	
Isolation due to the other side of the link E port isolated	The E port at the other end of the link is isolated.	

Reason Code (long version)	Description	Applicable Modes
Isolation due to invalid fabric reconfiguration	The port is isolated due to fabric reconfiguration.	
Isolation due to domain manager disabled	The fcdomain feature is disabled.	
Isolation due to zone merge failure	The zone merge operation failed.	
Isolation due to VSAN mismatch	The VSANs at both ends of an ISL are different.	
port channel administratively down	The interfaces belonging to the SAN port channel are down.	Only SAN port channel interfaces
Suspended due to incompatible speed	The interfaces belonging to the SAN port channel have incompatible speeds.	
Suspended due to incompatible mode	The interfaces belonging to the SAN port channel have incompatible modes.	
Suspended due to incompatible remote switch WWN	An improper connection is detected. All interfaces in a SAN port channel must be connected to the same switch.pair of switches.	
Bound physical interface down	The Ethernet interface bound to a virtual Fibre Channel interface is not operational.	Only virtual Fibre Channel interfaces
STP not forwarding in FCoE mapped VLAN	The Ethernet interface bound to a virtual Fibre Channel interface is not in an STP forwarding state for the VLAN associated with the virtual Fibre Channel interface	Only virtual Fibre Channel interfaces

Buffer-to-Buffer Credits

Buffer-to-buffer credits (BB_credits) are a flow-control mechanism to ensure that Fibre Channel interfaces do not drop frames. BB_credits are negotiated on a per-hop basis.

The BB_credit mechanism is used on Fibre Channel interfaces but not on virtual Fibre Channel interfaces. The receive BB_credit determines the receive buffering capability on the receive side without having to acknowledge the peer. This is important for links with large bandwidth-delays (long links with large latency) to be able to sustain line-rate traffic with increased latency.

For virtual Fibre Channel interfaces, BB_credits are not used. Virtual Fibre Channel interfaces provide flow control based on a class based pause mechanism named Priority Flow Control.Priority-Flow_control.

Note

- Buffer-to-buffer (B2B) credits are not configurable.
- Fill pattern in the 8G links must be IDLE. You must set the fill pattern in the 8G links to IDLE on both the peers. Use the command **switchport fill-pattern IDLE speed** *speed* to set the fill pattern to IDLE on Cisco Nexus 9000 switches.

```
switch (config)# interface fc1/1
switch (config-if)# switchport fill-pattern IDLE speed 8000
```

Note The receive BB_credit value is 64 in N9K-C93180YC-FX and 32 in N9K-C93360YC-FX2. This is applicable for all port modes (F,E) in both platforms and cannot be changed.

Licensing Requirements for Fibre Channel

Ensure that you have the correct license installed before using Fibre Channel interfaces and capabilities. For more information on licensing, see *Enabling FC/FCoE* chapter in this guide.

Enabling the Fibre Channel Port License

This section explains how to enable the licensing for SAN Switching.

Before you begin

To enable the port license, you must shut down the fibre channel (FC) ports.

Note For information about converting to FC ports, see Configuring Unified Ports.

SUMMARY STEPS

1. Enable the port license.

DETAILED STEPS

Enable the port license.

Example:

```
Switch(config)# int fc1/1
Switch(config-if)# port-license acquire
```

Configuring QoS for no-drop Support

A qos ingress policy is used to mark ingress FC/FCoE frames. The qos ingress policy must be applied to the interfaces that handle FC/FCoE traffic (such as, all ethernet/port-channel interfaces bound to vFCs).

Note Check to ensure that the port qos region has hardware TCAM space reserved. Whenever an ingress PACL TCAM threshold is seen in the syslog, increase the TCAM size and reload the switch.

This step is mandatory for FC/FCoE to work.

- Reserve TCAM space for the QoS region.

  You may need to acquire TCAM space reserved for other regions (such as the l3qos region).

- Save the configuration.
- Reload the line cards or switch.

  Reload the switch.

- Confirm the port qos region TCAM space.
- Example for TCAM carving on N9K-C93180YC-FX, :

```
hardware access-list tcam region ing-racl 1536
hardware access-list tcam region ing-ifacl 256
hardware access-list tcam region ing-redirect 256
```

Example:

```
switch# show hardware access-list tcam region |i i ifacl
Ingress PACL [ing-ifacl] size =  256
switch# config


switch(config)# hardware access-list tcam region ing-racl 1536
switch(config)# hardware access-list tcam region ing-ifacl 256
switch(config)# hardware access-list tcam region ing-redirect 256



switch# copy running-config startup-config
switch# reload

switch# show hardware access-list tcam region |i i ifacl
Ingress PACL [ing-ifacl] size =  256
```

Configuring QoS Policies for FC/FCoE

- There are four types of FC/FCoE default policies: network-qos, output queuing, input queuing, and input qos.
- To use a different queue or cos value for FC/FCoE traffic, create user-defined policies.
- You can configure a QoS policy by following one of these methods:

- Predefined policies—You can apply a predefined QoS policy: **default-fcoe-in-policy**.

Note
- No policy will be applied by default for FCoE.
- We recommend to apply **no-stats** to QoS policy.

- User-defined policy—You can create a QoS policy that conforms to one of the system-defined policies.

Configuring System-wide QoS Policy

Note The network-qos policy and output/input queuing policies should be applied at the system level and the qos policy should be applied at the interface level, for every interface that carries the FC/FCoE traffic.

```
switch(config)# system qos
switch(config-sys-qos)# service-policy type queuing input default-fcoe-in-que-policy
switch(config-sys-qos)# service-policy type queuing output { default-fcoe-8q-out-policy |
default-fcoe-out-policy }
switch(config-sys-qos)# service-policy type network-qos { default-fcoe-8q-nq-policy |
default-fcoe-nq-policy }
```

Configuration Example for user-defined policies

```
switch(config)# policy-map type network-qos fcoe_nq
switch(config-pmap-nqos)# class type network-qos c-nq1
switch(config-pmap-nqos-c)# pause pfc-cos 3
switch(config-pmap-nqos-c)# mtu 9216
switch(config-pmap-nqos-c)# class type network-qos c-nq2
switch(config-pmap-nqos-c)# mtu 1500
switch(config-pmap-nqos-c)# class type network-qos c-nq3
switch(config-pmap-nqos-c)# mtu 1500
switch(config-pmap-nqos-c)# class type network-qos c-nq-default
switch(config-pmap-nqos-c)# mtu 1500
switch(config-pmap-nqos-c)# exit
switch(config-pmap-nqos)# exit
switch(config)#
switch(config)# policy-map type queuing fcoe-in-policy
switch(config-pmap-que)# class type queuing c-in-q1
switch(config-pmap-c-que)# bandwidth percent 50
switch(config-pmap-c-que)# class type queuing c-in-q-default
switch(config-pmap-c-que)# bandwidth percent 50
switch(config-pmap-c-que)# exit
switch(config)
switch(config)# policy-map type queuing fcoe-out-policy
switch(config-pmap-que)# class type queuing c-out-q3
switch(config-pmap-c-que)# priority level 1
switch(config-pmap-c-que)# class type queuing c-out-q-default
switch(config-pmap-c-que)# bandwidth remaining percent 50
switch(config-pmap-c-que)# class type queuing c-out-q1
switch(config-pmap-c-que)# bandwidth remaining percent 50
switch(config-pmap-c-que)# class type queuing c-out-q2
switch(config-pmap-c-que)# bandwidth remaining percent 0
switch(config-pmap-c-que)# exit
```

```
switch(config)#
switch(config)# class-map type qos match-any fcoe
switch(config-cmap-qos)# match protocol fcoe
switch(config-cmap-qos)# match cos 3
switch(config-cmap-qos)# exit
switch(config)#
switch(config)# policy-map type qos fcoe_qos_policy
switch(config-pmap-qos)# class fcoe
switch(config-pmap-c-qos)# set cos 3
switch(config-pmap-c-qos)# set qos-group 1
switch(config-pmap-c-qos)# exit
switch(config-pmap-qos)# exit
switch(config)#
switch(config)# system qos
switch(config-sys-qos)# service-policy type queuing input fcoe-in-policy
switch(config-sys-qos)# service-policy type queuing output fcoe-out-policy
switch(config-sys-qos)# service-policy type network-qos fcoe_nq
```

Note The **set cos 3** command under the QOS policy is mandatory only when there are native fiber channel ports and the command is applicable only for N9K-C93180YC-FX platform , N9K-C93360YC-FX2 platforms . For all the other Cisco Nexus 9000 Platform switches, this step is optional.

Applying the ingress QoS policy to each Ethernet/port-channel interface that is bound to vFC interface for FC/FCoE.

```
switch(config)# interface ethernet 2/1
switch(config-if)# switchport mode trunk
switch(config-if)# mtu 9216 /* Or maximum allowed value */
switch(config-if)# service-policy type qos input { default-fcoe-in-policy | fcoe_qos_policy
 )  no-stats
switch(config-if)# exit
switch(config)#
```

- Configuring FC/FCoE QoS policies
  - There are four types of FC/FCoE default policies: network QoS, output queuing, input queuing, and QoS.
  - To use a different queue or cos value for FC/FCoE traffic, create user-defined policies.
- Configuring Network QoS Policies for FC/FCoE
  - You can configure a network QoS policy by following one of these methods:
    - Predefined policies—You can apply a predefined network QoS policy that fits your requirement. You have the option to choose either **default-fcoe-8q-nq-policy** or **default-fcoe-nq-policy**.

Note No policy will be applied by default for FC/FCoE.

    - User-defined policy—You can create a network QoS policy that conforms to one of the system-defined policies.
- Configuring Output Queuing Policies for FC/FCoE

- You can configure an output queuing policy by following one of these methods:
  - Predefined policies—You can apply a predefined output queuing policy that fits your requirement. You have the option to choose either **default-fcoe-8q-out-policy** or **default-fcoe-out-policy**.

Note No policy will be applied by default for FC/FCoE.

  - User-defined policy—You can create a output queuing policy that conforms to one of the system-defined policies.

- Configuring Input Queuing Policies for FC/FCoE
  - You can configure an input queuing policy by following one of these methods:
    - Predefined policies—You can apply a predefined input queuing policy: **default-fcoe-in-que-policy**.

Note No policy will be applied by default for FCoE.

    - User-defined policy—You can create a input queuing policy that conforms to one of the system-defined policies.

Note Whenever you see label allocation failure in the syslog, there is a possibility of FC/FCoE ACL not getting applied on interfaces. You must then check whether the QoS policy is applied with no-stats on the interfaces.

Physical Fibre Channel Interfaces

Cisco Nexus C93180YC-FX switches have up to 48 physical fibre channel (FC) interfaces as either uplinks connected to SAN network or as downlinks (connected to server or target).

Cisco Nexus C93180YC-FX and C93360YC-FX2 switches support up to 48 and 96 physical fibre channel (FC) interfaces respectively as either uplinks connected to SAN network or as downlinks (connected to server or target).

Each Fibre Channel port can be used as a downlink (connected to a server) or as an uplink (connected to the data center SAN network). The Fibre Channel interfaces support the following modes: E, F, TE, and TF.

Long-Distance ISLs

Beginning with Cisco NX-OS Release 9.3(5), the Cisco Nexus N9K-C93180YC-FX switch supports long distance on 32-Gbps Fibre Channel Inter-Switch Link (ISL).

When long-distance ISLs are required, you must have sufficient BB_credits configured to ensure that the ISL can run at maximum capacity. The formula for computing long-distance ISL BB_credits assumes a typical

Fibre Channel frame of 2 KB and factors in the interface speed. With fixed (64) buffer-to-buffer credits, the new switch now provide support for 32-Gbps Fibre Channel ISLs across distances of up to 3 kilometers.

Table 7: FC Long Distance across different speeds

Speed	Distance
32G	3 KM
16G	5 KM
8G	10 KM

Configuring Fibre Channel Interfaces

Configuring a Fibre Channel Interface

To configure a Fibre Channel interface, perform this task:

Note For information about creating FC ports or port conversion, see the Configuring Unified Ports section.

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface** {**fc** *slot*/*port*}|{**vfc** *vfc-id*}

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **interface** {**fc** *slot*/*port*}|{**vfc** *vfc-id*}	Selects a Fibre Channel interface and enters interface configuration mode. **Note** When a Fibre Channel interface is configured, it is automatically assigned a unique world wide name (WWN). If the interface's operational state is up, it is also assigned a Fibre Channel ID (FC ID).

Configuring a Range of Fibre Channel Interfaces

To configure a range of Fibre Channel interfaces, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface** { **fc** *slot*/*port* - *port* [, **fc** *slot*/*port* - *port*] | **vfc** *vfc-id* - *vfc-id* [, **vfc** *vfc-id* - *vfc-id*] }

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **interface** { **fc** *slot*/*port* - *port* [, **fc** *slot*/*port* - *port*] | **vfc** *vfc-id* - *vfc-id* [, **vfc** *vfc-id* - *vfc-id*] }	Selects the range of Fibre Channel interfaces and enters interface configuration mode.

Setting the Interface Administrative State

To gracefully shut down an interface, perform this task:

To enabledisable traffic flow, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface** {**fc** *slot*/*port*}|{**vfc** *vfc-id*}
3. switch(config-if)# **shutdown**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **interface** {**fc** *slot*/*port*}|{**vfc** *vfc-id*}	Selects a Fibre Channel interface and enters interface configuration mode.
Step 3	switch(config-if)# **shutdown**	Gracefully shuts down the interface and administratively disables traffic flow (default).

Configuring Interface Modes

SUMMARY STEPS

1. **configure terminal**
2. switch(config) # **interface vfc** *vfc-id*}
3. switch(config-if) # **switchport mode** {**F**}

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	switch(config) # **interface vfc** *vfc-id*}
	Example:
	`switch(config) # interface vfc 20`
	`switch(config-if) #`
		Selects a virtual Fibre Channel interface and enters interface configuration mode.

Step 3	switch(config-if) # **switchport mode** {**F**}
	Example:
	`switch(config-if) # switchport mode F`
	`switch(config-if) #`
		Sets the port mode.
		vFC interfaces support only F mode.

Example

This example shows the running configuration for vFC 20 bound to the Ethernet slot1,port 3 interface.

```
switch# show running-config
switch(config) # interface vfc20
switch(config-if) # bind interface Ethernet 1/3
switch(config-if) # switchport mode F
switch(config-if) # no shutdown
```

Configuring the Interface Description

Interface descriptions should help you identify the traffic or use for that interface. The interface description can be any alphanumeric string.

To configure a description for an interface, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface** {**fc** *slot*/*port*}|{**vfc** *vfc-id*}
3. switch(config-if)# **switchport description cisco-HBA2**
4. switch(config-if)# **no switchport description**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.

	Command or Action	Purpose
Step 2	switch(config)# **interface** {**fc** *slot*/*port*}|{**vfc** *vfc-id*}	Selects a Fibre Channel interface and enters interface configuration mode.
Step 3	switch(config-if)# **switchport description cisco-HBA2**	Configures the description of the interface. The string can be up to 80 characters long.
Step 4	switch(config-if)# **no switchport description**	Clears the description of the interface.

Configuring Unified Ports

Before you begin

Confirm that you have a supported Cisco Nexus switch. Unified Ports are available on the Cisco Nexus C93180YC-FX switch and C93360YC-FX2 switches.

Note For information about the C93180YC-FX or C93360YC-FX2 platform details, see the *Cisco Nexus 9000 Series Hardware Installation Guide*.

If you're configuring a unified port as Fibre Channel or FCoE, confirm that you have enabled the **install feature-set fcoe** and **feature-set fcoe** commands.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config) # **slot** *slot number*	Identifies the slot on the switch.
Step 3	switch(config-slot) # **port** *port number* **type** {**ethernet** | **fc**}	Configures a unified port as a native Fibre Channel port and an Ethernet port. • **type** —Specifies the type of port to configure on a slot in a chassis. • **ethernet** —Specifies an Ethernet port. • **fc** —Specifies a Fibre Channel (FC) port.

	Command or Action	Purpose
		Note • When you configure unified ports as Fibre Channel, the existing configuration for Fibre Channel interfaces and VSAN memberships are unaffected. • In N9K-C93180YC-FX switches, the FC port range must be in multiples of 4, and can be discontinuous also. Reload the switch for the change to take effect. • In N9K-C93360YC-FX2 switches, you must convert all the four front panel ports in a column to FC/Ethernet together. In this switch, four ports form a port group. For example, the first port group is 1,2,49,50; the second port group is 3,4,51,52 and so on.
Step 4	switch(config-slot)# **copy running-config startup-config**	Copies the running configuration to the startup configuration.
Step 5	switch(config-slot) # **reload**	Reboots the switch.
Step 6	switch(config) # **slot** *slot number*	Identifies the slot on the switch.
Step 7	switch(config-slot) # **no port** *port number* **type fc**	Change the port back as an ethernet port, after you perform copy r s and reload the switch.

Example

Note Individual ports cannot be converted to FC ports.

```
switch# configure terminal
switch(config)# slot 1
switch(config-slot)# port 1-24 type fc
Port type is changed. ACTION REQUIRED: Please save configurations and reload the switch
switch(config-slot)#
```

Configuring Port Speeds

Port speed can be configured on a physical Fibre Channel interface but not on a virtual Fibre Channel interface. The minimum supported speed is 4G and the maximum is 32G. By default, the port speed for an interface is automatically calculated by the switch.

Note 8G speed is not supported for server and target interfaces.

Caution Changing the interface speed is a disruptive operation.

To configure the port speed of the interface, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface fc** *slot*/*port*
3. switch(config-if)# **switchport speed 16000**
4. switch(config-if)# **no switchport speed**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **interface fc** *slot*/*port*	Selects the specified interface and enters interface configuration mode. **Note** You cannot configure the port speed of a virtual Fibre Channel interface.
Step 3	switch(config-if)# **switchport speed 16000**	Configures the port speed of the interface to 16G. The number indicates the speed in megabits per second (Mbps). You can set the speed to 4000 (for 4-Gbps interfaces), 8000 (for 8-Gbps interfaces), 16000 (for 16-Gbps interfaces), 32000 (for 32-Gbps interfaces), or auto (default). **Note** When you connect a 16G host adapter to a 32G SFP port on a Cisco Nexus 9000 switch, if the link does not come up when the speed is configured as auto speed or if it defaults to 8G speed, then, you must manually configure the port using the command **switchport speed 16000**.
Step 4	switch(config-if)# **no switchport speed**	Reverts to the factory default (auto) administrative speed of the interface.

Configuring Trunk Mode

To configure trunk mode perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface fc** *slot*/*port*
3. switch(config-if)# **switchport trunk mode on**
4. switch(config-if)# **switchport trunk mode off**
5. switch(config-if)# **switchport trunk mode auto**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **interface fc** *slot*/*port*	Configures the specified interface and enters interface configuration mode.
Step 3	switch(config-if)# **switchport trunk mode on**	Enables (default) the trunk mode for the specified interface.
Step 4	switch(config-if)# **switchport trunk mode off**	Disables the trunk mode for the specified interface.
Step 5	switch(config-if)# **switchport trunk mode auto**	Configures the trunk mode to auto mode, which provides automatic sensing for the interface.

Note

For FC ports with trunking mode on and SAN-PO links to come up between two switches, both switches should be configured with the OUI of each other.

Configure the OUI on the switches only if the OUI value is not registered by default on either of them. The OUI is found and configured as follows:

```
  N9K(config-if)# show wwn switch
    Switch WWN is 20:00:2c:d0:2d:50:ea:64
    N9K(config-if)#
```

On the switch, you can see the below output if the OUI (0x2cd02d) is already registered.

```
    MDS9710(config-if)# sh wwn oui | i 2cd02d
    0x2cd02d   Cisco               Default
    MDS9710(config-if) #
If the OUI is not registered, configure it manually.
 MDS9710(config-if)# wwn oui 0x2cd02d
```

Beginning with Cisco NX-OS Release 7.3(0)D1(1), the OUI is configurable on a Cisco MDS 9700 Series core switches.

Autosensing

Autosensing is enabled on all interfaces irrespective of the speed. If an 8G Small Form-Factor Pluggable (SFP) is inserted, the interface operates at 8G and 4G speed. If a 16G SFP is inserted, the interface operates only at 16G, 8G and 4G speeds and with a 32G SFP, the interface operates at 32G, 16G, and 8G speeds.

Understanding Bit Error Thresholds

The bit error rate threshold is used by the switch to detect an increased error rate before performance degradation seriously affects traffic.

The bit errors can occur for the following reasons:

- Faulty or bad cable.
- Faulty or bad GBIC or SFP.
- Short haul cable is used for long haul or long haul cable is used for short haul.
- Momentary synchronization loss.
- Loose cable connection at one or both ends.
- Improper GBIC or SFP connection at one or both ends.

A bit error rate threshold is detected when 15 error bursts occur in a 5-minute period. By default, the switch disables the interface when the threshold is reached.

You can enter the **shutdown**/**no shutdown** command sequence to reenable the interface.

You can configure the switch to not disable an interface when the threshold is crossed.

Note The switch generates a syslog message when bit error threshold events are detected, even if the interface is configured not to be disabled by bit-error threshold events.

To disable the bit error threshold for an interface, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface fc** *slot*/*port*
3. switch(config-if)# **switchport ignore bit-errors**
4. switch(config-if)# **no switchport ignore bit-errors**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **interface fc** *slot*/*port*	Selects a Fibre Channel interface and enters interface configuration mode.
Step 3	switch(config-if)# **switchport ignore bit-errors**	Prevents the detection of bit error threshold events from disabling the interface.
Step 4	switch(config-if)# **no switchport ignore bit-errors**	Prevents the detection of bit error threshold events from enabling the interface.

Configuring Global Attributes for Fibre Channel Interfaces

Configuring Switch Port Attribute Default Values

You can configure attribute default values for various switch port attributes. These attributes will be applied globally to all future switch port configurations, even if you do not individually specify them at that time.

To configure switch port attributes, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **no system default switchport shutdown san**
3. switch(config)# **system default switchport shutdown san**
4. switch(config)# **system default switchport trunk mode auto**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **no system default switchport shutdown san**	Configures the default setting for administrative state of an interface as Up. (The factory default setting is Down). **Tip** This command is applicable only to interfaces for which no user configuration exists for the administrative state.
Step 3	switch(config)# **system default switchport shutdown san**	Configures the default setting for administrative state of an interface as Down. This is the factory default setting. **Tip** This command is applicable only to interfaces for which no user configuration exists for the administrative state.
Step 4	switch(config)# **system default switchport trunk mode auto**	Configures the default setting for administrative trunk mode state of an interface as Auto. **Note** The default setting is trunk mode on.

Information About N Port Identifier Virtualization

N port identifier virtualization (NPIV) provides a means to assign multiple FC IDs to a single N port. This feature allows multiple applications on the N port to use different identifiers and allows access control, zoning, and port security to be implemented at the application level. The following figure shows an example application using NPIV.

Figure 6: NPIV Example

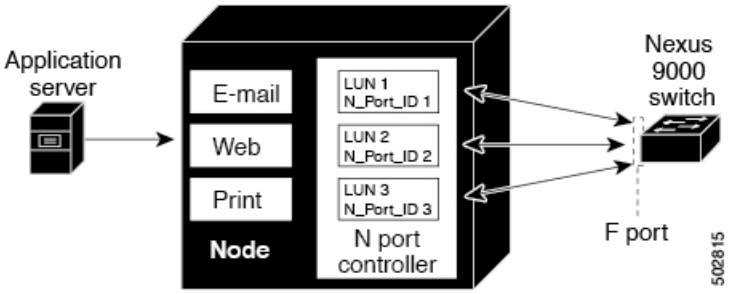

Enabling N Port Identifier Virtualization

You can enable or disable NPIV on the switch. Feature NPIV will be enabled by default when **feature-set fcoe** is enabled.

Before you begin

You must globally enable NPIV for all VSANs on the switch to allow the NPIV-enabled applications to use multiple N port identifiers.

Note All of the N port identifiers are allocated in the same VSAN.

SUMMARY STEPS

1. **configure terminal**
2. **feature npiv**
3. **no feature npiv**

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters configuration mode.

Step 2	**feature npiv** **Example:** `switch(config)# feature npiv`	Enables NPIV for all VSANs on the switch.
Step 3	**no feature npiv** **Example:** `switch(config)# no feature npiv`	Disables (default) NPIV on the switch.

Example Port Channel Configurations

This section shows examples on how to configure an F port channel in shared mode and how to bring up the link between F ports on the NPIV core switches and NP ports on the NPV switches. Before you configure the F port channel, ensure that F port trunking, F port channeling, and NPIV are enabled.

Example

This example shows how to create the port channel:

```
switch(config)# interface san-po-channel 2
switch(config-if)# switchport mode F
switch(config-if)# channel mode active
switch(config-if)# exit
```

This example shows how to configure the port channel member interfaces on the core switch:

```
switch(config)# interface fc1/4-6
switch(config-if)# shut
switch(config-if)# switchport mode F
switch(config-if)# switchport speed 32000
switch(config-if)# switchport trunk mode on
switch(config-if)# channel-group 2
switch(config-if)# no shut
switch(config-if)# exit
```

Verifying Fibre Channel Interfaces

Verifying SFP Transmitter Types

The SFP transmitter type can be displayed for a physical Fibre Channel interface (but not for a virtual Fibre Channel).

The small form-factor pluggable (SFP) hardware transmitters are identified by their acronyms when displayed in the **show interface brief** command. If the related SFP has a Cisco-assigned extended ID, then the **show interface** and **show interface brief** commands display the ID instead of the transmitter type. The **show interface transceiver** command and the **show interface fc** *slot*/*port* transceiver command display both values for Cisco supported SFPs.

Verifying Interface Information

The **show interface** command displays interface configurations. If no arguments are provided, this command displays the information for all the configured interfaces in the switch.

You can also specify arguments (a range of interfaces or multiple, specified interfaces) to display interface information. You can specify a range of interfaces by entering a command with the following example format: interface fc2/1 - 4 , fc3/2 - 3

The following example shows how to display all interfaces:

```
switch# show interface

fc3/1 is up
...
fc3/3 is up
...
Ethernet1/3 is up
...
mgmt0 is up
...
vethernet1/1 is up
...
vfc 1 is up
```

The following example shows how to display multiple specified interfaces:

```
switch# show interface fc3/1 , fc3/3
fc3/1 is up
...
fc3/3 is up
...
```

The following example shows how to display a specific interface:

```
switch# show interface vfc 1

vfc 1 is up

...
```

The following example shows how to display interface descriptions:

```
switch# show interface description
-------------------------------------------------
Interface          Description
-------------------------------------------------
fc3/1              test intest
Ethernet1/1              --
vfc 1              --
...
```

The following example shows how to display all interfaces in brief:

```
switch# show interface brief
```

The following example shows how to display interface counters:

```
switch# show interface counters
```

The following example shows how to display transceiver information for a specific interface:

```
switch# show interface fc3/1 transceiver
```

Note The **show interface transceiver** command is only valid if the SFP is present.

The **show running-configurationshow running-config** command displays the entire running configuration with information for all interfaces. The interfaces have multiple entries in the configuration files to ensure that the interface configuration commands execute in the correct order when the switch reloads. If you display

the running configuration for a specific interface, all the configuration commands for that interface are grouped together.

The following example shows the interface display when showing the running configuration for all interfaces:

```
switch# show running configurationshow running-config
...
interface fc3/5
  switchport speed 200016000
...
interface fc3/5
  switchport mode E
...
interface fc3/5
  channel-group 11 force
  no shutdown
```

The following example shows the interface display when showing the running configuration for a specific interface:

```
switch# show running configuration fc3/5show running-config fc3/5
interface fc3/5
  switchport speed 200016000
  switchport mode E
  channel-group 11 force
  no shutdown
```

Verifying BB_Credit Information

The following example shows how to display the BB_credit information for all Fibre Channel interfaces:

```
switch# show interface fc1/7
...
fc1/7 is up
Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
Port WWN is 20:07:2c:d0:2d:50:e5:24
Admin port mode is auto, trunk mode is off
snmp link state traps are enabled
Port mode is F, FCID is 0xe10280
Port vsan is 500
Operating Speed is 32 Gbps
Admin Speed is auto
Transmit B2B Credit is 12
Receive B2B Credit is 64
Receive data field Size is 2112
Beacon is turned off
fec state is enabled by default
5 minutes input rate 0 bits/sec,0 bytes/sec, 0 frames/sec
5 minutes output rate 0 bits/sec,0 bytes/sec, 0 frames/sec
16705 frames input,1225588 bytes
0 discards,0 errors
0 invalid CRC/FCS,0 unknown class
0 too long,0 too short
16714 frames output,1345676 bytes
0 discards,0 errors
0 input OLS,0 LRR,0 NOS,0 loop inits
7 output OLS,4 LRR, 0 NOS, 0 loop inits
Receive B2B Credit performance buffers is 0
12 transmit B2B credit remaining
0 low priority transmit B2B credit remaining
```

```
Interface last changed at Thu Nov 14 11:59:40 2019
```

Default Fibre Channel Interface Settings

The following table lists the default settings for native Fibre Channel interface parameters.

Table 8: Default Native Fibre Channel Interface Parameters

Parameters	Default
Interface mode	Auto
Interface speed	Auto
Administrative state	Shutdown (unless changed during initial setup)
Trunk mode	On (unless changed during initial setup)
Trunk-allowed VSANs	1 to 4093
Interface VSAN	Default VSAN (1)
Beacon mode	Off (disabled)
EISL encapsulation	Disabled
Data field size	2112 bytes

The following table lists the default settings for virtual Fibre Channel interface parameters.

Table 9: Default Virtual Fibre Channel Interface Parameters

Parameters	Default
Interface mode	F mode
Interface speed	n/a
Administrative state	Shutdown (unless changed during initial setup)
Trunk mode	On
Trunk-allowed VSANs	All VSANs
Interface VSAN	Default VSAN (1)
EISL encapsulation	n/a
Data field size	n/a

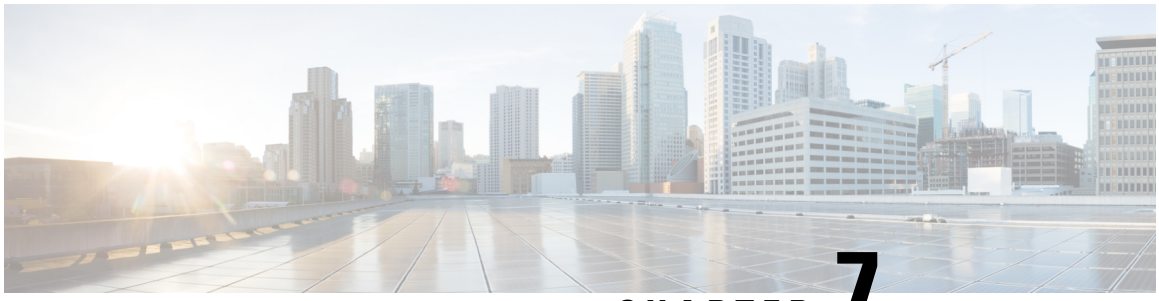

CHAPTER 7

Configuring and Managing VSANs

This chapter describes how to configure and manage VSANs.

This chapter includes the following sections:

- Configuring and Managing VSANs, on page 63
- Information About VSANs, on page 63
- Guidelines and Limitations for VSANs, on page 67
- Displaying the Static VSAN Configuration, on page 73
- Default Settings for VSANs, on page 73

Configuring and Managing VSANs

You can achieve higher security and greater stability in Fibre Channel fabrics by using virtual SANs (VSANs). 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.

Information About VSANs

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

VSAN Topologies

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, which increases 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.

The following figure 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 7: Logical VSAN Segmentation

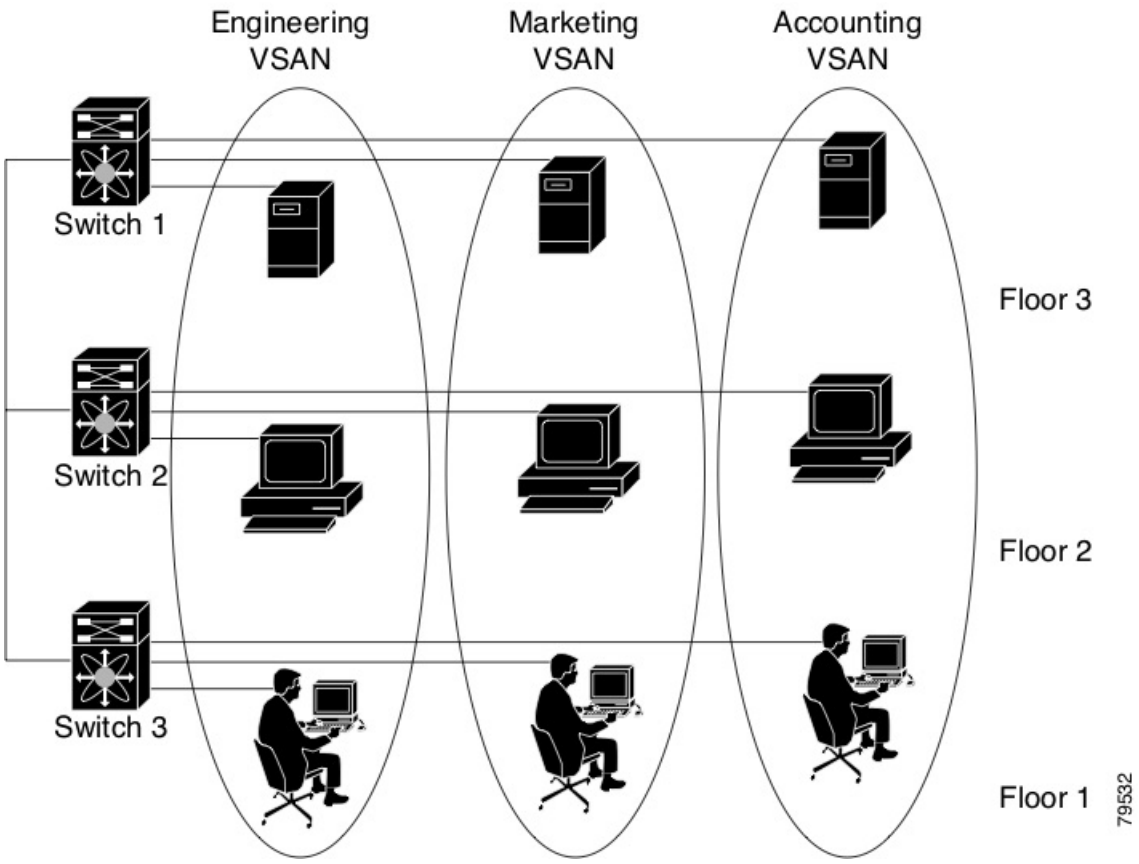

The application servers or storage arrays can be connected to the switch using Fibre Channel or virtual Fibre Channel interfaces. A VSAN can include a mixture of Fibre Channel and virtual Fibre Channel interfaces.

The following figure 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 8: Example of Two VSANs

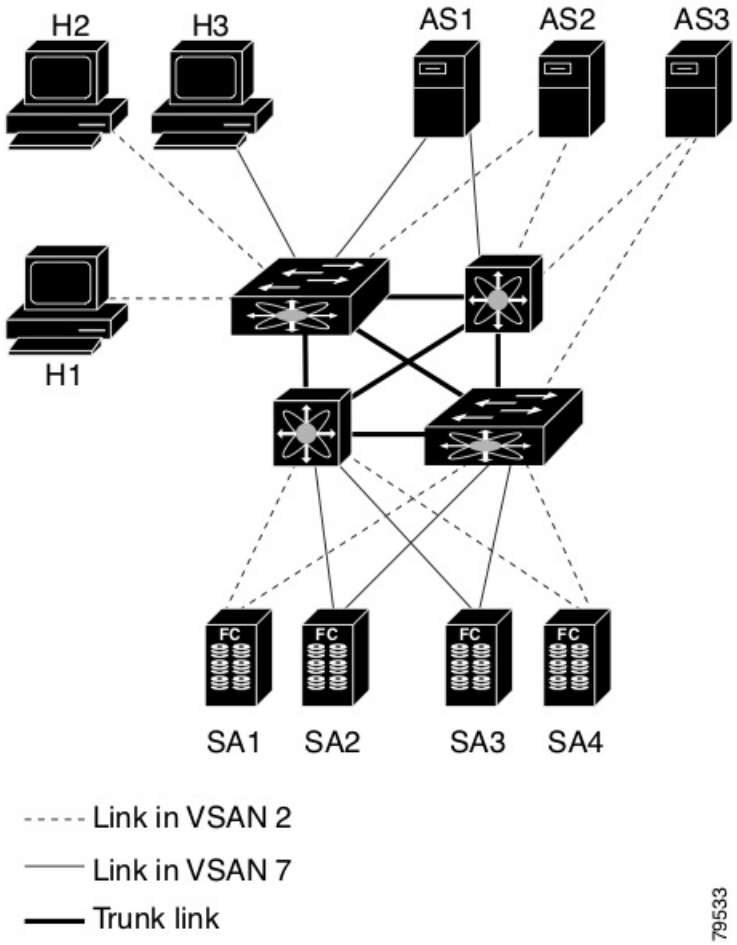

The four switches in this network are interconnected by VSAN trunk links that carry both VSAN 2 and VSAN 7 traffic. You can configure a different inter-switch topology for each VSAN. In the preceding figure, the inter-switch topology is identical for VSAN 2 and VSAN 7.

Without VSANs, a network administrator would need separate switches and links for separate SANs. By enabling VSANs, the same switches and links might be shared by multiple VSANs. VSANs allow SANs to be built on port granularity instead of switch granularity. The preceding figure 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 34 VSANs can be configured in a switch. Of these, one is a default VSAN (VSAN 1), and another is an isolated VSAN (VSAN 4094) and evfp isolated_vsan (vsan 4079). User-specified VSAN IDs range from 2 to 4078 and 4080-4093.

VSANs Versus Zones

Zones are always contained within a VSAN. 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. The following table lists the differences between VSANs and zones.

Table 10: 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.
VSANs limit unicast, multicast, and broadcast traffic.	Zones limit unicast traffic.
Membership is typically defined using the VSAN ID to F 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 F 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.

VSAN Characteristic	Zone Characteristic
VSANs encompass the entire fabric.	Zones are configured at the fabric edge.

The following figure 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. Zone A defined in VSAN 2 is different and separate from zone A defined in VSAN 7.

Figure 9: VSANS with Zoning

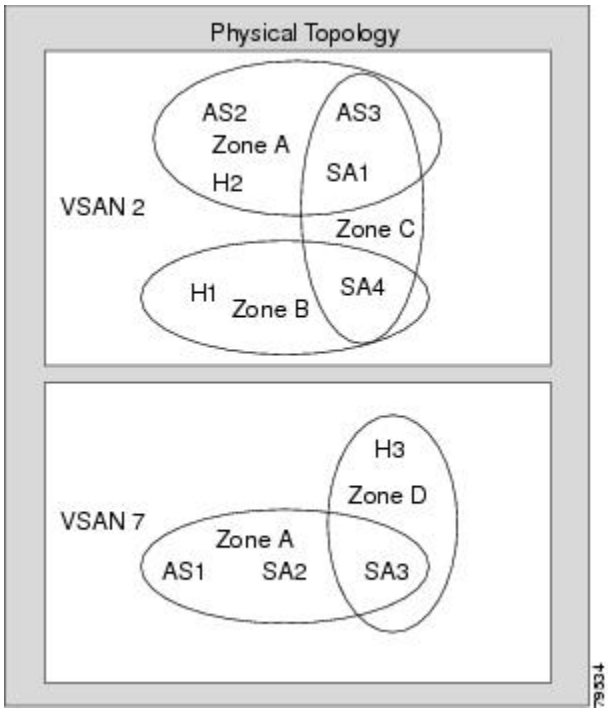

Guidelines and Limitations for VSANs

VSANs have the following configuration guidelines and limitations:

- VSAN ID—The VSAN ID identifies the VSAN as the default VSAN (VSAN 1), user-defined VSANs (VSAN 2-4078 and 4080-4093), evfp_isolated_vsan (vsan 4079), 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.
- 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.
- You can create only 32 VSANs in Cisco Nexus 9300-FX and 9700-FX platform switches, including the default VSAN 1.
- For 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 isolated VSAN:
  - The Isolated VSAN 4094 and Exchange Virtual Fabric Protocol (EVFP) isolated VSAN is 4079, are not available for user configuration.

About VSAN Creation

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

Creating VSANs Statically

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

SUMMARY STEPS

1. **configure terminal**
2. **vsan database**
3. **vsan** *vsan-id*
4. **vsan** *vsan-id* **name** *name*
5. **vsan** *vsan-id* **suspend**
6. switch(config-vsan-db)# **no vsan** *vsan-id* **suspend**
7. switch(config-vsan-db)# **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**configure terminal** **Example:**	Enters global configuration mode.

	Command or Action	Purpose
	`switch# configure terminal` `switch(config)#`	
Step 2	**vsan database** **Example:** `switch(config)# vsan database`	Configures the database for a VSAN. Application specific VSAN parameters cannot be configured from this prompt.
Step 3	**vsan** *vsan-id* **Example:** `switch(config-vsan-db)# vsan 360`	Creates a VSAN with the specified ID if that VSAN does not exist already.
Step 4	**vsan** *vsan-id* **name** *name* **Example:** `switch(config-vsan-db)# vsan 360 name test`	Updates the VSAN with the assigned name.
Step 5	**vsan** *vsan-id* **suspend** **Example:** `switch(config-vsan-db)# vsan 470 suspend`	Suspends the selected VSAN.
Step 6	switch(config-vsan-db)# **no vsan** *vsan-id* **suspend** **Example:** `switch(config-vsan-db)# no vsan 470 suspend`	Negates the **suspend** command issued in the previous step.
Step 7	switch(config-vsan-db)# **end** **Example:** `switch(config-vsan-db)# end`	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 statically (assigning VSANs to ports.) assign VSAN membership to ports.

VSAN trunking ports have an associated list of VSANs that are part of an allowed list.

Assigning Static Port VSAN Membership

You can statically assign VSAN membership for an interface port.

SUMMARY STEPS

1. **configure terminal**
2. **vsan database**
3. **vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**vsan database**
	Example:
	`switch(config)# vsan database`
	`switch(config-vsan-db)#`
		Configures the database for a VSAN.

Step 3	**vsan** *vsan-id* **Example:** `switch(config-vsan-db)# vsan 50`	Creates a VSAN with the specified ID if that VSAN does not exist already.

Default VSANs

The factory settings for Cisco SAN switches 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.

Up to 34 VSANs can be configured in a switch. Of these, one is a default VSAN (VSAN 1), and another is an isolated VSAN (VSAN 4094) and evfp isolated_vsan (vsan 4079). User-specified VSAN IDs range from 2 to 4078 and 4080-4093.

Isolated VSANs

VSAN 4094 is an isolated VSAN. When a VSAN is deleted, all nontrunking ports are transferred to the isolated VSAN to avoid an implicit transfer of ports to the default VSAN or to another configured VSAN. This action ensures that all ports in the deleted VSAN become 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 34 VSANs can be configured in a switch. Of these, one is a default VSAN (VSAN 1), and another is an isolated VSAN (VSAN 4094) and evfp isolated_vsan (vsan 4079). User-specified VSAN IDs range from 2 to 4078 and 4080-4093.

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 the figure below).

Figure 10: VSAN Port Membership Details

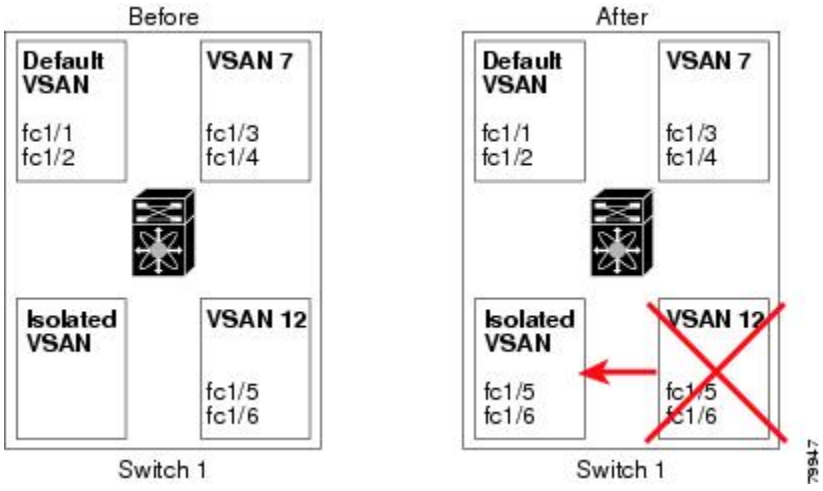

- VSAN-based runtime (name server), zoning, and configuration (static routes) information is removed when the VSAN is deleted.
- Configured VSAN interface information is removed when the VSAN is deleted.

Note The allowed VSAN list is not affected when a VSAN is deleted.

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

Deleting Static VSANs

You can delete a VSAN and its various attributes.

SUMMARY STEPS

1. **configure terminal**
2. **vsan database**
3. **vsan** *vsan-id*
4. switch(config-vsan-db)# **no vsan** *vsan-id*
5. switch(config-vsan-db)# **end**

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**vsan database**
	Example:
	`switch(config)# vsan database`
	`switch(config-vsan-db)#`
		Configures the VSAN database.

Step 3	**vsan** *vsan-id* **Example:** `switch(config-vsan-db)# vsan 2`	Places you in VSAN configuration mode.
Step 4	switch(config-vsan-db)# **no vsan** *vsan-id* **Example:** `switch(config-vsan-db)# no vsan 5`	Deletes VSAN 5 from the database and switch.
Step 5	switch(config-vsan-db)# **end** **Example:** `switch(config-vsan-db)# end`	Places you in EXEC mode.

Interop Mode

Interoperability enables the products of multiple vendors to connect with each other. Fibre Channel standards guide vendors to create common external Fibre Channel interfaces.

Displaying the Static VSAN Configuration

The following example shows how to display information about a specific VSAN:

```
switch# show vsan 100
```

The following example shows how to display VSAN usage:

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

The following example shows how to display all VSANs:

```
switch# show vsan
```

Default Settings for VSANs

The following table lists the default settings for all configured VSANs.

Table 11: 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).

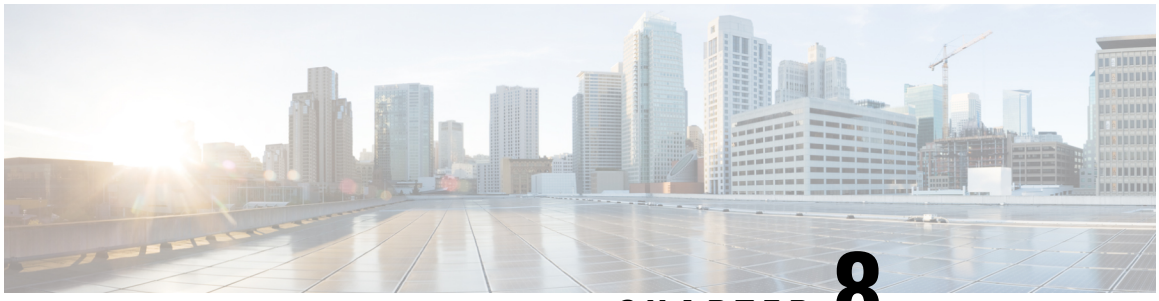

CHAPTER 8

Configuring SAN Port Channels

This chapter contains the following sections:

- Configuring SAN Port Channels, on page 75

Configuring SAN Port Channels

SAN port channels refer to the aggregation of multiple physical interfaces into one logical interface to provide higher aggregated bandwidth, load balancing, and link redundancy.

On Cisco Nexus 9000 switches, SAN port channels can include physical Fibre Channel interfaces. However virtual Fibre Channel interfaces are not supported. A SAN port channel can include up to 16 Fibre Channel interfaces.

Information About SAN Port Channels

About E and TE Port Channels

An E port channel refers to the aggregation of multiple E ports into one logical interface to provide higher aggregated bandwidth, load balancing, and link redundancy. Port channel can connect to interfaces across switching modules, so a failure of a switching module cannot bring down the port channel link. Cisco Nexus devices support a maximum of four SAN port channels in FC switch mode, which includes E/TE-port port channels.

A SAN port channel has the following functionality:

- Provides a point-to-point connection over ISL (E ports) or EISL (TE ports). Multiple links can be combined into a SAN port channel.
- Increases the aggregate bandwidth on an ISL by distributing traffic among all functional links in the channel.
- Load balances across multiple links and maintains optimum bandwidth utilization. Load balancing is based on the source ID, destination ID, and exchange ID (OX ID).
- Provides high availability on an ISL. If one link fails, traffic previously carried on this link is switched to the remaining links. If a link goes down in a SAN port channel, the upper layer protocol is not aware of it. To the upper layer protocol, the link is still there, although the bandwidth is diminished. The routing tables are not affected by link failure.

About F and TF Port Channels

An F port channel is also a logical interface that combines a set of F ports connected to the same Fibre Channel node and operates as one link between the F ports and the NP ports. The F port channels support bandwidth utilization and availability like the E port channels. F port channel are mainly used to connect Nexus 9000 core and NPV switches to provide optimal bandwidth utilization and transparent failover between the uplinks of a VSAN. An F port channel trunk combines the functionality and advantages of a TF port and an F port channel. This logical link uses the Cisco PTP and PCP protocols over Cisco EPP (ELS). Cisco Nexus devices support a maximum of four SAN port channels in FC switch mode, which includes F/TF-port port channels.

Understanding Port Channels and VSAN Trunking

Cisco Nexus devices implement VSAN trunking and port channels as follows:

- A SAN port channel enables several physical links to be combined into one aggregated logical link.
- An industry standard E port can link to other vendor switches and is referred to as inter-switch link (ISL), as shown on the left side of the figure below.
- VSAN trunking enables a link transmitting frames in the EISL format to carry traffic for multiple VSAN . When trunking is operational on an E port, that E port becomes a TE port. EISLs connects only between Cisco switches, as shown on the right side of the figure below.

Figure 11: VSAN Trunking Only

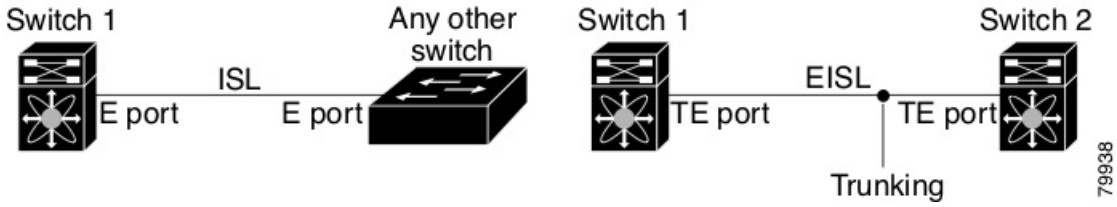

- You can create a SAN port channel with members that are E ports, as shown on the left side of the figure below. In this configuration, the port channel implements a logical ISL (carrying traffic for one VSAN).
- You can create a SAN port channel with members that are TE-ports, as shown on the right side of the figure below. In this configuration, the port channel implements a logical EISL (carrying traffic for multiple VSANs).

Figure 12: Port Channels and VSAN Trunking

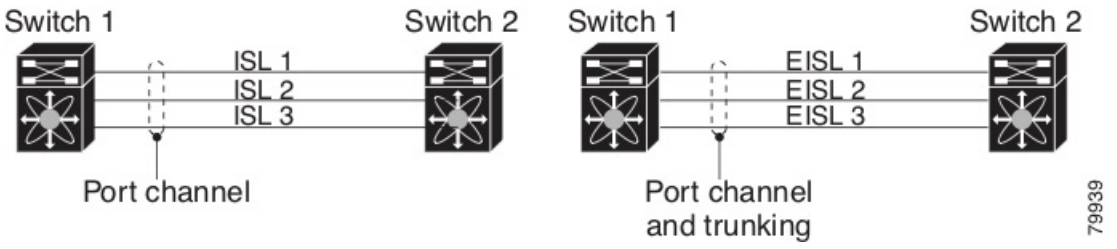

- Port channel interfaces can be channeled between the following port sets:
  - E ports and TE ports
  - F ports and NP ports

    - TF ports and TNP ports

- Trunking permits traffic on multiple VSANs between switches.
- Port channels and trunking can be used between TE ports over EISLs.

Understanding Load Balancing

Load-balancing functionality can be provided using the following methods:

- Flow based—All frames between source and destination follow the same links for a given flow. That is, whichever link is selected for the first exchange of the flow is used for all subsequent exchanges.
- Exchange based—The first frame in an exchange is assigned to a link, and then subsequent frames in the exchange follow the same link. However, subsequent exchanges can use a different link. This method provides finer granularity for load balancing while preserving the order of frames for each exchange.

The following figure illustrates how flow-based load balancing works. When the first frame in a flow is received on an interface for forwarding, link 1 is selected. Each subsequent frame in that flow is sent over the same link. No frame in SID1 and DID1 utilizes link 2.

Figure 13: SID1, DID1, and Flow-Based Load Balancing

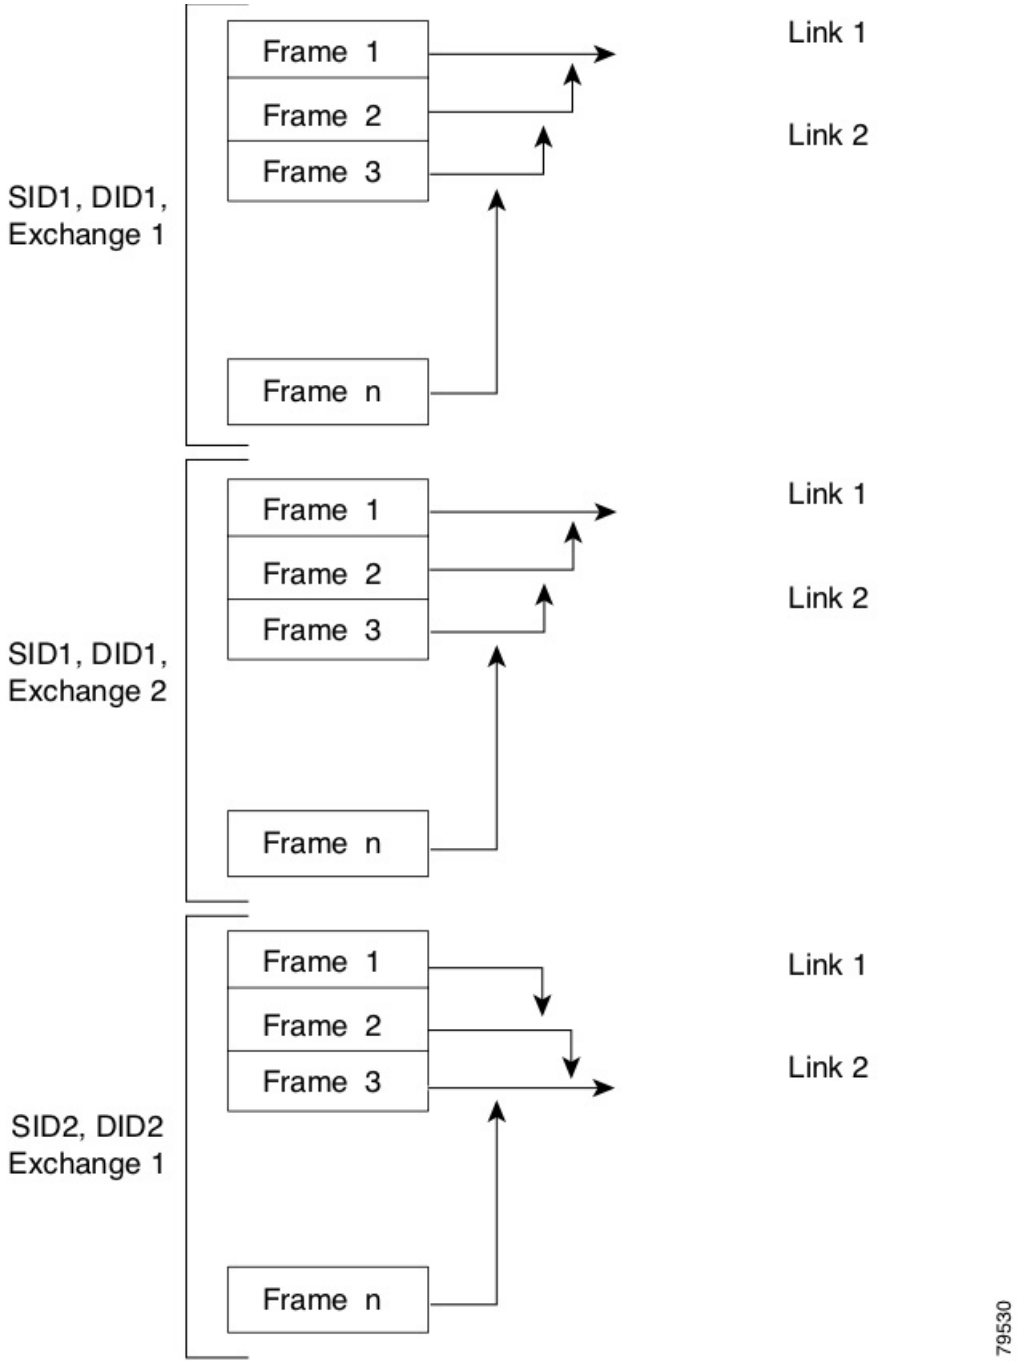

The following figure illustrates how exchange-based load balancing works. When the first frame in an exchange is received for forwarding on an interface, link 1 is chosen by a hash algorithm. All remaining frames in that particular exchange are sent on the same link. For exchange 1, no frame uses link 2. For the next exchange, link 2 is chosen by the hash algorithm. Now all frames in exchange 2 use link 2.

Figure 14: SID1, DID1, and Exchange-Based Load Balancing

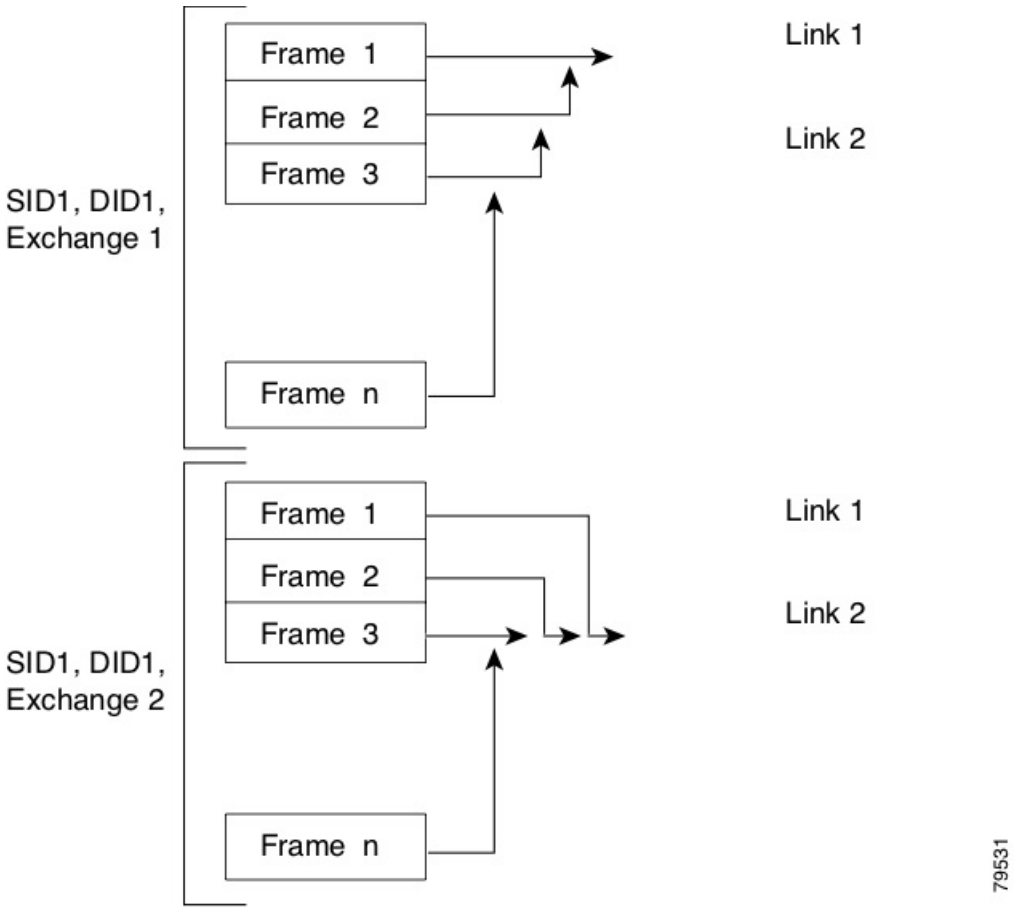

Configuring SAN Port Channels

SAN port channels are created with default values. You can change the default configuration just as any other physical interface.

The following figure provides examples of valid SAN port channel configurations.

Figure 15: Valid SAN Port Channel Configurations

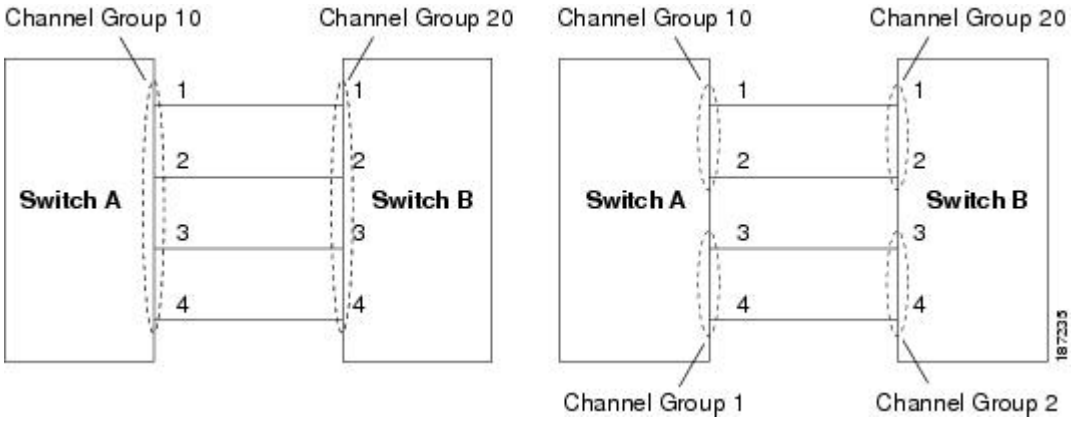

The following figure shows examples of invalid configurations. Assuming that the links are brought up in the 1, 2, 3, 4 sequence, links 3 and 4 will be operationally down as the fabric is misconfigured.

Figure 16: Misconfigured Configurations

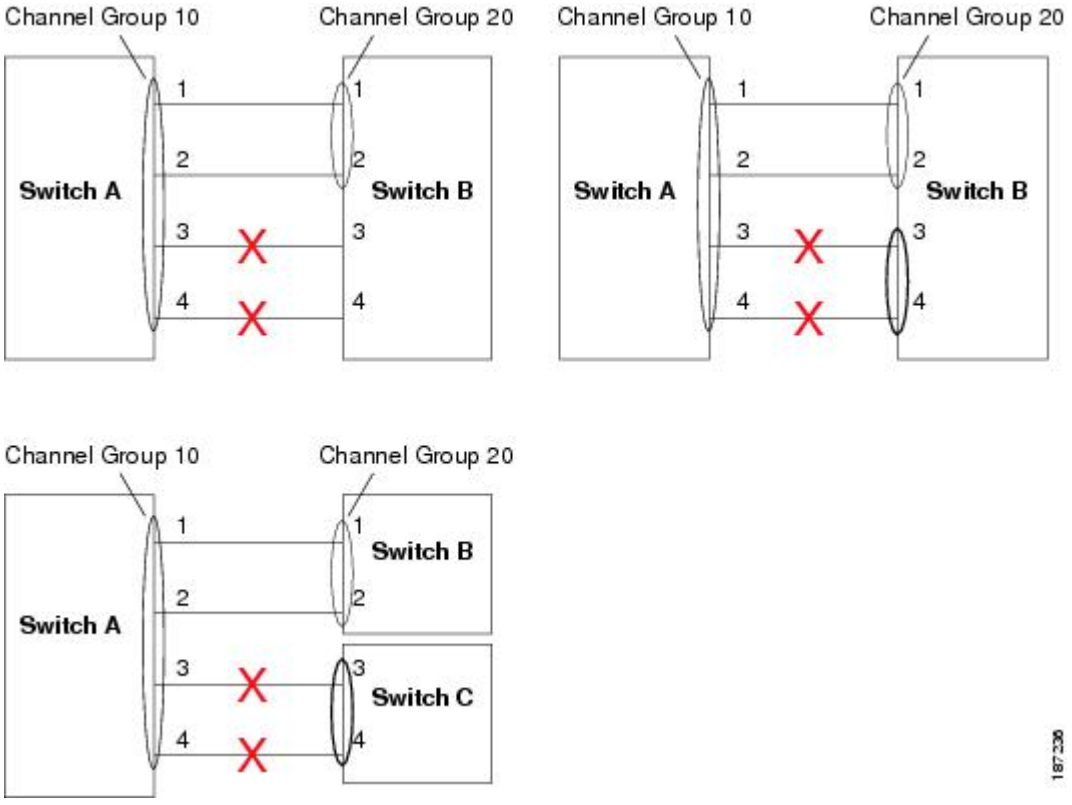

SAN Port Channel Configuration Guidelines

Before configuring a SAN port channel, consider the following guidelines:

- Port-channel mode is active by default. Port-channel **ON** mode is not supported.
- Configure SAN port Channel using Fibre Channel ports from different port groups.

- Ensure that one SAN port channel is not connected to different sets of switches. SAN port channels require point-to-point connections between the same set of switches.
- If you misconfigure SAN port channels, you may receive a misconfiguration message. If you receive this message, the port channel's physical links are disabled because an error has been detected.
- If the following requirements are not met, a SAN port channel error is detected:
  - Each switch on either side of a SAN port channel must be connected to the same number of interfaces.
  - Each interface must be connected to a corresponding interface on the other side.
  - Links in a SAN port channel cannot be changed after the port channel is configured. If you change the links after the port channel is configured, be sure to reconnect the links to interfaces within the port channel and reenable the links.

    If all three conditions are not met, the faulty link is disabled.

Enter the **show interface** command for that interface to verify that the SAN port channel is functioning as required.

F and TF Port Channel Guidelines

The guidelines for F and TF port channels are as follows:

- The ports must be in F mode.
- Automatic creation is not supported.
- ON mode is not supported. Only Active-Active mode is supported. By default, the mode is Active on the NPV switches.
- The name server registration of the N ports logging in through an F port channel will use the FWWN of the port channel interface.
- Before you configure F port channel, make sure that the feature fport-channel-trunk is enabled on the switch.
- For an NPV switch which is configured for trunking on any interface, or for 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 isolated VSAN:
  - If the trunk mode is enabled for any of the interfaces, or if the NP port channel is up, the reserved VSANs range from 3840 to 4078, which are not available for user configuration.
  - The Exchange Virtual Fabric Protocol (EVFP) isolated VSAN is 4079, and it is not available for user configuration.

Creating a SAN Port Channel

To create a SAN port channel, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface san-port-channel** *channel-number*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface san-port-channel** *channel-number*	Creates the specified SAN port channel using the default mode (on). The SAN port channel number is in the range of 1 to 256. **Note** Enter an unused channel number to create a new SAN port channel (for Fibre Channel ports). To view the range of used and unused channel numbers use the **show san-port-channel usage** command.

About Port Channel Modes

You can configure each SAN port channel with a channel group mode parameter to determine the port channel protocol behavior for all member ports in this channel group. The possible values for a channel group mode are as follows:

- On (default)—The member ports only operate as part of a SAN port channel or remain inactive. In this mode, the port channel protocol is not initiated. However, if a port channel protocol frame is received from a peer port, the software indicates its nonnegotiable status. Port channels configured in the On mode require you to explicitly enable and disable the port channel member ports at either end if you add or remove ports from the port channel configuration. You must physically verify that the local and remote ports are connected to each other.
- Active—The member ports initiate port channel protocol negotiation with the peer port(s) regardless of the channel group mode of the peer port. If the peer port, while configured in a channel group, does not support the port channel protocol, or responds with a nonnegotiable status, it will default to the On mode behavior. The Active port channel mode allows automatic recovery without explicitly enabling and disabling the port channel member ports at either end.

Note A F port channel is supported only in Active Mode.

The table below compares On and Active modes.

Table 12: Channel Group Configuration Differences

On Mode	Active Mode
No protocol is exchanged.	A port channel protocol negotiation is performed with the peer ports.
Moves interfaces to the suspended state if its operational values are incompatible with the SAN port channel.	Moves interfaces to the isolated state if its operational values are incompatible with the SAN port channel.

On Mode	Active Mode
When you add or modify a port channel member port configuration, you must explicitly disable (shut) and enable (no shut) the port channel member ports at either end.	When you add or modify a port channel interface, the SAN port channel automatically recovers.
Port initialization is not synchronized.	There is synchronized startup of all ports in a channel across peer switches.
All misconfigurations are not detected as no protocol is exchanged.	Consistently detect misconfigurations using a port channel protocol.
Transitions misconfigured ports to the suspended state. You must explicitly disable (shut) and enable (no shut) the member ports at either end.	Transitions misconfigured ports to the isolated state to correct the misconfiguration. Once you correct the misconfiguration, the protocol ensures automatic recovery.
This is the default mode.	You must explicitly configure this mode.

Configuring Active Mode SAN Port Channel

To configure active mode, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface san-port-channel** *channel-number*
3. switch(config-if)# **channel mode active**
4. switch(config-if)# **no channel mode active**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface san-port-channel** *channel-number*	Configures the specified port channel using the default On mode. The SAN port channel number is in the range of 1 to 256.
Step 3	switch(config-if)# **channel mode active**	Configures the Active mode.
Step 4	switch(config-if)# **no channel mode active**	Reverts to the default On mode.

Example of Configuring Active Modes

The following example shows how to configure active mode:

```
switch(config)# interface san-port-channel 1
switch(config-if)# channel mode active
```

About SAN Port Channel Deletion

When you delete the SAN port channel, the corresponding channel membership is also deleted. All interfaces in the deleted SAN port channel convert to individual physical links. After the SAN port channel is removed, regardless of the mode (active and on) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down.

If you delete the SAN port channel for one port, then the individual ports within the deleted SAN port channel retain the compatibility parameter settings (speed, mode, port VSAN, allowed VSAN, and port security). You can explicitly change those settings as required.

- If you use the default On mode to avoid inconsistent states across switches and to maintain consistency across switches, then the ports shut down. You must explicitly enable those ports again.
- If you use the Active mode, then the port channel ports automatically recover from the deletion.

Deleting SAN Port Channels

To delete a SAN port channel, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **no interface san-port-channel** *channel-number*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **no interface san-port-channel** *channel-number*	Deletes the specified port channel, its associated interface mappings, and the hardware associations for this SAN port channel.

Interfaces in a SAN Port Channel

You can add or remove a physical Fibre Channel interface (or a range of interfaces) to an existing SAN port channel. The compatible parameters on the configuration are mapped to the SAN port channel. Adding an interface to a SAN port channel increases the channel size and bandwidth of the SAN port channel. Removing an interface from a SAN port channel decreases the channel size and bandwidth of the SAN port channel.

Note Virtual Fibre Channel interfaces cannot be added to SAN port channels.

About Interface Addition to a SAN Port Channel

You can add a physical interface (or a range of interfaces) to an existing SAN port channel. The compatible parameters on the configuration are mapped to the SAN port channel. Adding an interface to a SAN port channel increases the channel size and bandwidth of the SAN port channel.

After the members are added, regardless of the mode (Active and On) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down.

For adding Fibre channel (FC) Breakout (BO) interfaces to a SAN Port channel on Cisco Nexus N9K-C9336C-FX2-E platform switch, see SAN Switching General Guidelines and Limitations.

Compatibility Check

A compatibility check ensures that the same parameter settings are used in all physical ports in the channel. Otherwise, they cannot become part of a SAN port channel. The compatibility check is performed before a port is added to the SAN port channel.

The check ensures that the following parameters and settings match at both ends of a SAN port channel:

- Capability parameters (type of interface, Fibre Channel at both ends).
- Administrative compatibility parameters (speed, mode, port VSAN and allowed VSAN).
- Operational parameters (speed and remote switch's WWN).

A port addition procedure fails if the capability and administrative parameters in the remote switch are incompatible with the capability and administrative parameters in the local switch. If the compatibility check is successful, the interfaces are operational and the corresponding compatibility parameter settings apply to these interfaces.

After you enable forcing a port to be added to a channel group by entering the **channel-group force** command, the following two conditions occur:

- When an interface joins a port channel the following parameters are removed and they are operationally replaced with the values on the port channel; however, this change will not be reflected in the running-configuration for the interface:
  - Bandwidth
  - Delay
  - Service policy
  - ACLs

  When an interface joins or leaves a port channel, the following parameters remain unaffected:
  - Beacon
  - Description
  - LACP port priority
  - Debounce
  - Shutdown
  - SNMP traps

Suspended and Isolated States

If the operational parameters are incompatible, the compatibility check fails and the interface is placed in a suspended or isolated state based on the configured mode:

- An interface enters the suspended state if the interface is configured in the On mode.
- An interface enters the isolated state if the interface is configured in the Active mode.

Adding an Interface to a SAN Port Channel

To add an interface to a SAN port channel, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *type slot*/port
3. switch(config-if)# **channel-group** *channel-number*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface** *type slot*/port	Enters configuration mode for the specified interface.
Step 3	switch(config-if)# **channel-group** *channel-number*	Adds the Fibre Channel interface to the specified channel group. If the channel group does not exist, it is created. The port is shut down.

Forcing an Interface Addition

You can force the port configuration to be overwritten by the SAN port channel. In this case, the interface is added to a SAN port channel.

- If you use the default On mode to avoid inconsistent states across switches and to maintain consistency across switches, then the ports shut down. You must explicitly enable those ports again.
- If you use the Active mode, then the port channel ports automatically recover from the addition.

Note When SAN port channels are created from within an interface, the **force** option cannot be used.

After the members are forcefully added, regardless of the mode (Active and On) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down.

To force the addition of a port to a SAN port channel, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *type slot*/*port*
3. switch(config-if)# **channel-group** *channel-number* **force**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface** *type slot*/*port*	Enters configuration mode for the specified interface.
Step 3	switch(config-if)# **channel-group** *channel-number* **force**	Forces the addition of the interface into the specified channel group. The E port is shut down.

About Interface Deletion from a SAN Port Channel

When a physical interface is deleted from the SAN port channel, the channel membership is automatically updated. If the deleted interface is the last operational interface, then the port channel status is changed to a down state. Deleting an interface from a SAN port channel decreases the channel size and bandwidth of the SAN port channel.

- If you use the default On mode to avoid inconsistent states across switches and to maintain consistency across switches, then the ports shut down. You must explicitly enable those ports again.
- If you use the Active mode, then the port channel ports automatically recover from the deletion.

After the members are deleted, regardless of the mode (Active and On) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down.

Deleting an Interface from a SAN Port Channel

To delete a physical interface (or a range of physical interfaces) from a SAN port channel, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** *type slot*/*portBO port*
3. switch(config-if)# **no channel-group** *channel-number*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface** *type slot*/*portBO port*	Enters configuration mode for the specified interface.
Step 3	switch(config-if)# **no channel-group** *channel-number*	Deletes the physical Fibre Channel interface from the specified channel group.

SAN Port Channel Protocol

The switch software provides robust error detection and synchronization capabilities. You can manually configure channel groups. The channel groups have the same capability and configurational parameters. Any change in configuration applied to the associated SAN port channel interface is propagated to all members of the channel group.

Cisco SAN switches support a protocol to exchange SAN port channel configurations, which simplifies port channel management with incompatible ISLs. An additional autocreation mode enables ISLs with compatible parameters to automatically form channel groups without manual intervention.

The port channel protocol is enabled by default.

The port channel protocol expands the port channel functional model in Cisco SAN switches. It uses the exchange peer parameters (EPP) services to communicate across peer ports in an ISL. Each switch uses the information received from the peer ports along with its local configuration and operational values to decide if it should be part of a SAN port channel. The protocol ensures that a set of ports are eligible to be part of the same SAN port channel. They are only eligible to be part of the same port channel if all the ports have a compatible partner.

The port channel protocol uses two subprotocols:

- Bringup protocol—Automatically detects misconfigurations so you can correct them. This protocol synchronizes the SAN port channel at both ends so that all frames for a given flow (as identified by the source FC ID, destination FC ID and OX_ID) are carried over the same physical link in both directions.

About Manually Configured Channel Groups

A user-configured channel group cannot be converted to an autocreated channel group. However, you can convert an autocreated channel group to a manual channel group. This task is irreversible. The channel group number does not change, but the member ports operate according to the properties of the manually configured channel group, and channel group autocreation is implicitly disabled for all the member ports.

If you enable persistence, be sure to enable it at both ends of the SAN port channel.

Example Port Channel Configurations

This section shows examples on how to configure an F port channel in shared mode and how to bring up the link between F ports on the NPIV core switches and NP ports on the NPV switches. Before you configure the F port channel, ensure that F port trunking, F port channeling, and NPIV are enabled.

Example

This example shows how to create the port channel:

```
switch(config)# interface san-po-channel 2
switch(config-if)# switchport mode F
switch(config-if)# channel mode active
switch(config-if)# exit
```

This example shows how to configure the port channel member interfaces on the core switch:

```
switch(config)# interface fc1/4-6
switch(config-if)# shut
switch(config-if)# switchport mode F
switch(config-if)# switchport speed 32000
switch(config-if)# switchport trunk mode on
switch(config-if)# channel-group 2
switch(config-if)# no shut
```

```
switch(config-if)# exit
```

Verifying SAN Port Channel Configuration

You can view specific information about existing SAN port channels at any time from EXEC mode. The following **show** commands provide further details on existing SAN port channels.

The **show san-port-channel summary** command displays a summary of SAN port channels within the switch. A one-line summary of each SAN port channel provides the administrative state, the operational state, the number of attached and active interfaces (up), and the first operational port (FOP), which is the primary operational interface selected in the SAN port channel to carry control-plane traffic (no load-balancing). The FOP is the first port that comes up in a SAN port channel and can change if the port goes down. The FOP is also identified by an asterisk (*) in show san-port-channel database cli.

To display VSAN configuration information, perform one of the following tasks:

SUMMARY STEPS

1. switch# **show san-port-channel summary** | **database** | **consistency** [**details**] | **usage** | **compatibility-parameters**
2. switch# **show san-port-channel database interface san-port-channel** *channel-number*
3. switch# switch# **show interface fc** *slot*/*port*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **show san-port-channel summary** | **database** | **consistency** [**details**] | **usage** | **compatibility-parameters**	Displays SAN port channel information.
Step 2	switch# **show san-port-channel database interface san-port-channel** *channel-number*	Displays information for the specified SAN port channel.
Step 3	switch# switch# **show interface fc** *slot*/*port*	Displays VSAN configuration information for the specified Fibre Channel interface.

Example of Verification Commands

The following example shows how to display a summary of SAN port channel information:

```
switch# show san-port-channel summary
--------------------------------------------------------------------------------
Interface                Total Ports        Oper Ports        First Oper Port
--------------------------------------------------------------------------------
san-port-channel 7             2                 0                 --
san-port-channel 8             2                 0                 --
san-port-channel 9             2                 2
```

The following example shows how to display SAN port channel consistency:

```
switch# show san-port-channel consistency
Database is consistent
```

The following example shows how to display details of the used and unused port channel numbers:

```
switch# show san-port-channel usage
Totally 3 port-channel numbers used
===================================
Used  :   77 - 79
Unused:   1 - 76 , 80 - 256
```

Default Settings for SAN Port Channels

The table below lists the default settings for SAN port channels.

Table 13: Default SAN Port Channel Parameters

Parameters	Default
Port channels	FSPF is enabled by default.
Create port channel	Administratively up.
Default port channel mode	On.
Autocreation	Disabled.

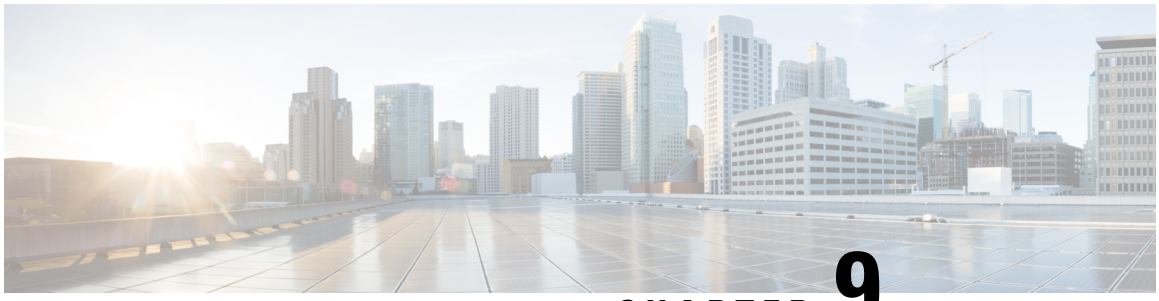

CHAPTER 9

Configuring Fibre Channel Domain Parameters

This chapter describes how to configure Fibre Channel domain parameters.

This chapter includes the following sections:

- Information About Domain Parameters, on page 91

Information About Domain Parameters

The Fibre Channel domain (fcdomain) feature performs principal switch selection, domain ID distribution, FC ID allocation, and fabric reconfiguration functions as described in the FC-SW-2 standards. The domains are configured on a per-VSAN basis. If you do not configure a domain ID, the local switch uses a random ID.

Caution Changes to fcdomain parameters should not be performed on a daily basis. These changes should be made by an administrator or individual who is completely familiar with switch operations.

When you change the configuration, be sure to save the running configuration. The next time you reboot the switch, the saved configuration is used. If you do not save the configuration, the previously saved startup configuration is used.

Fibre Channel Domains

The fcdomain has four phases:

- Principal switch selection—This phase guarantees the selection of a unique principal switch across the fabric.
- Domain ID distribution—This phase guarantees that each switch in the fabric obtains a unique domain ID.
- FC ID allocation—This phase guarantees a unique FC ID assignment to each device attached to the corresponding switch in the fabric.
- Fabric reconfiguration—This phase guarantees a resynchronization of all switches in the fabric to ensure they simultaneously restart a new principal switch selection phase.

The following figure shows an example fcdomain configuration.

Figure 17: Sample fcdomain Configuration

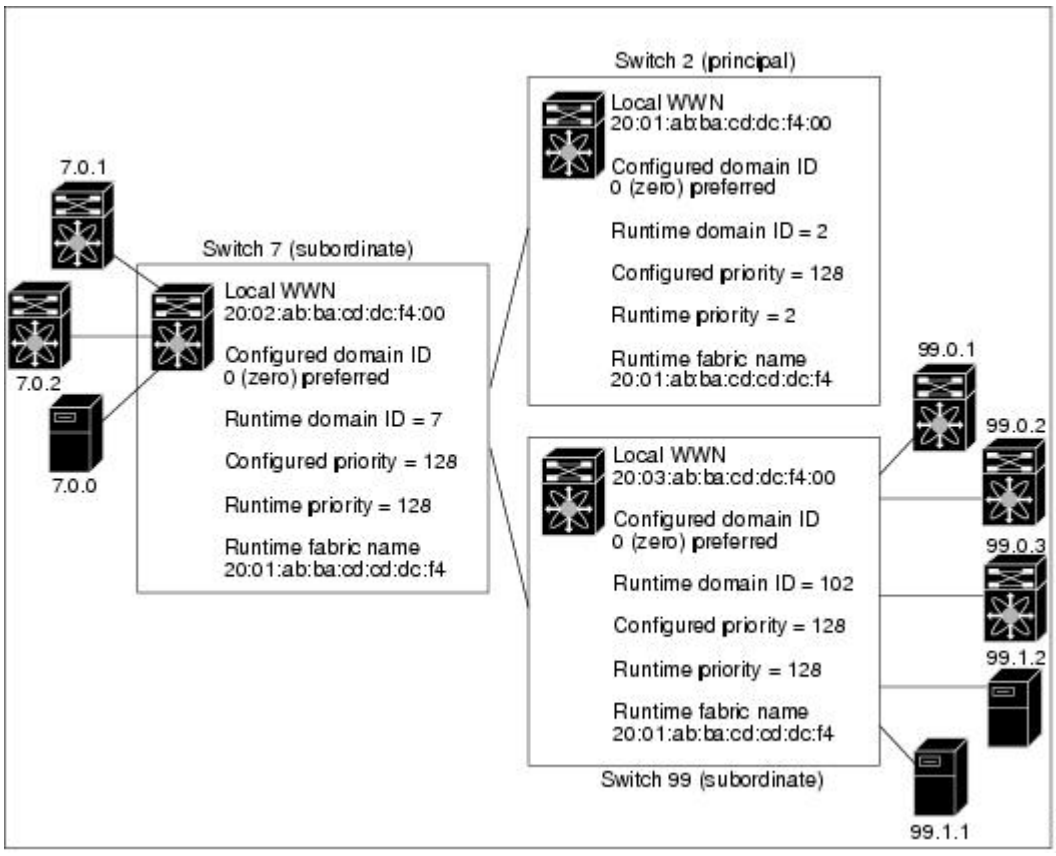

Domain Restarts

Fibre Channel domains can be started disruptively or nondisruptively. If you perform a disruptive restart, reconfigure fabric (RCF) frames are sent to other switches in the fabric and data traffic is disrupted on all the switches in the VSAN (including remotely segmented ISLs). If you perform a nondisruptive restart, build fabric (BF) frames are sent to other switches in the fabric and data traffic is disrupted only on the switch.

If you are attempting to resolve a domain ID conflict, you must manually assign domain IDs. A disruptive restart is required to apply most configuration changes, including manually assigned domain IDs. Nondisruptive domain restarts are acceptable only when changing a preferred domain ID into a static one (and the actual domain ID remains the same).

Note A static domain is specifically configured by the user and may be different from the runtime domain. If the domain IDs are different, the runtime domain ID changes to take on the static domain ID after the next restart.

Note Beginning Cisco NX-OS Release 9.3(5), Cisco Nexus C93180YC-FX switches support disruptive Restart Reconfigure Fabric (RCF).

You can apply most of the configurations to their corresponding runtime values. Each of the following sections provide further details on how the fcdomain parameters are applied to the runtime values.

The **fcdomain restart** command applies your changes to the runtime settings. The disruptive option is not supported.

Restarting a Domain

You can restart the fabric disruptively or nondisruptively.

SUMMARY STEPS

1. **configure terminal**
2. **fcdomain restart vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**fcdomain restart vsan** *vsan-id* **Example:** `switch (config)# fcdomain restart vsan 100`	Forces the VSAN to reconfigure without traffic disruption. The VSAN ID ranges from 1 to 4093.

Switch Priority

By default, the configured priority is 128. The valid range to set the priority is between 1 and 254. Priority 1 has the highest priority. Value 255 is accepted from other switches, but cannot be locally configured.

Any new switch cannot become the principal switch when it joins a stable fabric. During the principal switch selection phase, the switch with the highest priority becomes the principal switch. If two switches have the same configured priority, the switch with the lower world-wide name (WWN) becomes the principal switch.

The priority configuration is applied to runtime when the fcdomain is restarted. This configuration is applicable to both disruptive and nondisruptive restarts.

Configuring Switch Priority

You can configure the priority for the principal switch.

SUMMARY STEPS

1. **configure terminal**
2. **fcdomain priority** *number* **vsan** *vsan-id*
3. **no fcdomain priority** *number* **vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**fcdomain priority** *number* **vsan** *vsan-id* **Example:** `switch(config)# fcdomain priority 12 vsan 1`	Configures the specified priority for the local switch in the specified VSAN. The fcdomain priority ranges from 1 to 254. The VSAN ID ranges from 1 to 4093.
Step 3	**no fcdomain priority** *number* **vsan** *vsan-id* **Example:** `switch(config)# no fcdomain priority 12 vsan 1`	Reverts the priority to the factory default (128) in the specified VSAN. The fcdomain priority ranges from 1 to 254. The VSAN ID ranges from 1 to 4093.

Configuring Fabric Names

You can set the fabric name value for a disabled fcdomain.

SUMMARY STEPS

1. **configure terminal**
2. **fcdomain fabric-name 20:1:ac:16:5e:0:21:01 vsan** *vsan-id*
3. **no fcdomain fabric-name 20:1:ac:16:5e:0:21:01 vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**fcdomain fabric-name 20:1:ac:16:5e:0:21:01 vsan** *vsan-id* **Example:** `switch(config)# fcdomain fabric-name 20:1:ac:16:5e:0:21:01 vsan 1`	Assigns the configured fabric name value in the specified VSAN. The VSAN ID ranges from 1 to 4093.
Step 3	**no fcdomain fabric-name 20:1:ac:16:5e:0:21:01 vsan** *vsan-id* **Example:** `switch(config)# no fcdomain fabric-name 20:1:ac:16:5e:0:21:01 vsan 1`	Changes the fabric name value to the factory default (20:01:00:05:30:00:28:df) in VSAN 3010. The VSAN ID ranges from 1 to 4093.

Incoming RCFs

You can configure the rcf-reject option on a per-interface, per-VSAN basis. By default, the rcf-reject option is disabled (that is, RCF request frames are not automatically rejected).

The rcf-reject option takes effect immediately.

No fcdomain restart is required.

Note You do not need to configure the RCF reject option on virtual Fibre Channel interfaces.

Rejecting Incoming RCFs

You can reject incoming RCF request frames.

SUMMARY STEPS

1. **configure terminal**
2. **fcdomain rcf-reject vsan** *vsan-id*
3. **no fcdomain rcf-reject vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**fcdomain rcf-reject vsan** *vsan-id* **Example:** `switch(config-if)# fcdomain rcf-reject vsan 10`	Enables the RCF filter on the specified interface in the specified VSAN. The VSAN ID ranges from 1 to 4093.
Step 3	**no fcdomain rcf-reject vsan** *vsan-id* **Example:** `switch(config-if)# no fcdomain rcf-reject vsan 10`	Disables (default) the RCF filter on the specified interface in the specified VSAN. The VSAN ID ranges from 1 to 4093.

Autoreconfiguring Merged Fabrics

By default, the autoreconfigure option is disabled. When you join two switches belonging to two different stable fabrics that have overlapping domains, the following situations can occur:

- If the autoreconfigure option is enabled on both switches, a disruptive reconfiguration phase is started.
- If the autoreconfigure option is disabled on either or both switches, the links between the two switches become isolated.

The autoreconfigure option takes immediate effect at runtime. You do not need to restart the fcdomain. If a domain is currently isolated due to domain overlap, and you later enable the autoreconfigure option on both

switches, the fabric continues to be isolated. If you enabled the autoreconfigure option on both switches before connecting the fabric, a disruptive reconfiguration (RCF) will occur. A disruptive reconfiguration can affect data traffic. You can nondisruptively reconfigure the fcdomain by changing the configured domains on the overlapping links and eliminating the domain overlap.

Enabling Autoreconfiguration

You can enable automatic reconfiguration in a specific VSAN (or range of VSANs).

SUMMARY STEPS

1. **configure terminal**
2. **fcdomain auto-reconfigure vsan** *vsan-id*
3. **no fcdomain auto-reconfigure vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**fcdomain auto-reconfigure vsan** *vsan-id* **Example:** `switch(config)# fcdomain auto-reconfigure vsan 1`	Enables the automatic reconfiguration option in the specified VSAN. The VSAN ID ranges from 1 to 4093.
Step 3	**no fcdomain auto-reconfigure vsan** *vsan-id* **Example:** `switch(config)# no fcdomain auto-reconfigure vsan 1`	Disables the automatic reconfiguration option and reverts it to the factory default in the specified VSAN. The VSAN ID ranges from 1 to 4093.

Domain IDs

Domain IDs uniquely identify a switch in a VSAN. A switch may have different domain IDs in different VSANs. The domain ID is part of the overall FC ID.

Domain IDs - Guidelines

The configured domain ID can be preferred or static. By default, the configured domain ID is 0 (zero) and the configured type is preferred.

Note The 0 (zero) value can be configured only if you use the preferred option.

If you do not configure a domain ID, the local switch sends a random ID in its request. We recommend that you use static domain IDs.

When a subordinate switch requests a domain, the following process takes place (see the figure below):

- The local switch sends a configured domain ID request to the principal switch.
- The principal switch assigns the requested domain ID if available. Otherwise, it assigns another available domain ID.

Figure 18: Configuration Process Using the Preferred Option

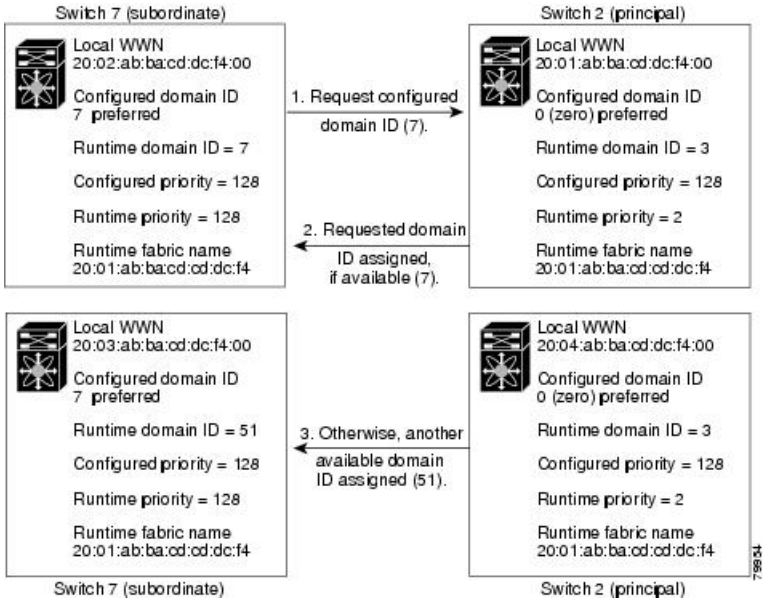

The operation of a subordinate switch changes based on three factors:

- The allowed domain ID lists
- The configured domain ID
- The domain ID that the principal switch has assigned to the requesting switch

In specific situations, the changes are as follows:

- When the received domain ID is not within the allowed list, the requested domain ID becomes the runtime domain ID and all interfaces on that VSAN are isolated.
- When the assigned and requested domain IDs are the same, the preferred and static options are not relevant, and the assigned domain ID becomes the runtime domain ID.
- When the assigned and requested domain IDs are different, the following cases apply:
  - If the configured type is static, the assigned domain ID is discarded, all local interfaces are isolated, and the local switch assigns itself the configured domain ID, which becomes the runtime domain ID.
  - If the configured type is preferred, the local switch accepts the domain ID assigned by the principal switch and the assigned domain ID becomes the runtime domain ID.

If you change the configured domain ID, the change is only accepted if the new domain ID is included in all the allowed domain ID lists currently configured in the VSAN. Alternatively, you can also configure zero-preferred domain ID.

Caution You must enter the fcdomain restart command if you want to apply the configured domain changes to the runtime domain.

Note If you have configured an allow domain ID list, the domain IDs that you add must be in that range for the VSAN.

Related Topics

Allowed Domain ID Lists, on page 99

Configuring Static or Preferred Domain IDs

You can specify a static or preferred domain ID.

SUMMARY STEPS

1. **configure terminal**
2. **fcdomain domain** *domain-id* **static vsan** *vsan-id*
3. **no fcdomain domain** *domain-id* **static vsan** *vsan-id*
4. **fcdomain domain** *domain-id* **preferred vsan** *vsan-id*
5. **no fcdomain domain** *domain-id* **preferred vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**fcdomain domain** *domain-id* **static vsan** *vsan-id* **Example:** `switch(config)# fcdomain domain 1 static vsan 3`	Configures the switch in the specified VSAN to accept only a specific value and moves the local interfaces in the specified VSAN to an isolated state if the requested domain ID is not granted. The domain ID range is 1 to 239. The VSAN ID range is 1 to 4093.
Step 3	**no fcdomain domain** *domain-id* **static vsan** *vsan-id* **Example:** `switch(config)# no fcdomain domain 1 static vsan 3`	Resets the configured domain ID to factory defaults in the specified VSAN. The configured domain ID becomes 0 preferred.
Step 4	**fcdomain domain** *domain-id* **preferred vsan** *vsan-id* **Example:**	Configures the switch in the specified VSAN to request a preferred domain ID 3 and accepts any value assigned by

	Command or Action	Purpose
	`switch(config)# fcdomain domain 1 preferred vsan 5`	the principal switch. The domain ID range is 1 to 239. The VSAN ID range is 1 to 4093.
Step 5	**no fcdomain domain** *domain-id* **preferred vsan** *vsan-id* **Example:** `switch(config)# no fcdomain domain 1 preferred vsan 5`	Resets the configured domain ID to 0 (default) in the specified VSAN. The configured domain ID becomes 0 preferred.

Allowed Domain ID Lists

By default, the valid range for an assigned domain ID list is from 1 to 239. You can specify a list of ranges to be in the allowed domain ID list and separate each range with a comma. The principal switch assigns domain IDs that are available in the locally configured allowed domain list.

Use allowed domain ID lists to design your VSANs with nonoverlapping domain IDs. This helps you in the future if you need to implement IVR without the NAT feature.

If you configure an allowed list on one switch in the fabric, we recommend that you configure the same list in all other switches in the fabric to ensure consistency or use CFS to distribute the configuration.

Configuring Allowed Domain ID Lists

You can configure the allowed domain ID list.

SUMMARY STEPS

1. **configure terminal**
2. **fcdomain allowed** *domain-id range* **vsan** *vsan-id*
3. **no fcdomain allowed** *domain-id range* **vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**fcdomain allowed** *domain-id range* **vsan** *vsan-id* **Example:** `switch(config)# fcdomain allowed 3 vsan 10`	Configures the list to allow switches with the domain ID range in the specified VSAN. The domain ID range is from 1 to 239. The VSAN ID range is from 1 to 4093.
Step 3	**no fcdomain allowed** *domain-id range* **vsan** *vsan-id* **Example:** `switch(config)# no fcdomain allowed 3 vsan 10`	Reverts to the factory default of allowing domain IDs from 1 through 239 in the specified VSAN.

CFS Distribution of Allowed Domain ID Lists

You can enable the distribution of the allowed domain ID list configuration information to all Cisco SAN switches in the fabric using the Cisco Fabric Services (CFS) infrastructure. This feature allows you to synchronize the configuration across the fabric from the console of a single switch. Because the same configuration is distributed to the entire VSAN, you can avoid a possible misconfiguration and the possibility that two switches in the same VSAN have configured incompatible allowed domains.

Use CFS to distribute the allowed domain ID list to ensure consistency in the allowed domain ID lists on all switches in the VSAN.

Note We recommend configuring the allowed domain ID list and committing it on the principal switch.

Enabling Distribution

You can enable (or disable) allowed domain ID list configuration distribution.

CFS distribution of allowed domain ID lists is disabled by default. You must enable distribution on all switches to which you want to distribute the allowed domain ID lists.

Before you begin

CFS has the following prerequisites:

CFS is enabled by default. All devices in the fabric must have CFS enabled, or they do not receive distributions. If CFS is disabled for an application, that application does not distribute any configuration, and it does not accept a distribution from other devices in the fabric. To enable the CFS use **cfs distribute** command.

SUMMARY STEPS

1. **configure terminal**
2. **fcdomain distribute**
3. **no fcdomain distribute**

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**fcdomain distribute** **Example:** `switch(config)# fcdomain distribute`	Enables domain configuration distribution.
Step 3	**no fcdomain distribute** **Example:** `switch(config)# no fcdomain distribute`	Disables (default) domain configuration distribution.

Locking the Fabric

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

- No other user can make any configuration changes to this feature.
- A pending configuration is created by copying the active configuration. Subsequent modifications are made to the pending configuration and remain there until you commit the changes to the active configuration (and other switches in the fabric) or discard them.

Committing Changes

You can commit pending domain configuration changes and release the lock.

To apply the pending domain configuration changes to other SAN switches in the VSAN, you must commit the changes. The pending configuration changes are distributed and, on a successful commit, the configuration changes are applied to the active configuration in the SAN switches throughout the VSAN and the fabric lock is released.

SUMMARY STEPS

1. **configure terminal**
2. **fcdomain commit vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**fcdomain commit vsan** *vsan-id* **Example:** `switch(config)# fcdomain commit vsan 45`	Commits the pending domain configuration changes.

Discarding Changes

You can discard pending domain configuration changes and release the lock.

At any time, you can discard the pending changes to the domain configuration and release the fabric lock. If you discard (abort) the pending changes, the configuration remains unaffected and the lock is released.

SUMMARY STEPS

1. **configure terminal**
2. **fcdomain abort vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**fcdomain abort vsan** *vsan-id* **Example:** `switch(config)# fcdomain abort vsan 30`	Discards the pending domain configuration changes.

Clearing a Fabric Lock

If you have performed a domain configuration task and have not released 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 pending changes are discarded and the fabric lock is released.

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

To release a fabric lock, enter the **clear fcdomain session vsan** command in EXEC mode using a login ID that has administrative privileges:

```
switch# clear fcdomain session vsan 10
```

Displaying CFS Distribution Status

You can display the status of CFS distribution for allowed domain ID lists by using the **show fcdomain status** command:

```
switch# show fcdomain status
CFS distribution is enabled
```

Displaying Pending Changes

You can display the pending configuration changes by using the **show fcdomain pending** command:

```
switch# show fcdomain pending vsan 10
Pending Configured Allowed Domains
---------------------------------
VSAN 10
Assigned or unallowed domain IDs: 1-9,24,100,231-239.
[User] configured allowed domain IDs: 10-230.
```

You can display the differences between the pending configuration and the current configuration by using the **show fcdomain pending-diff** command:

```
switch# show fcdomain pending-diff vsan 10
Current Configured Allowed Domains
---------------------------------
VSAN 10
```

```
Assigned or unallowed domain IDs: 24,100.
[User] configured allowed domain IDs: 1-239.
Pending Configured Allowed Domains
---------------------------------
VSAN 10
Assigned or unallowed domain IDs: 1-9,24,100,231-239.
[User] configured allowed domain IDs: 10-230.
```

Displaying Session Status

You can display the status of the distribution session by using the **show fcdomain session-status vsan** command:

```
switch# show fcdomain session-status vsan 1
Last Action: Distribution Enable
Result: Success
```

Contiguous Domain ID Assignments

By default, the contiguous domain assignment is disabled. When a subordinate switch requests the principal switch for two or more domains and the domains are not contiguous, the following situations can occur:

- If the contiguous domain assignment is enabled in the principal switch, the principal switch locates contiguous domains and assigns them to the subordinate switches. If contiguous domains are not available, the switch software rejects this request.
- If the contiguous domain assignment is disabled in the principal switch, the principal switch assigns the available domains to the subordinate switch.

Enabling Contiguous Domain ID Assignments

You can enable contiguous domains in a specific VSAN (or a range of VSANs).

SUMMARY STEPS

1. **configure terminal**
2. **fcdomain contiguous-allocation vsan** *vsan-id - vsan-id*
3. **no fcdomain contiguous-allocation vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

	Command or Action	Purpose
Step 2	**fcdomain contiguous-allocation vsan** *vsan-id - vsan-id* **Example:** `switch(config)# fcdomain contiguous-allocation vsan 22-30`	Enables the contiguous allocation option in the specified VSAN range. **Note** The **contiguous-allocation** option takes immediate effect at runtime. You do not need to restart the fcdomain.
Step 3	**no fcdomain contiguous-allocation vsan** *vsan-id* **Example:** `switch(config)# no fcdomain contiguous-allocation vsan 7`	Disables the contiguous allocation option and reverts it to the factory default in the specified VSAN.

FC IDs

When an N port logs into a SAN switch, it is assigned an FC ID. By default, the persistent FC ID feature is enabled. If this feature is disabled, the following situations can occur:

- An N port logs into a SAN switch. The WWN of the requesting N port and the assigned FC ID are retained and stored in a volatile cache. The contents of this volatile cache are not saved across reboots.
- The switch is designed to preserve the binding FC ID to the WWN on a best-effort basis. For example, if one N port disconnects from the switch and its FC ID is requested by another device, this request is granted and the WWN with the initial FC ID association is released.
- The volatile cache stores up to 4000 entries of WWN to FC ID binding. If this cache is full, a new (more recent) entry overwrites the oldest entry in the cache. In this case, the corresponding WWN to FC ID association for the oldest entry is lost.
- N ports receive the same FC IDs if disconnected and reconnected to any port within the same switch (as long as it belongs to the same VSAN).

Persistent FC IDs

When persistent FC IDs are enabled, the following occurs:

- The current FC IDs in use in the fcdomain are saved across reboots.
- The fcdomain automatically populates the database with dynamic entries that the switch has learned about after a device (host or disk) is plugged into a port interface.

Note If you connect to the switch from an AIX or HP-UX host, be sure to enable the persistent FC ID feature in the VSAN that connects these hosts.

Note When persistent FC IDs are enabled, FC IDs cannot be changed after a reboot. FC IDs are enabled by default, but can be disabled for each VSAN.

A persistent FC ID assigned to an F port can be moved across interfaces and can continue to maintain the same persistent FC ID.

Enabling the Persistent FC ID Feature

You can enable the persistent FC ID feature.

SUMMARY STEPS

1. **configure terminal**
2. **fcdomain fcid persistent vsan** *vsan-id*
3. **no fcdomain fcid persistent vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**fcdomain fcid persistent vsan** *vsan-id* **Example:** `switch(config)# fcdomain fcid persistent vsan 78`	Activates (default) persistency of FC IDs in the specified VSAN.
Step 3	**no fcdomain fcid persistent vsan** *vsan-id* **Example:** `switch(config)# no fcdomain fcid persistent vsan 33`	Disables the FC ID persistency feature in the specified VSAN.

Persistent FC ID Configuration Guidelines

When the persistent FC ID feature is enabled, you can enter the persistent FC ID submode and add static or dynamic entries in the FC ID database. By default, all added entries are static. Persistent FC IDs are configured on a per-VSAN basis.

When manually configuring a persistent FC ID, follow these requirements:

- Ensure that the persistent FC ID feature is enabled in the required VSAN.
- Ensure that the required VSAN is an active VSAN. Persistent FC IDs can only be configured on active VSANs.
- Verify that the domain part of the FC ID is the same as the runtime domain ID in the required VSAN. If the software detects a domain mismatch, the command is rejected.
- Verify that the port field of the FC ID is 0 (zero) when configuring an area.

Configuring Persistent FC IDs

You can configure persistent FC IDs.

SUMMARY STEPS

1. **configure terminal**
2. **fcdomain fcid database**
3. **vsan** *vsan-id* **wwn 33:e8:00:05:30:00:16:df fcid** *fcid*
4. **vsan** *vsan-id* **wwn 11:22:11:22:33:44:33:44 fcid** *fcid* **dynamic**
5. **vsan** *vsan-id* **wwn 11:22:11:22:33:44:33:44 fcid** *fcid* **area**

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**fcdomain fcid database** **Example:** `switch(config)# fcdomain fcid database`	Enters FC ID database configuration submode.
Step 3	**vsan** *vsan-id* **wwn 33:e8:00:05:30:00:16:df fcid** *fcid* **Example:** `switch(config-fcid-db)# vsan 26 wwn 33:e8:00:05:30:00:16:df fcid 4`	Configures a device WWN (33:e8:00:05:30:00:16:df) with the FC ID 0x070128 in the specified VSAN. **Note** To avoid assigning a duplicate FC ID, use the **show fcdomain address-allocation vsan** command to display the FC IDs in use.
Step 4	**vsan** *vsan-id* **wwn 11:22:11:22:33:44:33:44 fcid** *fcid* **dynamic** **Example:** `switch(config-fcid-db)# vsan 13 wwn 11:22:11:22:33:44:33:44 fcid 6 dynamic`	Configures a device WWN (11:22:11:22:33:44:33:44) with the FC ID 0x070123 in the specified VSAN in dynamic mode.
Step 5	**vsan** *vsan-id* **wwn 11:22:11:22:33:44:33:44 fcid** *fcid* **area** **Example:** `switch(config-fcid-db)# vsan 88 wwn 11:22:11:22:33:44:33:44 fcid 4 area`	Configures a device WWN (11:22:11:22:33:44:33:44) with the FC IDs 0x070100 through 0x701FF in the specified VSAN. **Note** To secure the entire area for this fcdomain, assign 00 as the last two characters of the FC ID.

Unique Area FC IDs for HBAs

Note Read this section only if the Host Bus Adapter (HBA) port and the storage port are connected to the same switch.

Some HBA ports require a different area ID than for the storage ports when they are both connected to the same switch. For example, if the storage port FC ID is 0x6f7704, the area for this port is 77. In this case, the HBA port's area can be anything other than 77. The HBA port's FC ID must be manually configured to be different from the storage port's FC ID.

Cisco SAN switches facilitate this requirement with the FC ID persistence feature. You can use this feature to preassign an FC ID with a different area to either the storage port or the HBA port.

Configuring Unique Area FC IDs for an HBA

You can configure a different area ID for the HBA port.

The following task uses an example configuration with a switch domain of 111(6f hex). The server connects to the switch over FCoE. The HBA port connects to interface vfc20.

Step 1 Obtain the port WWN (Port Name field) ID of the HBA using the **show flogi database** command.

```
switch# show flogi database
--------------------------------------------------------------------------------
INTERFACE VSAN  FCID           PORT NAME                NODE NAME
--------------------------------------------------------------------------------
 vfc20     3   0x6f7703  50:05:08:b2:00:71:c8:c2   50:05:08:b2:00:71:c8:c0
```

Step 2 Shut down the HBA interface in the SAN switch.

```
switch# configure terminal
switch(config)# interface vfc 20
switch(config-if)# shutdown
switch(config-if)# end
```

Step 3 Verify that the FC ID feature is enabled using the **show fcdomain vsan** command.

```
switch# show fcdomain vsan 3
...
Local switch configuration information:
        State: Enabled
        FCID persistence: Disabled
```

If this feature is disabled, continue to the next step to enable the persistent FC ID.

If this feature is already enabled, skip to the following step.

Step 4 Enable the persistent FC ID feature in the SAN switch.

```
switch# configure terminal
switch(config)# fcdomain fcid persistent vsan 3
switch(config)# end
```

Step 5 Assign a new FC ID with a different area allocation. In this example, replace *77* with *ee*.

```
switch# configure terminal
switch(config)# fcdomain fcid database
```

```
switch(config-fcid-db)# vsan 3 wwn 50:05:08:b2:00:71:c8:c2
fcid 0x6fee00 area
```

Step 6 Enable the HBA interface in the SAN switch.

```
switch# configure terminal
switch(config)# interface vfc 20
switch(config-if)# no shutdown
switch(config-if)# end
```

Step 7 Verify the pWWN ID of the HBA by using the **show flogi database** command.

```
switch# show flogi database
--------------------------------------------------------------------------------
INTERFACE VSAN   FCID      PORT NAME                 NODE NAME
--------------------------------------------------------------------------------
 vfc20    3    0x6fee00    50:05:08:b2:00:71:c8:c2   50:05:08:b2:00:71:c8:c0
```

Persistent FC ID Selective Purging

Persistent FC IDs can be purged selectively. Static entries and FC IDs currently in use cannot be deleted. The table below identifies the FC ID entries that are deleted or retained when persistent FC IDs are purged.

Table 14: Purged FC IDs

Persistent FC ID state	Persistent Usage State	Action
Static	In use	Not deleted
Static	Not in use	Not deleted
Dynamic	In use	Not deleted
Dynamic	Not in use	Deleted

Purging Persistent FC IDs

You can purge persistent FC IDs.

SUMMARY STEPS

1. **purge fcdomain fcid vsan** *vsan-id*
2. **purge fcdomain fcid vsan** *vsan-id* - *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	**purge fcdomain fcid vsan** *vsan-id* **Example:** `switch# purge fcdomain fcid vsan 667`	Purges all dynamic and unused FC IDs in the specified VSAN.
Step 2	**purge fcdomain fcid vsan** *vsan-id* **-** *vsan-id* **Example:** `switch# purge fcdomain fcid vsan 50-100`	Purges dynamic and unused FC IDs in the specified VSAN range.

Verifying the fcdomain Configuration

Note If the fcdomain feature is disabled, the runtime fabric name in the display is the same as the configured fabric name.

This example shows how to display information about fcdomain configurations:

```
switch# show fcdomain vsan 2
```

Use the **show fcdomain domain-list** command to display the list of domain IDs of all switches belonging to a specified VSAN. This list provides the WWN of the switches owning each domain ID. The next example uses the following values:

- A switch with WWN of 20:01:00:05:30:00:47:df is the principal switch and has domain 200.
- A switch with WWN of 20:01:00:0d:ec:08:60:c1 is the local switch (the one where you typed the CLI command to show the domain-list) and has domain 99.
- The IVR manager obtained virtual domain 97 using 20:01:00:05:30:00:47:df as the WWN for a virtual switch.

```
switch# show fcdomain domain-list vsan 76
Number of domains: 3
Domain ID              WWN
---------    -----------------------
0xc8(200)    20:01:00:05:30:00:47:df [Principal]
 0x63(99)    20:01:00:0d:ec:08:60:c1 [Local]
 0x61(97)    50:00:53:0f:ff:f0:10:06 [Virtual (IVR)]
```

Use the **show fcdomain allowed vsan** command to display the list of allowed domain IDs configured on this switch..

```
switch# show fcdomain allowed vsan 1
Assigned or unallowed domain IDs: 1-96,100,111-239.
[Interoperability Mode 1] allowed domain IDs: 97-127.
[User] configured allowed domain IDs: 50-110.
```

Ensure that the requested domain ID passes the switch software checks, if interop 1 mode is required in this switch.

The following example shows how to display all existing, persistent FC IDs for a specified VSAN. You can also specify the unused option to view only persistent FC IDs that are still not in use.

```
switch# show fcdomain fcid persistent vsan 1000
```

The following example shows how to display frame and other fcdomain statistics for a specified VSAN or SAN port channel:

```
switch# show fcdomain statistics vsan 1
VSAN Statistics
        Number of Principal Switch Selections: 5
        Number of times Local Switch was Principal: 0
        Number of 'Build Fabric's: 3
        Number of 'Fabric Reconfigurations': 0
```

The following example shows how to display FC ID allocation statistics including a list of assigned and free FC IDs:

```
switch# show fcdomain address-allocation vsan 1
```

The following example shows how to display the valid address allocation cache. The cache is used by the principal switch to reassign the FC IDs for a device (disk or host) that exited and reentered the fabric. In the cache content, VSAN refers to the VSAN that contains the device, WWN refers to the device that owned the FC IDs, and mask refers to a single or entire area of FC IDs.

```
switch# show fcdomain address-allocation cache
```

Default Settings for Fibre Channel Domains

The following table lists the default settings for all fcdomain parameters.

Table 15: Default fcdomain Parameters

Parameters	Default
fcdomain feature	Enabled
Configured domain ID	0 (zero)
Configured domain	Preferred
auto-reconfigure option	Disabled
contiguous-allocation option	Disabled
Priority	128
Allowed list	1 to 239
Fabric name	20:01:00:05:30:00:28:df
rcf-reject	Disabled

Parameters	Default
Persistent FC ID	Enabled
Allowed domain ID list configuration distribution	Disabled

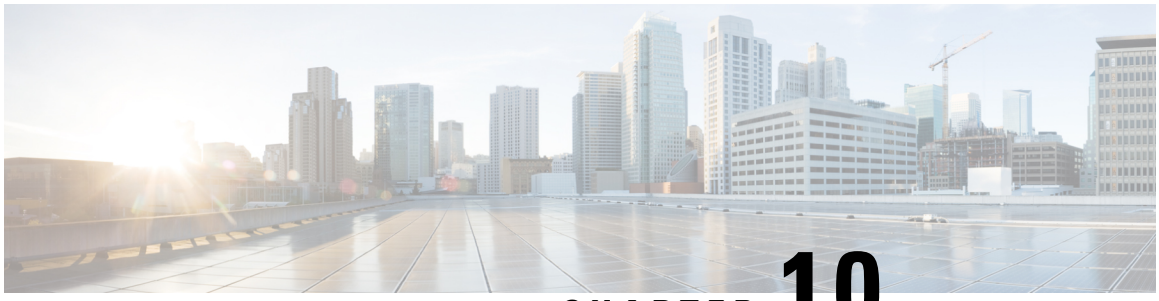

CHAPTER 10

Configuring FCoE VLANs and Virtual Interfaces

This chapter contains the following sections:

- Information About Virtual Interfaces, on page 113
- Guidelines and Limitations for FCoE VLANs and Virtual Interfaces, on page 113
- Configuring Virtual Interfaces, on page 115
- Verifying the Virtual Interface , on page 120
- Mapping VSANs to VLANs Example Configuration , on page 123

Information About Virtual Interfaces

Cisco Nexus devices support Fibre Channel over Ethernet (FCoE), which allows Fibre Channel and Ethernet traffic to be carried on the same physical Ethernet connection between the switch and the servers.

The Fibre Channel portion of FCoE is configured as a virtual Fibre Channel interface. Logical Fibre Channel features (such as interface mode) can be configured on virtual Fibre Channel interfaces.

Guidelines and Limitations for FCoE VLANs and Virtual Interfaces

FCoE VLANs and Virtual Fiber Channel (vFC) interfaces have these guidelines and limitations:

- Each vFC interface must be bound to an FCoE-enabled Ethernet or EtherChannel interface or to the MAC address of a remotely connected adapter. FCoE is supported on 10-Gigabit and 40-Gigabit Ethernet interfaces.
- A virtual Fibre Channel interface must be bound to an interface before it can be used. The binding is to a physical Ethernet interface (when the converged network adapter (CNA) is directly connected to the Cisco Nexus device), a MAC address (when the CNA is remotely connected over a Layer 2 bridge), or an EtherChannel.
- The Ethernet or EtherChannel interface that you bind to the vFC interface must be configured as follows:
  - The Ethernet or EtherChannel interface must be a trunk port (use the **switchport mode trunk** command).
  - The FCoE VLAN that corresponds to a vFC's VSAN must be in the allowed VLAN list.

- Set the MTU 9216 and QoS polices to the interface. You can use default (service-policy type qos input default-fcoe-in-policy) or custom QoS policies.
- You must not configure an FCoE VLAN as the native VLAN of the trunk port.

Note The native VLAN is the default VLAN on a trunk. Any untagged frames transit the trunk as native VLAN traffic.

- You should use an FCoE VLAN only for FCoE.
- Do not use the default VLAN, VLAN1, as an FCoE VLAN.
- You must configure the Ethernet interface as PortFast (use the **spanning-tree port type edge trunk** command).

Note You are not required to configure trunking on the server interface even if the switch interface is configured with trunking enabled. All non-FCoE traffic from the server is passed on the native VLAN.

- The vFC interface can be bound to Ethernet port channels with multiple member ports connected to FCoE Initialization Protocol (FIP) snooping bridges.
- Each vFC interface is associated with only one VSAN.
- You must map any VSAN with associated vFC interfaces to a dedicated FCoE-enabled VLAN.
- FCoE is not supported on private VLANs.
- If the converged access switches (in the same SAN fabric or in another) need to be connected to each other over Ethernet links for a LAN alternate path, then you must explicitly configure such links to exclude all FCoE VLANs from membership.
- You must use separate FCoE VLANs for FCoE in SAN-A and SAN-B fabrics.
- FCoE connectivity to pre-FIP CNAs over virtual port channels (vPCs) is not supported.

Note Virtual interfaces are created with the administrative state set to down. You must explicitly configure the administrative state to bring the virtual interface into operation.

Configuring Virtual Interfaces

Mapping a VSAN to a VLAN

A unique, dedicated VLAN must be configured at every converged access switch to carry traffic for each VSAN in the SAN (for example, VLAN 1002 for VSAN 1, VLAN 1003 for VSAN 2, and so on). If Multiple Spanning Tree (MST) is enabled, a separate MST instance must be used for FCoE VLANs.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **vlan** *vlan-id*
3. switch(config-vlan)# **fcoe** [**vsan** *vsan-id*]
4. switch(config-vlan)# **exit**
5. (Optional) switch(config)# **show vlan fcoe**
6. (Optional) switch(config-if)# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vlan** *vlan-id*	Enters VLAN configuration mode. The VLAN number range is from 1 to 4096.
Step 3	switch(config-vlan)# **fcoe** [**vsan** *vsan-id*]	Enables FCoE for the specified VLAN. If you do not specify a VSAN number, a mapping is created from this VLAN to the VSAN with the same number. Configures the mapping from this VLAN to the specified VSAN.
Step 4	switch(config-vlan)# **exit**	Exits VLAN configuration mode. You must exit this mode to execute the configured commands on your Cisco Nexus device.
Step 5	(Optional) switch(config)# **show vlan fcoe**	Displays information about the FCoE configuration for a VLAN.
Step 6	(Optional) switch(config-if)# **copy running-config startup-config**	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to map VLAN 200 to VSAN 2:

```
switch(config)# vlan 200
```

```
switch(config-vlan)# fcoe vsan 2
```

Creating a Virtual Fibre Channel Interface

You can create a virtual Fibre Channel interface. You must bind the virtual Fibre Channel interface to a physical interface before it can be used.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface vfc** *vfc-id*
3. switch(config-if)# **bind** {**interface** {**ethernet** *slot*/*port* | **port-channel** *channel-number*} | **mac-address** *MAC-address*}
4. (Optional) switch(config-if)# **no bind** {**interface** {**ethernet** *slot*/*port* | **port-channel** *channel-number*} | **mac-address** *MAC-address*}
5. (Optional) switch(config)# **no interface vfc** *vfc-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface vfc** *vfc-id*	Creates a virtual Fibre Channel interface (if it does not already exist) and enters interface configuration mode. The virtual Fibre Channel interface ID range is from 1 to 8192.
Step 3	switch(config-if)# **bind** {**interface** {**ethernet** *slot*/*port* | **port-channel** *channel-number*} | **mac-address** *MAC-address*}	Binds the virtual Fibre Channel interface to the specified interface.
Step 4	(Optional) switch(config-if)# **no bind** {**interface** {**ethernet** *slot*/*port* | **port-channel** *channel-number*} | **mac-address** *MAC-address*}	Unbinds the virtual Fibre Channel interface from the specified interface.
Step 5	(Optional) switch(config)# **no interface vfc** *vfc-id*	Deletes a virtual Fibre Channel interface.

Example

This example shows how to bind a virtual Fibre Channel interface to an Ethernet interface:

```
switch# configure terminal
switch(config)# interface vfc 4
switch(config-if)# bind interface ethernet 1/4
```

This example shows how to bind a virtual Fibre Channel interface to port-channel.:

```
switch# configure terminal
switch(config)# interface vfc 3
switch(config-if)# bind interface port-channel 1
```

This example shows how to bind a virtual Fibre Channel interface to a MAC address:

```
switch# configure terminal
switch(config)# interface vfc 2
switch(config-if)# bind mac-address 00:0a:00:00:00:36
```

This example shows how to delete a virtual Fibre Channel interface:

```
switch# configure terminal
switch(config)# no interface vfc 4
```

This example shows how to unbind a virtual Fibre Channel interface from an ethernet interface:

```
switch# configure terminal
switch(config)# int vfc17
switch(config-if)# no bind interface ethernet 1/17
switch(config-if)# exit
```

Associating a Virtual Fibre Channel Interface to a VSAN

A unique, dedicated VLAN must be configured at every converged access switch to carry traffic for each Virtual Fabric (VSAN) in the SAN (for example, VLAN 1002 for VSAN 1, VLAN 1003 for VSAN 2, and so on). If MST is enabled, a separate MST instance must be used for FCoE VLANs.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **vsan database**
3. switch(config-vsan)# **vsan** *vsan-id* **interface vfc** *vfc-id*
4. (Optional) switch(config-vsan)# **no vsan** *vsan-id* **interface vfc** *vfc-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **vsan database**	Enters VSAN configuration mode.
Step 3	switch(config-vsan)# **vsan** *vsan-id* **interface vfc** *vfc-id*	Configures the association between the VSAN and virtual Fibre Channel interface. The VSAN number must map to a VLAN on the physical Ethernet interface that is bound to the virtual Fibre Channel interface.
Step 4	(Optional) switch(config-vsan)# **no vsan** *vsan-id* **interface vfc** *vfc-id*	Disassociates the connection between the VSAN and virtual Fibre Channel interface.

Example

This example shows how to associate a virtual Fibre Channel interface to a VSAN:

```
switch# configure terminal
```

```
switch(config)# vsan database
switch(config-vsan)# vsan 2 interface vfc 4
```

Creating an Implicit Virtual Fibre Channel Port Channel Interface

You can create a virtual Fibre Channel (vFC), and implicitly bind it to an Ethernet interface or a port-channel using a single command. For this, the vFC identifier must match the Ethernet interface or port-channel identifier. The Ethernet interface can be a module (slot or port) interface (slot/QSFP-module/port).

Configuring virtual Fibre Channel Interface

Before you begin

- Ensure you have installed the correct license for FCoE.
- Ensure you have enabled FCoE.

Step 1 Enter global configuration mode:

switch# **configure terminal**

Step 2 Create a VFC (if it does not already exist):

Additionally, *vfc slot/port* binds the vFC to an Ethernet *slot/port* interface. The vFC *slot/QSFP-module/port* binds the vFC to a breakout interface.

switch(config) # **interface vfc** {id | *slot/port* | *slot/QSFP-module/port* }

Step 3 Bring up the vFC interface:

switch(config-if) # **no shutdown**

Step 4 Required: Exit the interface configuration mode:

switch(config-if) # **exit**

Configuring virtual Fibre Channel Interface

This example shows how to implicitly bind a virtual Fibre Channel interface to an Ethernet interface:

```
switch# configure terminal
switch(config)# interface eth1/11
switch(config-if)# switchport mode trunk
switch(config-if)# mtu 9216
switch(config-if)# service-policy type qos input default-fcoe-in-policy
switch(config-if)# no shutdown

switch(config)# interface vfc1/11
switch(config-if)# no shutdown
switch(config-if)# exit
switch(config)#

switch(config)# vsan database
switch(config-vsan-db)# vsan 10
```

```
switch(config-vsan-db)# exit
switch(config)#

switch(config)# vlan 10
switch(config-vlan)# fcoe vsan 10
switch(config-vlan)# exit
switch(config)#

switch(config)# vsan database
switch(config-vsan-db)# vsan 10 interface vfc1/11
switch(config-vsan-db)# exit
switch(config)#
switch(config)# show interface vfc1/11
vfc1/11 is trunking (Not all VSANs UP on the trunk)
Bound interface is Ethernet1/11
Hardware is Ethernet
Port WWN is 20:0b:00:de:fb:9d:0e:a0
Admin port mode is F, trunk mode is on
snmp link state traps are enabled
Port mode is TF
Port vsan is 10
Operating Speed is 10 Gbps
Admin Speed is auto
Trunk vsans (admin allowed and active) (1,10)
Trunk vsans (up) (10)
Trunk vsans (isolated) ()
Trunk vsans (initializing) (1)
11 fcoe in packets
1692 fcoe in octets
0 fcoe out packets
0 fcoe out octets
Interface last changed at Mon Dec 16 09:03:33 2019

switch(config)#
```

Configuring virtual Fibre Channel – Port Channel Interface

Step 1 Enter global configuration mode:

switch# **configure terminal**

Step 2 Create a vFC that implicitly binds to the Ethernet port-channel based on its number:

The port number range is from 1 to 4096.

switch(config) # **interface vfc-port-channel** *port number*

Step 3 Bring up the vFC port:

switch(config-if) # **no shutdown**

Step 4 Required: Exit from the current interface configuration mode:

switch(config-if) # **exit**

Configuring virtual Fibre Channel - Port Channel Interface

The example shows how you can create a vFC-port-channel that implicitly binds to Ethernet port-channel:

```
switch# configure terminal
switch(config)# interface port-channel 10
switch(config-if)# switchport
switch(config-if)# switchport mode trunk
switch(config-if)# mtu 9216
switch(config-if)# service-policy type qos input default-fcoe-in-policy
switch(config-if)# no shutdown
switch(config-if)# exit

switch(config)# interface eth1/49
switch(config-if)# channel-group 10 force
switch(config-if)# no shutdown
switch(config-if)# exit

switch# configure terminal
switch(config)# interface vfc-port-channel 10
switch(config-if)# no shutdown
switch(config-if)# exit

switch(config)# vlan 10
switch(config-vlan)# fcoe vsan 10
switch(config-vlan)# exit
switch(config)#

switch(config)# vsan database
switch(config-vsan-db)# vsan 10 interface vfc-port-channel 10
switch(config-vsan-db)# exit

switch(config)# show interface vfc-port-channel 10
vfc-po10 is trunking (Not all VSANs UP on the trunk)
Bound interface is port-channel10
Hardware is Ethernet
Port WWN is 25:1b:00:de:fb:9d:0e:a0
Admin port mode is F, trunk mode is on
snmp link state traps are enabled
Port mode is TF
Port vsan is 10
Operating Speed is 40 Gbps
Admin Speed is auto
Trunk vsans (admin allowed and active) (1,10)
Trunk vsans (up) (10)
Trunk vsans (isolated) ()
Trunk vsans (initializing) (1)
11 fcoe in packets
1236 fcoe in octets
0 fcoe out packets
0 fcoe out octets
Interface last changed at Mon Dec 16 08:56:13 2019
```

Verifying the Virtual Interface

To display configuration information about virtual interfaces, perform one of the following tasks:

Command	Purpose
switch# **show interface vfc** *vfc-id*	Displays the detailed configuration of the specified Fibre Channel interface.
switch# **show interface brief**	Displays the status of all interfaces.
switch# **show vlan fcoe**	Displays the mapping of FCoE VLANs to VSANs.

This example shows how to display a virtual Fibre Channel interface bound to an Ethernet interface:

```
switch# show interface vfc 11

vfc11 is trunking (Not all VSANs UP on the trunk)


Bound interface is Ethernet1/11
Hardware is Ethernet
Port WWN is 20:0a:00:de:fb:9d:0e:df
Admin port mode is F, trunk mode is on
snmp link state traps are enabled
Port mode is TF
Port vsan is 10
Operating Speed is 10 Gbps
Admin Speed is auto
Trunk vsans (admin allowed and active) (1,10)
Trunk vsans (up) (10)
Trunk vsans (isolated) ()
Trunk vsans (initializing) (1)
2 fcoe in packets
152 fcoe in octets
0 fcoe out packets
0 fcoe out octets
Interface last changed at Wed Dec 18 10:36:58 2019
```

This example shows how to display a virtual Fibre Channel interface bound to a MAC address:

```
switch# show interface vfc 11


vfc11 is trunking (Not all VSANs UP on the trunk)
Bound MAC is 0090.faf8.7513
Hardware is Ethernet
Port WWN is 20:0a:00:de:fb:9d:0e:df
Admin port mode is F, trunk mode is on
snmp link state traps are enabled
Port mode is TF
Port vsan is 10
Operating Speed is 10 Gbps
Admin Speed is auto
Trunk vsans (admin allowed and active) (1,10)
Trunk vsans (up) (10)
Trunk vsans (isolated) ()
Trunk vsans (initializing) (1)
3 fcoe in packets
228 fcoe in octets
0 fcoe out packets
0 fcoe out octets
Interface last changed at Mon Dec 16 09:09:02 2019
```

This example shows how to display the status of all the interfaces on the switch (some output has been removed for brevity):

```
switch# show interface brief

--------------------------------------------------------------------------------
Port   VRF          Status IP Address                              Speed    MTU
--------------------------------------------------------------------------------
mgmt0  --           up     10.127.119.70                           1000    1500

--------------------------------------------------------------------------------
Ethernet        VLAN    Type Mode   Status  Reason                 Speed     Port

Interface                                                                    Ch#
--------------------------------------------------------------------------------
Eth1/1          --      eth  routed up      none                    10G(D) --
Eth1/2          --      eth  routed down    Administratively down   auto(D) --
Eth1/3          --      eth  routed down    Link not connected      auto(D) --
Eth1/4          --      eth  routed down    Link not connected      auto(D) --
Eth1/5          --      eth  routed down    XCVR not inserted       auto(D) --
Eth1/6          --      eth  routed down    XCVR not inserted       auto(D) --
Eth1/7          --      eth  routed up      none                    10G(D) --
Eth1/8          --      eth  routed down    Link not connected      auto(D) --
Eth1/9          --      eth  routed down    XCVR not inserted       auto(D) --
Eth1/10         --      eth  routed down    Link not connected      auto(D) --
Eth1/11         1       eth  trunk  up      none                    10G(D) --
Eth1/12         --      eth  routed down    XCVR not inserted       auto(D) --
Eth1/49         1       eth  trunk  up      none                    40G(D) 10
Eth1/50         1       eth  trunk  up      none                    40G(D) --
Eth1/51         --      eth  routed down    XCVR not inserted       auto(D) --
Eth1/52         --      eth  routed down    XCVR not inserted       auto(D) --
Eth1/53         --      eth  routed down    XCVR not inserted       auto(D) --
Eth1/54         --      eth  routed down    XCVR not inserted       auto(D) --

--------------------------------------------------------------------------------
Port-channel VLAN    Type Mode   Status  Reason   Speed   Protocol Interface

--------------------------------------------------------------------------------

Po10         1       eth  trunk  up      none     a-40G(D)  none

--------------------------------------------------------------------------------
Interface  Vsan   Admin  Admin   Status       SFP    Oper  Oper   Port
                  Mode   Trunk                       Mode  Speed  Channel
                         Mode                               (Gbps)
--------------------------------------------------------------------------------
fc1/13      1      auto   on      sfpAbsent    --     --    --     --
fc1/14      1      auto   on      sfpAbsent    --     --    --     --
fc1/15      1      auto   on      sfpAbsent    --     --    --     --
fc1/16      1      auto   on      sfpAbsent    --     --    --     --
fc1/17      1      auto   on      trunking     swl    TE    32     --
fc1/18      1      auto   on      trunking     swl    TE    32     --
fc1/19      1      auto   on      trunking     swl    TE    32     --
fc1/20      1      auto   on      trunking     swl    TE    32     --
fc1/21      1      auto   on      sfpAbsent    --     --    --     --
fc1/22      1      auto   on      sfpAbsent    --     --    --     --
fc1/23      1      auto   on      notConnected swl    --    --     --
fc1/24      1      auto   on      notConnected swl    --    --     --
fc1/25      1      auto   on      sfpAbsent    --     --    --     --
fc1/26      1      auto   on      sfpAbsent    --     --    --     --
fc1/27      1      auto   on      sfpAbsent    --     --    --     --
fc1/28      1      auto   on      sfpAbsent    --     --    --     --
fc1/29      1      auto   on      up           swl    F     16     --
fc1/30      1      auto   on      errDisabled  --     --    --     --
fc1/31      1      auto   on      up           swl    F     32     --
fc1/32      1      auto   on      notConnected swl    --    --     --
fc1/33      1      auto   on      sfpAbsent    --     --    --     --
```

```
fc1/34      1      auto   on      errDisabled  h10g-  --     --     --
fc1/35      1      auto   on      sfpAbsent    --     --     --     --
fc1/36      1      auto   on      sfpAbsent    --     --     --     --
fc1/37      1      auto   on      sfpAbsent    --     --     --     --
fc1/38      1      auto   on      sfpAbsent    --     --     --     --
fc1/39      1      auto   on      sfpAbsent    --     --     --     --
fc1/40      1      auto   on      sfpAbsent    --     --     --     --
fc1/41      1      auto   on      sfpAbsent    --     --     --     --
fc1/42      1      auto   on      sfpAbsent    --     --     --     --
fc1/43      1      auto   on      sfpAbsent    --     --     --     --
fc1/44      1      auto   on      sfpAbsent    --     --     --     --
fc1/45      1      auto   on      errDisabled  --     --     --     --
fc1/46      1      auto   on      errDisabled  --     --     --     --
fc1/47      1      auto   on      errDisabled  --     --     --     --
fc1/48      1      auto   on      errDisabled  --     --     --     --

--------------------------------------------------------------------------------
Interface  Vsan   Admin  Admin   Status       Bind             Oper    Oper
                  Mode   Trunk                Info              Mode    Speed
                         Mode                                          (Gbps)
--------------------------------------------------------------------------------
vfc11       10     F      on      trunking     0090.faf8.7513   TF     10
vfc1/50     10     F      on      trunking     Eth1/50          TF     40
```

This example shows how to display the mapping between the VLANs and VSANs on the switch:

```
switch# show vlan fcoe

VLAN      VSAN      Status

--------  --------  --------

15        15        Operational

20        20        Operational

25        25        Operational

30        30        Non-operational
```

Mapping VSANs to VLANs Example Configuration

The following example shows how to configure the FCoE VLAN and a virtual Fibre Channel interface:

SUMMARY STEPS

1. Enable the associated VLAN and map the VLAN to a VSAN.
2. Configure the VLAN on a physical Ethernet interface.
3. Create a virtual Fibre Channel interface and bind it to a physical Ethernet interface.
4. Associate the virtual Fibre Channel interface to the VSAN.
5. (Optional) Display membership information for the VSAN.
6. (Optional) Display the interface information for the virtual Fibre Channel interface.

DETAILED STEPS

Step 1 Enable the associated VLAN and map the VLAN to a VSAN.

```
switch(config)# vlan 200
```

```
switch(config-vlan)# fcoe vsan 2
switch(config-vlan)# exit
```

Step 2 Configure the VLAN on a physical Ethernet interface.

```
switch(config)# interface eth1/11
switch(config)# spanning-tree port type edge trunk
switch(config-if)# switchport mode trunk
switch(config-if)# switchport trunk allowed vlan 1,200
switch(config-if)# mtu 9216
switch(config-if)# service-policy type qos input default-fcoe-in-policy
switch(config-if)# exit
```

Step 3 Create a virtual Fibre Channel interface and bind it to a physical Ethernet interface.

```
switch(config)# interface vfc 11
switch(config-if)# bind interface ethernet 1/4
switch(config-if)# no shutdown
switch(config-if)# exit
```

Note By default, all virtual Fibre Channel interfaces reside on VSAN 1. If the VLAN to VSAN mapping is to a VSAN other than VSAN 1, then proceed to Step 4.

Step 4 Associate the virtual Fibre Channel interface to the VSAN.

```
switch(config)# vsan database
switch(config-vsan-db)# vsan 2
switch(config-vsan-db)# vsan 2 interface vfc 11
switch(config-vsan)# exit
```

Step 5 (Optional) Display membership information for the VSAN.

```
switch# show vsan 2 membership
vsan 2 interfaces
          vfc 11
```

Step 6 (Optional) Display the interface information for the virtual Fibre Channel interface.

```
switch# show interface vfc 11

vfc11 is trunking (Not all VSANs UP on the trunk)
Bound interface is Ethernet1/11
Hardware is Ethernet
Port WWN is 20:0a:00:de:fb:9d:0e:df
Admin port mode is F, trunk mode is on
snmp link state traps are enabled
Port mode is TF
Port vsan is 2
Operating Speed is 10 Gbps
Admin Speed is auto
Trunk vsans (admin allowed and active) (1-2,10)
Trunk vsans (up) (2)
```

```
Trunk vsans (isolated) ()
Trunk vsans (initializing) (1,10)
2 fcoe in packets
152 fcoe in octets
0 fcoe out packets
0 fcoe out octets
Interface last changed at Mon Dec 16 09:22:25 2019
```

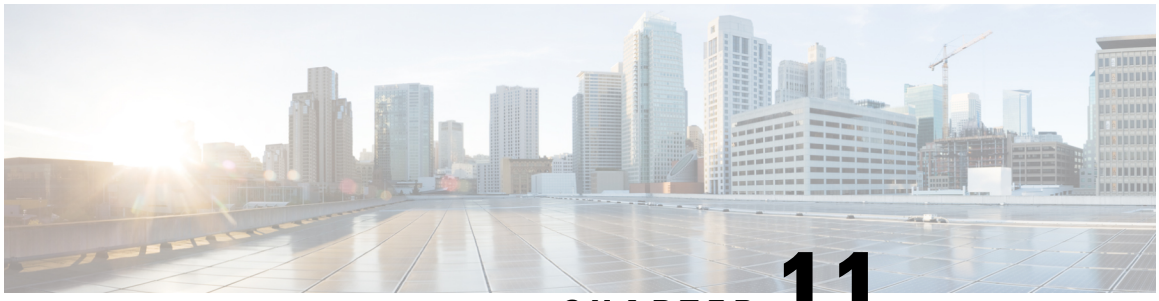

CHAPTER 11

Managing FLOGI, Name Server, and RSCN Databases

This chapter describes how to configure and manage FLOGI, name server and RSCN databases.

This chapter includes the following sections:

- Managing FLOGI, Name Server and RSCN Databases, on page 127

Managing FLOGI, Name Server and RSCN Databases

Fabric Login

In a Fibre Channel fabric, each host or disk requires an FC ID. Use the **show flogi** command to verify if a storage device is displayed in the fabric login (FLOGI) table as in the following examples. 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. The maximum number of FLOGIs or FDISCs per port is 256 and the maximum number of FLOGIs or FDISC per switch = 1000.

This example shows how to verify the storage devices in the fabric login (FLOGI) table:

```
switch# show flogi database
--------------------------------------------------------------------------------
INTERFACE  VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------------
fc2/3      1     0xb200e2  21:00:00:04:cf:27:25:2c  20:00:00:04:cf:27:25:2c
fc2/3      1     0xb200e1  21:00:00:04:cf:4c:18:61  20:00:00:04:cf:4c:18:61
fc2/3      1     0xb200d1  21:00:00:04:cf:4c:18:64  20:00:00:04:cf:4c:18:64
fc2/3      1     0xb200ce  21:00:00:04:cf:4c:16:fb  20:00:00:04:cf:4c:16:fb
fc2/3      1     0xb200cd  21:00:00:04:cf:4c:18:f7  20:00:00:04:cf:4c:18:f7
vfc3/1     2     0xb30100  10:00:00:05:30:00:49:63  20:00:00:05:30:00:49:5e
Total number of flogi = 6.
```

This example shows how to verify the storage devices attached to a specific interface:

```
switch# show flogi database interface vfc1/1

INTERFACE  VSAN    FCID           PORT NAME               NODE NAME

--------------------------------------------------------------------------------

vfc1/1     1     0x870000  20:00:00:1b:21:06:58:bc  10:00:00:1b:21:06:58:bc

Total number of flogi = 1.
```

This example shows how to verify the storage devices associated with VSAN 1:

```
switch# show flogi database vsan 1

show flogi database vsan 1
--------------------------------------------------------------------------------
INTERFACE VSAN FCID PORT NAME NODE NAME
--------------------------------------------------------------------------------
fc1/17 1 0xee0000 21:00:00:24:ff:17:08:2e 20:00:00:24:ff:17:08:2e
fc1/18 1 0xee0020 10:00:00:90:fa:dc:0f:08 20:00:00:90:fa:dc:0f:08
fc1/37 1 0xee00ef 50:06:01:6a:08:60:7c:67 50:06:01:60:88:60:7c:67
Total number of flogi = 3.
```

Name Server Proxy

The name server functionality maintains a database that contains 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 need to modify (update or delete) the contents of a database entry that was previously registered by a different device.

All name server registration requests come from the same port whose parameter is registered or changed. If it does not, then the request is rejected.

This authorization enables WWNs to register specific parameters for another node.

About Registering Name Server Proxies

All name server registration requests come from the same port whose parameter is registered or changed. If it does not, then the request is rejected.

This authorization enables WWNs to register specific parameters for another node.

Registering Name Server Proxies

You can register the name server proxy.

SUMMARY STEPS

1. **configure terminal**
2. **fcns proxy-port** *wwn-id* **vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**fcns proxy-port** *wwn-id* **vsan** *vsan-id* **Example:** `switch(config)# fcns proxy-port 11:22:11:22:33:44:33:44 vsan 300`	Configures a proxy port for the specified VSAN.

Rejecting Duplicate pWWNs

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, same fcdomain.

By default, any future flogi (with duplicate pwwn) on different switch in the same vsan, will be rejected and earlier 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,will be allowed to succeed by deleting earlier FCNS entry.

To reject duplicate pWWNs, follow these steps:

SUMMARY STEPS

1. **configure terminal**
2. **fcns reject-duplicate-pwwn vsan** *vsan-id*
3. **no fcns reject-duplicate-pwwn vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**fcns reject-duplicate-pwwn vsan** *vsan-id* **Example:** `switch(config)# fcns reject-duplicate-pwwn vsan 100`	Any future flogi (with duplicate pwwn) on different switch, will be rejected and earlier FLOGI retained (default).
Step 3	**no fcns reject-duplicate-pwwn vsan** *vsan-id* **Example:** `switch(config)# no fcns reject-duplicate-pwwn vsan 256`	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.

Displaying Name Server Database Entries

This example shows how to display the name server database for all VSANs:

```
switch# show fcns database

VSAN 1:
```

```
--------------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0xe90000    N     20:00:00:6b:f1:70:08:ec (Cisco)         scsi-fcp:init fc-gs
0xec0020    N     21:00:00:24:ff:7f:37:05 (Company A)     scsi-fcp:target
0xec0040    N     50:08:01:60:01:59:49:33                 scsi-fcp:init
0xec0060    N     20:12:00:11:0d:9d:06:00                 scsi-fcp:init
0xec0080    N     50:08:01:60:08:df:19:11                 scsi-fcp:init
0xec00a0    N     20:00:d8:b1:90:41:1d:d1 (Cisco)
0xec00ef    N     50:06:01:61:08:60:7a:ab (Company B)     scsi-fcp:both
0xee0000    N     50:08:01:60:08:df:19:10                 scsi-fcp
0xee0020    N     20:13:00:11:0d:9d:07:00                 scsi-fcp:target
0xee0040    N     10:00:00:90:fa:d1:ef:12 (Company C)     scsi-fcp:init
0xee0060    N     20:00:00:6b:f1:70:08:ed (Cisco)         scsi-fcp:init fc-gs
0xef0020    N     50:08:01:60:01:59:49:32                 scsi-fcp
0xef0040    N     20:11:00:11:0d:96:e7:00                 scsi-fcp:init

Total number of entries = 13

VSAN 2:
--------------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0x5e0020    N     25:6b:28:6f:7f:21:03:f6 (Cisco)         npv
0x5e0040    N     25:6b:e0:0e:da:49:c2:2a (Cisco)         npv
0x5e0080    N     21:ed:00:2a:10:7a:89:1d (Cisco)         npv
0x840000    N     20:0f:2c:d0:2d:50:d3:48 (Cisco)         npv
0x840040    N     25:52:2c:d0:2d:50:d3:48 (Cisco)         npv

Total number of entries = 5
```

This example shows how to display the name server database and the statistical information for the specified VSANs:

```
switch# show fcns database vsan 1

VSAN 1:
--------------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0xe90000    N     20:00:00:6b:f1:70:08:ec (Cisco)         scsi-fcp:init fc-gs
0xec0020    N     21:00:00:24:ff:7f:37:05 (Company A)     scsi-fcp:target
0xec0040    N     50:08:01:60:01:59:49:33                 scsi-fcp:init
0xec0060    N     20:12:00:11:0d:9d:06:00                 scsi-fcp:init
0xec0080    N     50:08:01:60:08:df:19:11                 scsi-fcp:init
0xec00a0    N     20:00:d8:b1:90:41:1d:d1 (Cisco)
0xec00ef    N     50:06:01:61:08:60:7a:ab (Company B)     scsi-fcp:both
0xee0000    N     50:08:01:60:08:df:19:10                 scsi-fcp
0xee0020    N     20:13:00:11:0d:9d:07:00                 scsi-fcp:target
0xee0040    N     10:00:00:90:fa:d1:ef:12 (Company C)     scsi-fcp:init
0xee0060    N     20:00:00:6b:f1:70:08:ed (Cisco)         scsi-fcp:init fc-gs
0xef0020    N     50:08:01:60:01:59:49:32                 scsi-fcp
0xef0040    N     20:11:00:11:0d:96:e7:00                 scsi-fcp:init

Total number of entries = 13
```

This example shows how to display the name server database for all VSANs:

```
switch# show fcns database detail

show fcns database detail
------------------------
VSAN:200 FCID:0xee0000
------------------------
```

```
port-wwn (vendor) :21:00:00:24:ff:17:08:2e (Qlogic)
node-wwn :20:00:00:24:ff:17:08:2e
class :3
node-ip-addr :0.0.0.0
ipa :ff ff ff ff ff ff ff ff
fc4-types:fc4_features :scsi-fcp:init
symbolic-port-name :
symbolic-node-name :QLE2742 FW:v8.05.44 DVR:v2.1.73.0
port-type :N
port-ip-addr :0.0.0.0
fabric-port-wwn :20:11:00:de:fb:53:a3:a0
hard-addr :0x000000
permanent-port-wwn (vendor) :21:00:00:24:ff:17:08:2e (Qlogic)
connected interface :fc1/17
switch name (IP address) :sw (192.168.1.1)
------------------------
VSAN:200 FCID:0xee0020
```

This example shows how to display the name server database statistics for all VSANs:

```
switch# show fcns statistics

show fcns statistics
Name server statistics for vsan 1
==================================
registration requests received = 0
deregistration requests received = 0
queries received = 0
queries sent = 0
reject responses sent = 0
RSCNs received = 0
RSCNs sent = 0

Name server statistics for vsan 200
==================================
registration requests received = 18
deregistration requests received = 0
queries received = 78
queries sent = 0
reject responses sent = 0
RSCNs received = 0
RSCNs sent = 8

Name server statistics for vsan 201
==================================
registration requests received = 0
deregistration requests received = 0
queries received = 0
queries sent = 0
reject responses sent = 0
RSCNs received = 0
RSCNs sent = 0

Name server statistics for vsan 202
==================================
registration requests received = 0
deregistration requests received = 0
queries received = 0
queries sent = 0
reject responses sent = 0
RSCNs received = 0
RSCNs sent = 0
```

FDMI

Cisco Nexus N9K-C93180YC-FX and N9K-C93360YC-FX2 switches support 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 switch software can extract the following management information about attached HBAs and host operating systems without installing proprietary host agents:

- Manufacturer, model, and serial number
- Node name and node symbolic name
- Hardware, driver, and firmware versions
- Host operating system (OS) name and version number

All FDMI entries are stored in persistent storage and are retrieved when the FDMI process is started.

Displaying FDMI

This examples shows how to display all HBA details for a specified VSAN:

```
switch# show fdmi database detail vsan 1
```

This example displays HBA list for all VSANs:

```
switch# sh fdmi database
Registered HBA List for VSAN 10
  10:00:00:90:fa:c7:e1:f6
Registered HBA List for VSAN 108
  20:04:00:11:0d:dd:00:00
  20:05:00:11:0d:dd:00:00
```

This example displays the HBA list for a specific VSAN:

```
switch# sh fdmi database vsan 10
Registered HBA List for VSAN 10
  10:00:00:90:fa:c7:e1:f6
```

This example displays all the details of the HBA list:

```
switch# sh fdmi database detail
Registered HBA List for VSAN 10
-------------------------------
HBA-ID: 10:00:00:90:fa:c7:e1:f6
-------------------------------
Node Name        :20:00:00:90:fa:c7:e1:f6
Manufacturer     :Emulex Corporation
Serial Num       :FC61659139
Model            :LPe32002-M2
Model Description:Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
Hardware Ver     :0000000c
Driver Ver       :11.4.33.1
ROM Ver          :11.4.204.25
Firmware Ver     :11.4.204.25
OS Name/Ver      :VMware ESXi 6.7.0 Releasebuild-8169922
CT Payload Len   :245760
  Port-id: 10:00:00:90:fa:c7:e1:f6
    Supported FC4 types:1 scsi-fcp fc-gs
    Supported Speed    :8G 16G 32G
```

```
    Current Speed      :16G
    Maximum Frame Size :2048
    OS Device Name     :vmhba8
    Host Name          :localhost
Registered HBA List for VSAN 108
--------------------------------
HBA-ID: 20:04:00:11:0d:dd:00:00
--------------------------------
Node Name        :20:04:00:11:0d:23:b4:00
Manufacturer     :QLogic Corporation
Serial Num       :RFD1743U70327
Model            :QLE2742
Model Description:Cisco QLE2742 Dual Port 32Gb FC to PCIe Gen3 x8 Adapter
Hardware Ver     :BK3210407-43  B
Driver Ver       :8.07.00.34.Trunk-SCST.18-k
ROM Ver          :3.60
Firmware Ver     :8.08.204 (785ad0
  Port-id: 20:04:00:11:0d:dd:00:00
    Supported FC4 types:scsi-fcp 40 fc-av
    Supported Speed    :8G 16G 32G
    Current Speed      :32G
    Maximum Frame Size :2112
    OS Device Name     :qla2xxx:host7
    Host Name          :VirtuaLUN
--------------------------------
HBA-ID: 20:05:00:11:0d:dd:00:00
--------------------------------
Node Name        :20:05:00:11:0d:23:b5:00
Manufacturer     :QLogic Corporation
Serial Num       :RFD1743U70327
Model            :QLE2742
Model Description:Cisco QLE2742 Dual Port 32Gb FC to PCIe Gen3 x8 Adapter
Hardware Ver     :BK3210407-43  B
Driver Ver       :8.07.00.34.Trunk-SCST.18-k
ROM Ver          :3.60
Firmware Ver     :8.08.204 (785ad0
  Port-id: 20:05:00:11:0d:dd:00:00
    Supported FC4 types:scsi-fcp 40 fc-av
    Supported Speed    :8G 16G 32G
    Current Speed      :32G
    Maximum Frame Size :2112
    OS Device Name     :qla2xxx:host8
    Host Name          :VirtuaLUN
```

This example displays all the details of the HBA list for a specific VSAN:

```
switch# sh fdmi database detail vsan 10
Registered HBA List for VSAN 10
--------------------------------
HBA-ID: 10:00:00:90:fa:c7:e1:f6
--------------------------------
Node Name        :20:00:00:90:fa:c7:e1:f6
Manufacturer     :Emulex Corporation
Serial Num       :FC61659139
Model            :LPe32002-M2
Model Description:Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
Hardware Ver     :0000000c
Driver Ver       :11.4.33.1
ROM Ver          :11.4.204.25
Firmware Ver     :11.4.204.25
OS Name/Ver      :VMware ESXi 6.7.0 Releasebuild-8169922
CT Payload Len   :245760
  Port-id: 10:00:00:90:fa:c7:e1:f6
    Supported FC4 types:1 scsi-fcp fc-gs
    Supported Speed    :8G 16G 32G
```

```
Current Speed      :16G
Maximum Frame Size :2048
OS Device Name     :vmhba8
Host Name          :localhost
```

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 a State Change Registration (SCR) request). These notifications provide a timely indication of one or more of the following events:

- Disks joining or leaving the fabric
- A name server registration change
- A new zone enforcement
- IP address change
- Any other similar event that affects the operation of the host

A switch RSCN (SW-RSCN) is sent to registered hosts and to all reachable switches in the fabric.

Note The switch sends an RSCN to notify registered nodes that a change has occurred. It is up to the nodes to query the name server again to obtain the new information. The details of the changed information are not delivered by the switch in the RSCN sent to the nodes.

Displaying RSCN Information

The following example shows how to display registered device information:

```
switch# show rscn scr-table vsan 1

show rscn scr-table vsan 1
SCR table for VSAN: 1
---------------------------------------------
FC-ID REGISTERED FOR
---------------------------------------------
0xee0000 fabric and nport detected rscns
0xee0020 fabric and nport detected rscns
0xee00ef fabric and nport detected rscns

Total number of entries = 3
```

Note The SCR table is not configurable. It is populated when hosts send SCR frames with RSCN information. If hosts do not receive RSCN information, then the **show rscn scr-table** command will not return entries.

Multi-pid Option

If the RSCN multi-pid option is enabled, RSCNs generated to the registered Nx ports might 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, 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, one of the following actions 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 support multi-pid RSCN payloads. If so, disable the RSCN multi-pid option.

Note For the PORT_OFFLINE events, irrespective of whether the multi-pid option is enabled or disabled, multiple RSCNs are generated (depending on the number of ports) and sent immediately.

For the PORT_ONLINE events,

- if the multi-pid option is enabled, a single RSCN is generated irrespective of the number of ports and sent immediately. This RSCN includes multiple pages containing information about all the ports coming UP.
- if the multi-pid option is disabled, multiple RSCNs are generated (depending on the number of ports) and sent immediately.

Configuring the multi-pid Option

You can configure the **multi-pid** option.

SUMMARY STEPS

1. **configure terminal**
2. **rscn multi-pid vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**rscn multi-pid vsan** *vsan-id* **Example:** `switch(config)# rscn multi-pid vsan 405`	Sends RSCNs in a multi-pid format for the specified VSAN.

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 SAN switches.

You can suppress the transmission of these SW-RSCNs over an ISL.

SUMMARY STEPS

1. **configure terminal**
2. **rscn suppress domain-swrscn vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**rscn suppress domain-swrscn vsan** *vsan-id* **Example:** `switch(config)# rscn suppress domain-swrscn vsan 250`	Suppresses transmission of domain format SW-RSCNs for the specified VSAN.

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.

This example shows how 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 entering the **show rscn statistics** command:

```
switch# show rscn statistics vsan 1
```

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. When a timeout occurs, all the events are dequeued and coalesced RSCNs are sent to registered users. The default timer values minimize the number of coalesced RSCNs that are 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.

Note CFS is enabled by default. All devices in the fabric must have CFS enabled, or they do not receive distributions. If CFS is disabled for an application, that application does not distribute any configuration, and it does not accept a distribution from other devices in the fabric. Use **cfs distribute** command to enable the CFS

Note Before performing a downgrade, make sure that you revert the RSCN timer value in your network to the default value. Failure to do so will disable the links across your VSANs and other devices.

You can configure the RSCN timer.

SUMMARY STEPS

1. **configure terminal**
2. **rscn distribute**
3. **rscn event-tov** *timeout* **vsan** *vsan-id*
4. **no rscn event-tov** *timeout* **vsan** *vsan-id*
5. **rscn commit vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**rscn distribute** **Example:** `switch(config)# rscn distribute`	Enables RSCN timer configuration distribution.
Step 3	**rscn event-tov** *timeout* **vsan** *vsan-id* **Example:** `switch(config)# rscn event-tov 1000 vsan 501`	Sets the event time-out value in milliseconds for the specified VSAN. The range is 0 to 2000 milliseconds. Setting a zero (0) value disables the timer.
Step 4	**no rscn event-tov** *timeout* **vsan** *vsan-id* **Example:** `switch(config)# no rscn event-tov 1100 vsan 245`	Reverts to the default value (2000 milliseconds for Fibre Channel VSANs).
Step 5	**rscn commit vsan** *vsan-id* **Example:** `switch(config)# rscn commit vsan 25`	Commits the RSCN timer configuration to be distributed to the switches in the specified VSAN.

Verifying the RSCN Timer Configuration

You verify the RSCN timer configuration using the **show rscn event-tov vsan** command. This example shows how to clear the RSCN statistics for VSAN 10:

```
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. 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, which also reduces the number of SW-RSCNs.

RSCN supports two modes, distributed and nondistributed. In distributed mode, RSCN uses Cisco Fabric Services (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.

Caution 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.

Enabling RSCN Timer Configuration Distribution

You can enable RSCN timer configuration distribution.

SUMMARY STEPS

1. **configure terminal**
2. **rscn distribute**
3. **no rscn distribute**

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**rscn distribute** **Example:**	Enables RSCN timer distribution.

	Command or Action	Purpose
	`switch(config)# rscn distribute`	
Step 3	**no rscn distribute** **Example:** `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 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.

You can commit RSCN timer configuration changes.

SUMMARY STEPS

1. **configure terminal**
2. **rscn commit vsan** *timeout*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**rscn commit vsan** *timeout* **Example:** `switch(config)# rscn commit vsan 500`	Commits the RSCN timer changes.

Discarding the RSCN Timer Configuration Changes

If you discard (abort) the changes made to the pending database, the configuration database remains unaffected and the lock is released.

You can discard RSCN timer configuration changes.

SUMMARY STEPS

1. **configure terminal**

2. **rscn abort vsan** *timeout*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**rscn abort vsan** *timeout* **Example:** `switch(config)# rscn abort vsan 800`	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.

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 RSCN session, use the **clear rscn session vsan** command in EXEC mode. This example shows how to clear the RSCN session for VSAN 10:

```
switch# clear rscn session vsan 10
```

Displaying RSCN Configuration Distribution Information

This example shows how to display the registration status for RSCN configuration distribution:

```
switch# show cfs application name rscn
 Enabled        : Yes
 Timeout        : 5s
 Merge Capable  : Yes
 Scope          : Logical
```

Note A merge failure results when the RSCN timer values are different on the merging fabrics.

This example shows how 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 vsan 1
rscn event-tov 2000 ms vsan 1
```

This example shows how to display the difference between pending and active configurations:

```
switch# show rscn pending-diff vsan 10
- rscn event-tov 2000
+ rscn event-tov 1001
```

Default Settings for RSCN

The following table lists the default settings for RSCN.

Table 16: Default RSCN Settings

Parameters	Default
RSCN timer value	2000 milliseconds for Fibre Channel VSANs
RSCN timer configuration distribution	Disabled

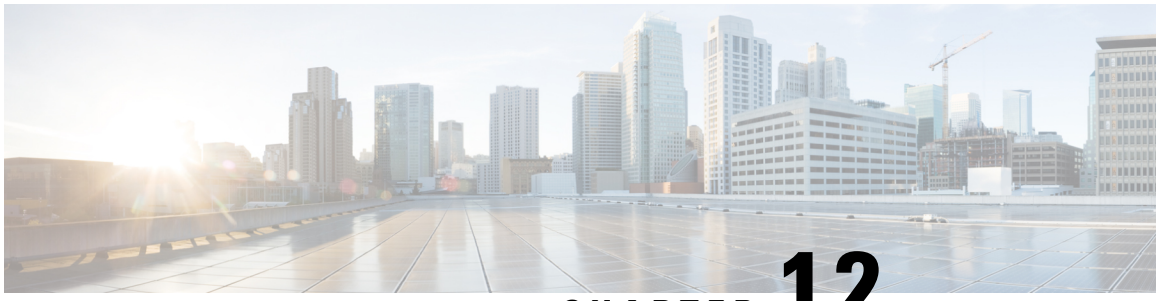

CHAPTER 12

Distributing Device Alias Services

This chapter describes how to distribute device alias services.

This chapter contains the following sections:

- Distributing Device Alias Services, on page 143

Distributing Device Alias Services

Cisco SAN switches support Distributed Device Alias Services (device aliases) on a fabric-wide basis.

Information About Device Aliases

Cisco SAN switches support Distributed Device Alias Services (device aliases) on a fabric-wide basis.

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

Device Alias Features

Device aliases have the following features:

- The device alias information is independent of the VSAN configuration.
- The device alias configuration and distribution is independent of the zone server and the zone server database.
- 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.
- Basic and enhanced modes.
- Device aliases used to configure zones are displayed automatically with their respective pWWNs in the **show** command output.

Related Topics

Device Alias Modes, on page 145

Device Alias Requirements

Device aliases have the following requirements:

- You can only assign device aliases to pWWNs.
- There must be a one-to-one relationship between the pWWN and the device alias that maps to it.
- A device alias name is restricted to 64 alphanumeric characters and may include one or more of the following characters:
  - a to z and A to Z
  - Device alias names must begin with an alphabetic character (a to z or A to Z).
  - 1 to 9
  - - (hyphen) and _ (underscore)
  - $ (dollar sign) and ^ (up caret)

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.

Device alias database changes are validated with the applications. If any of the applications cannot accept the device alias database changes, then those changes are rejected; this applies to device alias database changes resulting from either a commit or merge operation.

Creating Device Aliases

You can create a device alias in the pending database.

SUMMARY STEPS

1. **configure terminal**
2. **device-alias database**
3. **device-alias name** *device-name* **pwwn** *pwwn-id*
4. **no device-alias name** *device-name*
5. **device-alias rename** *old-device-name new-device-name*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**device-alias database**
	Example:
	`switch(config)# device-alias database`
	`switch(config-device-alias-db)#`
		Enters the pending database configuration submode.

Step 3	**device-alias name** *device-name* **pwwn** *pwwn-id* **Example:** `switch(config-device-alias-db)# device-alias name mydevice pwwn 21:01:00:e0:8b:2e:80:93`	Specifies a device name 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	**no device-alias name** *device-name* **Example:** `switch(config-device-alias-db)# no device-alias name mydevice`	Removes the device name for the device that is identified by its pWWN.
Step 5	**device-alias rename** *old-device-name new-device-name* **Example:** `switch(config-device-alias-db)# device-alias rename mydevice mynewdevice`	Renames an existing device alias with a new name.

EXAMPLES

This example shows how to display the device alias configuration.

```
switch# show device-alias name x

device-alias name x pwwn 21:01:00:e0:8b:2e:80:93
```

Device Alias Modes

When operating in basic mode, which is the default mode, the device alias is immediately expanded to a pWWN. In basic mode, when device aliases are changed to point to a new HBA, for example, that change is not reflected in the zone server. Users must remove the previous HBA's pWWN, add the new HBA's pWWN and then reactivate the zoneset.

Note Both basic and enhanced device alias modes are supported in Cisco NX-OS Release 10.1(1).

Note Cisco NX-OS Release 9.3(3) supports only the basic mode.

Beginning Cisco NX-OS Release 9.3(5), both basic and enhanced device alias modes are supported.

When operating in enhanced mode, applications accept a device alias name in its native format. Instead of expanding the device alias to a pWWN, the device alias name is stored in the configuration and distributed in its native device alias format. So application such as zone server can automatically keep track of the device alias membership changes and enforce them accordingly. The primary benefit of operating in enhanced mode is that you have a single point of change.

Whenever you change device alias modes, the change is distributed to other switches in the network only if device alias distribution is enabled or on. Otherwise, the mode change only takes place on the local switch.

Note Enhanced mode, or native device alias-based configurations, are not accepted in interop mode VSANs. IVR zoneset activation fails in interop mode VSANs if the corresponding zones have native device alias-based members.

Device Alias Mode Guidelines and Limitations for Device Alias Services

Device Alias services have these configuration guidelines and limitations:

- If two fabrics running in different device alias modes are joined together, the device alias merge fails. There is no automatic conversion to one mode or the other during the merge process. In this situation, you must select one mode over the other.
- Before changing from enhanced to basic mode, you must first explicitly remove all native device alias-based configurations from both local and remote switches, or replace all device alias-based configuration members with the corresponding pWWN.
- If you remove a device alias from the device alias database, all applications automatically stop enforcing the corresponding device alias. If that corresponding device alias is part of an active zone set, all the traffic to and from that pWWN is disrupted.
- Renaming the device alias not only changes the device alias name in the device alias database, but also replaces the corresponding device alias configuration in all of the applications.
- When a new device alias is added to the device alias database, and the application configuration is present on that device alias, it automatically takes effect. For example, if the corresponding device alias is part of the active zoneset and the device is online, then zoning is enforced automatically. You do not have to reactivate the zone set.
- If a device alias name is mapped to a new HBA's pWWN, the application's enforcement changes accordingly. In this case, the zone server automatically enforces zoning based on the new HBA's pWWN.

Configuring Device Alias Modes

You can configure device aliases to operate in enhanced mode.

SUMMARY STEPS

1. **configure terminal**
2. **device-alias mode enhanced**
3. **no device-alias mode enhance**

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**device-alias mode enhanced** **Example:** `switch(config)# device-alias mode enhanced`	Assigns the device alias to operate in enhanced mode.
Step 3	**no device-alias mode enhance** **Example:** `switch(config)# no device-alias mode enhance`	Assigns the device alias to operate in basic mode.

EXAMPLES

This example shows how to display the current device alias mode setting.

```
switch# show device-alias status
Fabric Distribution: Enabled
Database:- Device Aliases 0 Mode: Basic
Locked By:- User "admin" SWWN 20:00:00:0d:ec:30:90:40
Pending Database:- Device Aliases 0 Mode: Basic
```

Device Alias Distribution

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

If device alias distribution is disabled, database changes are not distributed to the switches in the fabric. The same changes would have to be performed manually on all switches in the fabric to keep the device alias database up-to-date. Database changes immediately take effect, so there would also not be any pending database and commit or abort operations. If you have not committed the changes and you disable distribution, a commit task fails.

Note CFS is enabled by default. All devices in the fabric must have CFS enabled, or they do not receive distributions.If CFS is disabled for an application, that application does not distribute any configuration, and it does not accept a distribution from other devices in the fabric. To enable CFS use the **cfs distribute** command.

This example shows how to display a failed device alias 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.)
```

Locking the Fabric

When you perform any device alias configuration 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. Subsequent modifications are made to the pending database. The pending database remains in use until you commit the modifications to the pending database or discard (**abort**) the changes to the pending database.

Committing Changes

You can commit changes.

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

- The pending database content overwrites the effective database content.
- The pending database is distributed to the switches in the fabric and the effective database on those switches is overwritten with the new changes.
- The pending database is emptied of its contents.
- The fabric lock is released for this feature.

SUMMARY STEPS

1. **configure terminal**
2. **device-alias commit**

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

	Command or Action	Purpose
Step 2	**device-alias commit** **Example:** `switch(config)# device-alias commit`	Commits the changes made to the currently active session.

Discarding Changes

You can discard the device alias session changes.

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

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

SUMMARY STEPS

1. **configure terminal**
2. **device-alias abort**

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**device-alias abort** **Example:** `switch(config)# device-alias abort`	Discards the currently active session.

EXAMPLES

This example shows how to display the status of the discard operation:

```
switch(config)# show device-alias status

Fabric Distribution: Enabled
Database:- Device Aliases 2 Mode: Basic
Checksum: 0x22a1d11a2762bdb3cae50f16a21a1e1
Locked By:- User "CLI/SNMPv3:admin" SWWN 20:00:00:de:fb:9d:0e:a0
Pending Database:- Device Aliases 3 Mode: Basic
```

This example shows how to display the status of the abort operation:

```
switch(config)# device-alias abort
switch(config)#
```

```
switch(config)#  show device-alias session status
Last Action Time Stamp : Mon Nov 4 09:10:11 2019
Last Action : Abort
Last Action Result : Success
Last Action Failure Reason : none
switch(config)#
```

Overriding the Fabric Lock

You can use locking operations (clear, commit, abort) only when device alias distribution is enabled. 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.

The changes are only available in the volatile directory and may be discarded if the switch is restarted.

To use administrative privileges and release a locked device alias session, use the **clear device-alias session** command in EXEC mode.

```
switch# clear device-alias session
```

This example shows how to display the status of the clear operation:

```
switch# show device-alias status

Fabric Distribution: Enabled

Database:- Device Aliases 24

Status of the last CFS operation issued from this switch:

==========================================================

Operation: Clear Session<--------------------Lock released by administrator

Status: Success<-----------------------------Successful status of the operation
```

Disabling and Enabling Device Alias Distribution

You can disable or enable the device alias distribution.

SUMMARY STEPS

1. **configure terminal**
2. **no device-alias distribute**
3. **device-alias distribute**

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**no device-alias distribute** **Example:** `switch(config)# no device-alias distribute`	Disables the distribution.

	Command or Action	Purpose
Step 3	**device-alias distribute** **Example:** `switch(config)# device-alias distribute`	Enables the distribution (default).

EXAMPLES

This example shows how to display the status of device alias distribution:

```
switch# show device-alias status

Fabric Distribution: Disabled
Database:- Device Aliases 3 Mode: Basic
Checksum: 0x284031ab5aade498a7e89cef1b04d7f
switch(config)#
```

This example shows the device alias display when distribution is disabled:

```
switch# show device-alias status

Fabric Distribution: Disabled
Database:- Device Aliases 3 Mode: Basic
Checksum: 0x284031ab5aade498a7e89cef1b04d7f
switch(config)#
```

Device Alias Database Merge Guidelines

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 identical pWWNs are not mapped to two different device aliases.

If the combined number of device entries in both databases exceeds the supported configuration limit, then the merge will fail.

Verifying the Device Alias Configuration

To display device alias information, perform one of the following tasks:

Command	Purpose
show zoneset [**active**]	Displays the device aliases in the zone set information.
show device-alias database [**pending** | **pending-diffs**]	Displays the device alias database.
show device-alias {**pwwn** *pwwn-id* | **name** *device-name* } [**pending**]	Displays the device alias information for the specified pwwn or alias.
show flogi database [**pending**]	Displays device alias information in the flogi database.
show fcns database [**pending**]	Displays device alias information in the fcns database.

Default Settings for Device Alias Services

The following table lists the default settings for device alias parameters.

Table 17: Default Device Alias Parameters

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

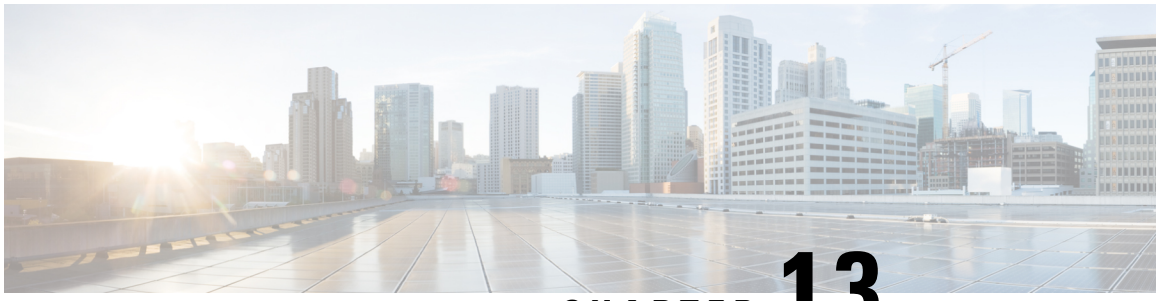

CHAPTER 13

Configuring and Managing Zones

This chapter describes how to configure and manage zones.

This chapter contains the following sections:

- Information About Zones, on page 153

Information About 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.

Note Cisco NX-OS Release 9.3(3) supports only basic zoning features.

Cisco NX-OS Release 10.1(1) supports basic, enhanced, and smart zoning. Advanced zoning capabilities specified in the FC-GS-4 and FC-SW-3 standards are supported. You can use either the existing basic zoning capabilities or the advanced, standards-compliant zoning capabilities.

Information About Zoning

Zoning Features

Zoning includes the following features:

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

- A physical fabric can have a maximum of 16,000 members. This includes all VSANs in the fabric.

- A zone set consists of one or more zones.
  - A zone set can be activated or deactivated as a single entity across all switches in the fabric.
  - Only one zone set can be activated at any time.
  - A zone can be a member of more than one zone set.
  - A zone switch can have a maximum of 1000 zone sets.
- Zoning can be administered from any switch in the fabric.
  - When you activate a zone (from any switch), all switches in the fabric receive the active zone set. Additionally, full zone sets are distributed to all switches in the fabric, if this feature is enabled in the source switch in basic zoning mode and default in enhanced zoning mode.
  - 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 can be specified using the following device alias members:
  - 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.
  - Device Alias—Specifies a device alias name.
  - FC Alias—Specifies a FC alias name.

Note For N ports attached to the switch over a virtual Fibre Channel interface, you can specify zone membership using the device alias of the logged-in device, pWWN of the N port, the FC ID of the N port, or the fabric pWWN of the virtual Fibre Channel interface.

- Default zone membership includes all ports or WWNs that do not have a specific membership association. Access between default zone members is controlled by the default zone policy.
- You can configure up to 8000 zones per VSAN and a maximum of 8000 zones for all VSANs on the switch.
- The maximum of 4000 zone ACL entries are supported.

- If the number of zone ACL entries exceed 4000, the zone may transition to soft zoning mode.

Zoning Example

The following figure shows a zone set with two zones, zone 1 and zone 2, in a fabric. Zone 1 provides access from all three hosts (H1, H2, H3) to the data residing on storage systems S1 and S2. Zone 2 restricts the data on S3 to access only by H3. H3 resides in both zones.

Figure 19: Fabric with Two Zones

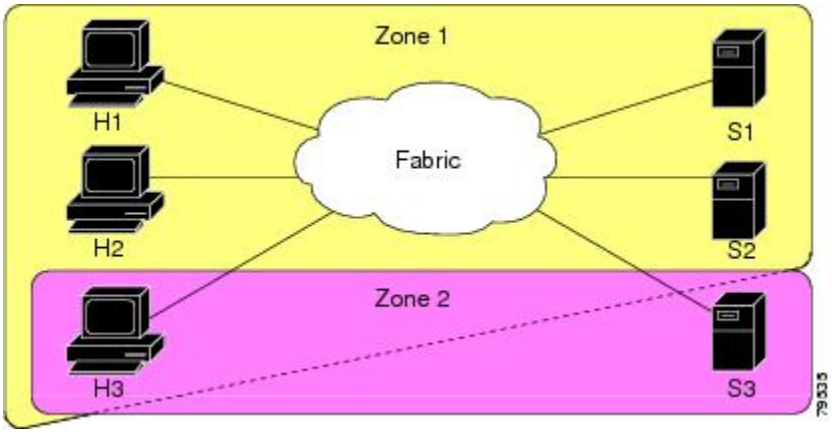

You can use other ways to partition this fabric into zones. The following figure shows another possibility. Assume that there is a need to isolate storage system S2 for the purpose of testing new software. To achieve this, zone 3 is configured, which contains only host H2 and storage S2. You can restrict access to only H2 and S2 in zone 3, and to H1 and S1 in zone 1.

Figure 20: Fabric with Three Zones

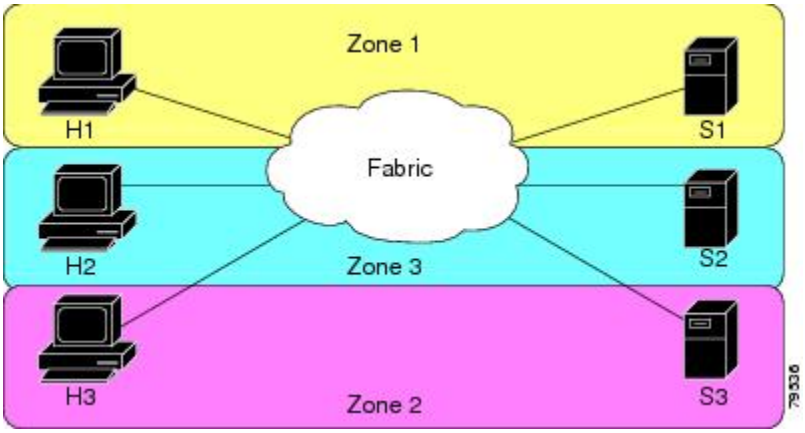

Zone Implementation

Cisco SAN switches automatically support the following basic zone features (no additional configuration is required):

- Zones are contained in a VSAN.
- Hard zoning cannot be manually disabled.

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

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

- Propagate full zone sets to all switches per VSAN.
- Change the default policy for unzoned members.
- Bring E ports out of isolation.

Active and Full Zone Sets

Before configuring a zone set, consider the following guidelines:

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

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

The following figure shows a zone being added to an activated zone set.

Figure 21: Active and Full Zone Sets

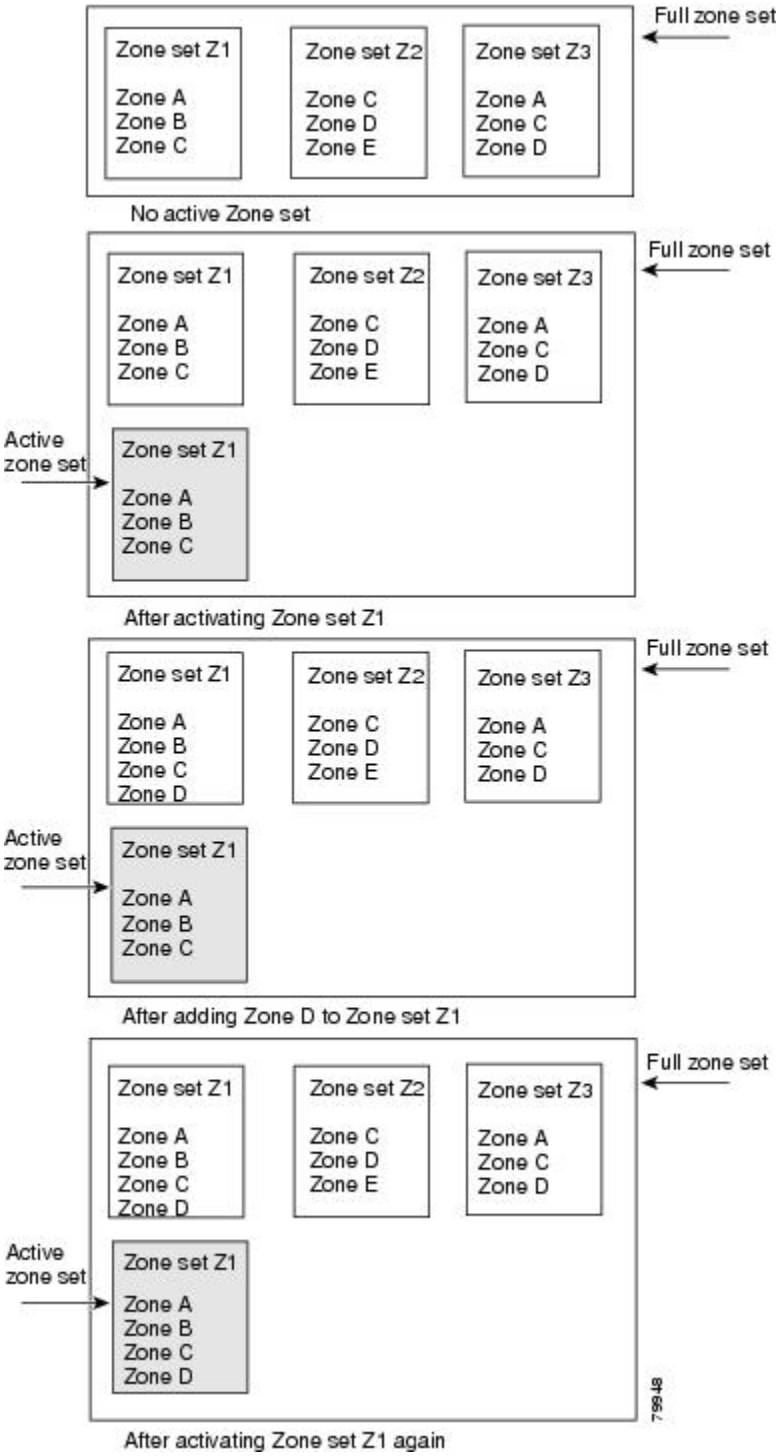

Configuring a Zone

You can configure a zone and assign a zone name.

SUMMARY STEPS

1. **configure terminal**
2. **zone name** *zone-name* **vsan** *vsan-id*
3. **member** *type value*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**zone name** *zone-name* **vsan** *vsan-id* **Example:** `switch(config)# zone name test vsan 5`	Configures a zone in the specified VSAN. **Note** All alphanumeric characters or one of the following symbols ($, -, ^, _) are supported.
Step 3	**member** *type value* **Example:** `switch(config-zone)# member interface 4`	Configures a member for the specified zone based on the type (pWWN, fabric pWWN, FC ID, fcalias, device alias, domain ID, or interface) and value specified.

Configuration Examples

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.

The following examples show how to configure zone members:

```
switch(config)# zone name MyZone vsan 2
```

pWWN example:

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

Fabric pWWN example:

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

FC ID example:

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

FC alias example:

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

Device alias example:

```
switch(config-zone)# member device-alias finance
```

Domain ID example:

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

Show WWN example:

```
switch# show wwn switch
switch(config-zone)# member interface fc 2/1
switch(config-zone)# member interface fc 2/1 swwn 20:00:00:05:30:00:4a:de
switch(config-zone)# member interface fc 2/1 domain-id 25
```

Note The zone's default system settings such as **system default zone default-zone permit** and **system default zone distribute full** takes effect only on the newly created VSANs after you manually apply the settings. These settings may not be applied on VSAN 1 even though when they are set as part of FC setup script.

The following example shows how to configure different types of member alias:

```
switch(config)# fcalias name AliasSample vsan 3
```

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
```

Device alias example:

```
switch(config-fcalias)# member device-alias devName
```

Zone Sets

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

Figure 22: Hierarchy of Zone Sets, Zones, and Zone Members

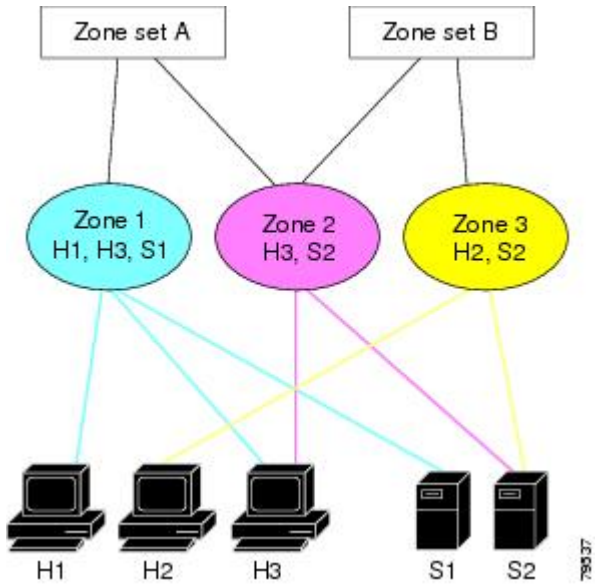

Zones provide a method for specifying access control, while zone sets are a grouping of zones to enforce access control in the fabric. Either zone set A or zone set B can be activated (but not together).

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

Activating a Zone Set

You can activate or deactivate an existing zone set.

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

SUMMARY STEPS

1. **configure terminal**
2. **zoneset activate name** *zoneset-name* **vsan** *vsan-id*
3. **no zoneset activate name** *zoneset-name* **vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**zoneset activate name** *zoneset-name* **vsan** *vsan-id* **Example:**	Activates the specified zone set.

	Command or Action	Purpose
	`switch(config)# zoneset activate name test vsan 34`	
Step 3	**no zoneset activate name** *zoneset-name* **vsan** *vsan-id* **Example:** `switch(config)# no zoneset activate name test vsan 30`	Deactivates the specified zone set.

Default Zone

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

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

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

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

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

Note The default settings for default zone configurations can be changed.

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

Configuring the Default Zone Access Permission

To permit or deny traffic to members in the default zone, perform this task:

SUMMARY STEPS

1. **configure terminal**
2. **zone default-zone permit vsan** *vsan-id*
3. **no zone default-zone permit vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**zone default-zone permit vsan** *vsan-id* **Example:** `switch(config)# zone default-zone permit vsan 13`	Permits traffic flow to default zone members.
Step 3	**no zone default-zone permit vsan** *vsan-id* **Example:** `switch(config)# no zone default-zone permit vsan 40`	Denies (default) traffic flow to default zone members.

FC Alias Creation

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

- pWWN—The WWN of the N 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.
- 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—Specifies a device alias name.

Tip The switch supports a maximum of 2048 aliases per VSAN.

Creating FC Aliases

You create an alias.

SUMMARY STEPS

1. **configure terminal**
2. **fcalias name** *alias-name* **vsan** *vsan-id*
3. **member** *type value*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**fcalias name** *alias-name* **vsan** *vsan-id* **Example:** `switch(config)# fcalias name testname vsan 50`	Configures an alias name. The alias name can be any case-sensitive, alphanumeric string up to 64 characters.
Step 3	**member** *type value* **Example:** `switch(config-fcalias)# member pwwn 4`	Configures a member for the specified fcalias based on the type (pWWN, fabric pWWN, FC ID, domain ID, or interface) and value specified. **Note** Multiple members can be specified on multiple lines.

Creating FC Aliases Example

Table 18: Type and Value Syntax for the member Command

Device alias	**member device-alias** *device-alias*
Domain ID	**member domain-id** *domain-id* **portnumber** *number*
FC ID	**member fcid** *fcid*
Fabric pWWN	**member fwwn** *fwwn-id*
Local sWWN interface	**member interface** *type slot*/*port*
Domain ID interface	**member interface** *type slot*/*port* **domain-id** domain-id
Remote sWWN interface	**member interface** *type slot*/*port* **swwn** *swwn-id*
pWWN	**member pwwn** *pwwn-id*

The following example shows how to configure different types of member alias:

```
switch(config)# fcalias name AliasSample vsan 3
```

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
switch(config-fcalias)# member interface fc 2/1
switch(config-fcalias)# member interface fc2/1 domain-id 25
```

Device alias example:

```
switch(config-fcalias)# member device-alias devName
```

Creating Zone Sets and Adding Member Zones

You can create a zone set to include several zones.

SUMMARY STEPS

1. **configure terminal**
2. **zone set name** *zoneset-name* **vsan** *vsan-id*
3. **member** *name*
4. **zone name** *zone-name*
5. **member fcid** *fcid*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**zone set name** *zoneset-name* **vsan** *vsan-id* **Example:** `switch(config)# zone set name new vsan 23`	Configures a zone set with the configured zoneset-name. **Tip** To activate a zone set, you must first create the zone and a zone set.
Step 3	**member** *name* **Example:** `switch(config-zoneset)# member new`	Adds a zone as a member of the previously specified zone set. **Tip** If the specified zone name was not previously configured, this command will return a "zone not present" error message:
Step 4	**zone name** *zone-name* **Example:** `switch(config-zoneset)# zone name trial`	Adds a zone to the specified zone set. **Tip** Execute this step only if you need to create a zone from a zone set prompt.
Step 5	**member fcid** *fcid* **Example:** `switch(config-zoneset-zone)# member fcid 0x222222`	Adds a new member to the new zone. **Tip** Execute this step only if you need to add a member to a zone from a zone set prompt.

Tip You do not have to copy the running configuration to the startup configuration to store the active zone set. However, you need to copy the running configuration to the startup configuration to explicitly store full zone sets.

Zone Enforcement

Zoning can be enforced in two ways: soft and hard. Each end device (N 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 N port does not know about the FC IDs of other devices outside its zone, it cannot access those devices.

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

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

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

Cisco SAN switches support both hard and soft zoning.

Zone Set Distribution

You can distribute full zone sets using one of two methods: one-time distribution using the **zoneset distribute vsan** command at the EXEC mode level or full zone set distribution using the **zoneset distribute full vsan** command at the configuration mode level. The following table lists the differences between the methods.

Table 19: Zone Set Distribution Differences

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

Enabling Full Zone Set Distribution

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

You can enable full zone set and active zone set distribution to all switches on a per VSAN basis.

SUMMARY STEPS

1. **configure terminal**
2. **zoneset distribute full vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**zoneset distribute full vsan** *vsan-id* **Example:** `switch(config)# zoneset distribute full vsan 12`	Enables sending a full zone set along with an active zone set.

Enabling a One-Time Distribution

You can perform a one-time distribution of inactive, unmodified zone sets throughout the fabric.

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 command only distributes the full zone set information, as it does not save the information to the startup configuration. You must explicitly enter the **copy running-config startup-config** command to save the full zone set information to the startup configuration.

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

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

```
switch# show zone status vsan 3
VSAN: 3 default-zone: permit distribute: active only Interop: 100
    mode:basic merge-control:allow
    session:none
    hard-zoning:enabled
Default zone:
    qos:none broadcast:disabled ronly:disabled
Full Zoning Database :
    Zonesets:0  Zones:0 Aliases: 0
Active Zoning Database :
    Name: nozoneset  Zonesets:1  Zones:2
```

```
Status: Zoneset distribution completed at 04:01:06 Aug 28 2010
```

Recovering from Link Isolation

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

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

Figure 23: Importing and Exporting the Database

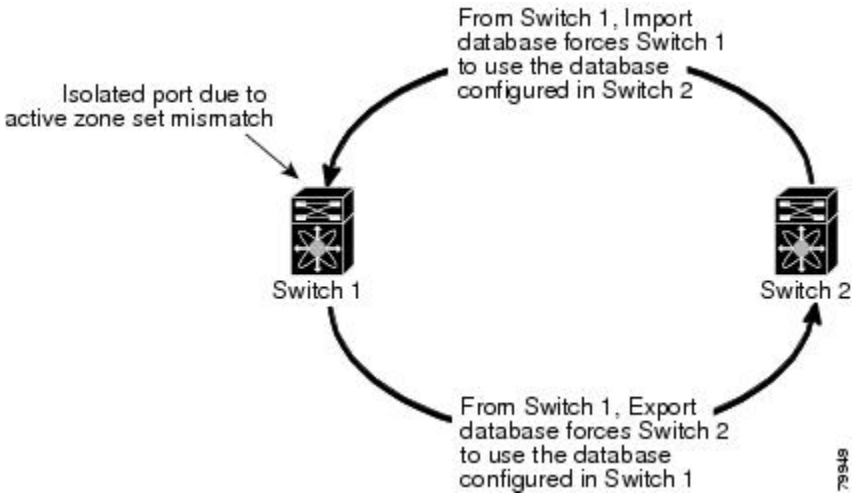

Importing and Exporting Zone Sets

You can import or export the zone set information from or to an adjacent switch.

SUMMARY STEPS

1. **zoneset export vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	**zoneset export vsan** *vsan-id* **Example:** `switch# zoneset export vsan 5`	Exports the zone set to the adjacent switch connected through the specified VSAN or range of VSANs.

Zone Set Duplication

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

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

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

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

Copying Zone Sets

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

SUMMARY STEPS

1. **zone copy active-zoneset full-zoneset vsan** *vsan-id*
2. **zone copy vsan** *vsan-id* **active-zoneset scp://guest@myserver/tmp/active_zoneset.txt**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**zone copy active-zoneset full-zoneset vsan** *vsan-id* **Example:** `switch# zone copy active-zoneset full-zoneset vsan 301`	Makes a copy of the active zone set in the specified VSAN to the full zone set.
Step 2	**zone copy vsan** *vsan-id* **active-zoneset scp://guest@myserver/tmp/active_zoneset.txt** **Example:** `switch# zone copy vsan 55 active-zoneset scp://guest@myserver/tmp/active_zoneset.txt`	Copies the active zone in the specified VSAN to a remote location using SCP.

Renaming Zones, Zone Sets, and Aliases

You can rename a zone, zone set, fcalias, or zone-attribute-group.

SUMMARY STEPS

1. **configure terminal**
2. **zoneset rename** *oldname newname* **vsan** *vsan-id*
3. **zone rename** *oldname newname* **vsan** *vsan-id*

4. **fcalias rename** *oldname newname* **vsan** *vsan-id*
5. **zoneset activate name** *newname* **vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**zoneset rename** *oldname newname* **vsan** *vsan-id* **Example:** `switch(config)# zoneset rename test myzoneset vsan 60`	Renames a zone set in the specified VSAN.
Step 3	**zone rename** *oldname newname* **vsan** *vsan-id* **Example:** `switch(config)# zone rename test myzone vsan 50`	Renames a zone in the specified VSAN.
Step 4	**fcalias rename** *oldname newname* **vsan** *vsan-id* **Example:** `switch(config)# fcalias rename test myfc vsan 200`	Renames a fcalias in the specified VSAN.
Step 5	**zoneset activate name** *newname* **vsan** *vsan-id* **Example:** `switch(config)# zoneset activate name myzone vsan 50`	Activates the zone set and updates the new zone name in the active zone set.

Cloning Zones, Zone Sets and FC Aliases

You can clone a zone, zone set and fcalias.

SUMMARY STEPS

1. **configure terminal**
2. **zoneset clone** *oldname newname* **vsan** *vsan-id*
3. **zone clone** *oldname newname* **vsan** *number*
4. **fcalias clone** *oldname newname* **vsan** *vsan-id*
5. **zoneset activate name** *newname* **vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	**configure terminal** **Example:**	Enters global configuration mode.

	Command or Action	Purpose
	`switch# configure terminal` `switch(config)#`	
Step 2	**zoneset clone** *oldname newname* **vsan** *vsan-id* **Example:** `switch(config)# zoneset clone test myzoneset2 vsan 2`	Clones a zone set in the specified VSAN.
Step 3	**zone clone** *oldname newname* **vsan** *number* **Example:** `switch(config)# zone clone test myzone3 vsan 3`	Clones a zone in the specified VSAN.
Step 4	**fcalias clone** *oldname newname* **vsan** *vsan-id* **Example:** `switch(config)# fcalias clone test myfcalias vsan 30`	Clones a fcalias in the specified VSAN.
Step 5	**zoneset activate name** *newname* **vsan** *vsan-id* **Example:** `switch(config)# zoneset activate name myzonetest1 vsan 3`	Activates the zone set and updates the new zone name in the active zone set.

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

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

Verifying the Zone Configuration

You can view any zone information by using the **show** command. If you request information for a specific object (for example, a specific zone, zone set, VSAN, or alias, or keywords such as brief or active), only information for the specified object is displayed.

Command	Purpose
show zone	Displays zone information for all VSANs.

Command	Purpose
show zone vsan *vsan-id*	Displays zone information for a specific VSAN.
show zoneset vsan *vsan-id*	Displays the configured zone sets for a range of VSANs.
show zone name *zone-name*	Displays the members of a specific zone.
show fcalias vsan *vsan-id*	Displays the fcalias configuration.
show zone member pwwn *pwwn-id*	Displays all zones to which a member belongs.
show zone statistics	Displays the number of control frames exchanged with other switches.
show zoneset active	Displays the active zone set.
show zone active	Displays the active zones.
show zone status	Displays the zone status.

Enhanced Zoning

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.

Note In case the scale zone configuration replays with enhanced zone mode, you must manually clear the local zone database before applying the saved scale zone configuration to the running configuration.

Note Broadcast zoning is not supported on the Cisco Nexus 9000 Series switches.

The following table compares the differences between basic and enhanced zoning:

Table 20: Advantages of Enhanced Zoning

Basic Zoning	Enhanced Zoning	Enhanced Zoning Advantages
Administrators can make simultaneous configuration changes. Upon activation, one administrator can overwrite another administrator's changes.	Performs all configurations within a single configuration session. When you begin a session, the switch locks the entire fabric to implement the change.	One configuration session for the entire fabric to ensure consistency within the fabric.
If a zone is part of multiple zone sets, you create an instance of this zone in each zone set.	References to the zone are used by the zone sets as required once you define the zone.	Reduced payload size as the zone is referenced. The size is more significant 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 per switch basis, the managing switch provides a combined status about the activation. It does not identify the failure switch.	Retrieves the activation results and the nature of the problem from each remote switch.	Enhanced error reporting eases the troubleshooting process.
To distribute the zoning database, you must reactivate the same zone set. The reactivation may affect hardware changes for hard zoning on the local switch and on remote switches.	Implements changes to the zoning database and distributes it without reactivation.	Distribution of zone sets without activation avoids hardware changes for hard zoning in the switches.

Changing from Basic Zoning to Enhanced Zoning

You can change to the enhanced zoning mode from the basic mode.

Step 1 Verify that all switches in the fabric can operate in the enhanced mode.

Step 2 If one or more switches cannot operate in the enhanced mode, then your request to move to enhanced mode is rejected.

Step 3 Set the operation mode to enhanced zoning mode.

Changing from Enhanced Zoning to Basic Zoning

Cisco Nexus 9000 switches allow you to change from enhanced zoning to basic zoning to enable you to downgrade and upgrade to other Cisco NX-OS releases.

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

Step 2 If such configurations exist, delete them before proceeding with this procedure. If you do not delete the existing configuration, the switch software automatically removes them.

Step 3 Set the operation mode to basic zoning mode.

Enabling Enhanced Zoning

You can enable enhanced zoning in a VSAN.

By default, the enhanced zoning feature is disabled in Cisco Nexus 9000 switches.

SUMMARY STEPS

1. **configure terminal**
2. **zone mode enhanced vsan** *vsan-id*
3. **no zone mode enhanced vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**zone mode enhanced vsan** *vsan-id* **Example:** `switch(config)# zone mode enhanced vsan 22`	Enables enhanced zoning in the specified VSAN.
Step 3	**no zone mode enhanced vsan** *vsan-id* **Example:** `switch(config)# no zone mode enhanced vsan 30`	Disables enhanced zoning in the specified VSAN.

Modifying the Zone Database

You can commit or discard changes to the zoning database in a VSAN.

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.

SUMMARY STEPS

1. **configure terminal**

2. **zone commit vsan** *vsan-id*
3. switch(config)# **zone commit vsan** *vsan-id* **force**
4. switch(config)# **no zone commit vsan** *vsan-id*
5. **no zone commit vsan** *vsan-id* **force**

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**zone commit vsan** *vsan-id* **Example:** `switch(config)# zone commit vsan 679`	Applies the changes to the enhanced zone database and closes the session.
Step 3	switch(config)# **zone commit vsan** *vsan-id* **force** **Example:** `switch(config)# zone commit vsan 34 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** *vsan-id* **Example:** `switch(config)# no zone commit vsan 22`	Discards the changes to the enhanced zone database and closes the session.
Step 5	**no zone commit vsan** *vsan-id* **force** **Example:** `switch(config)# no zone commit vsan 34 force`	Forcefully discards the changes to the enhanced zone database and closes the session created by another user.

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.

Verifying Enhanced Zone Information

This example shows how to display the zone status for a specified VSAN:

```
switch# show zone status vsan 2
```

Merging the Database

The merge method 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 following table.

Table 21: Database Zone Merge Status

Local Database	Adjacent Database	Merge Status	Results of the Merge
The databases contain zone sets with the same name. In the enhanced zoning mode, the active zone set does not have a name in interop mode 3. The zone set names are only present for full zone sets but are different zones, aliases, and attributes groups.		Successful.	ISLs are not isolated if the database mege is successful.
The databases contain a zone, FC alias, or zone attribute group object with same name1 but different members.		Failed.	The adjacent database information populates the local database. ISLs are isolated.
Empty.	Contains data.	Successful.	The merging of the local and adjacent databases.
Contains data.	Empty.	Successful.	The local database information populates the adjacent database.

The merge process operates as follows:

- The software compares the protocol versions. If the protocol versions differ, the ISL is isolated.
- If the protocol versions are the same, then the zone policies are compared. If the zone policies (includes **Default zoning**: *permit/deny*, **Smart-zoning**: *enable/disable* and **Merge policy** - *allow/restrict*) differ, the ISL is isolated.
- If the zone merge options are the same, the comparison is implemented based on the merge control setting.
  - If the setting is restrict, the active zone set and the full zone set should be identical. Otherwise, the link is isolated.
  - If the setting is allow, the merge rules are used to perform the merge.

Configuring Zone Merge Control Policies

You can configure merge control policies.

SUMMARY STEPS

1. **configure terminal**
2. **zone merge-control restrict vsan** *vsan-id*
3. **no zone merge-control restrict vsan** *vsan-id*
4. **zone commit vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**zone merge-control restrict vsan** *vsan-id* **Example:** `switch(config)# zone merge-control restrict vsan 24`	Configures a restricted merge control setting for this VSAN.
Step 3	**no zone merge-control restrict vsan** *vsan-id* **Example:** `switch(config)# no zone merge-control restrict vsan 33`	Defaults to using the allow merge control setting for this VSAN.
Step 4	**zone commit vsan** *vsan-id* **Example:** `switch(config)# zone commit vsan 20`	Commits the changes made to the specified VSAN.

Default Zone Policies

You can permit or deny traffic in the default zone.

SUMMARY STEPS

1. **configure terminal**
2. **zone default-zone permit vsan** *vsan-id*
3. **no zone default-zone permit vsan** *vsan-id*
4. **zone commit vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**zone default-zone permit vsan** *vsan-id* **Example:** `switch(config)# zone default-zone permit vsan 12`	Permits traffic flow to default zone members.
Step 3	**no zone default-zone permit vsan** *vsan-id* **Example:** `switch(config)# no zone default-zone permit vsan 12`	Denies traffic flow to default zone members and reverts to factory default.
Step 4	**zone commit vsan** *vsan-id* **Example:** `switch(config)# zone commit vsan 340`	Commits the changes made to the specified VSAN.

Configuring System Default Zoning Settings

You can configure default settings for default zone policies and full zone distribution for new VSANs on the switch.

Note The zone's default system settings such as system default zone default-zone permit and system default zone distribute full takes effect only on the newly created VSANs after you manually apply the settings. These settings may not be applied on VSAN 1 even though when they are set as part of FC setup script.

You can configure zone settings using FC script also. For more information about configuring default zone settings using FC script, see: *Cisco Nexus 9000 Series NX-OS Fundamentals Configuration Guide*.

SUMMARY STEPS

1. **configure terminal**
2. **system default zone default-zone permit**
3. **no system default zone default-zone permit**
4. **system default zone distribute full**
5. **no system default zone distribute full**

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**system default zone default-zone permit** **Example:** `switch(config)# system default zone default-zone permit`	Configures permit as the default zoning policy for new VSANs on the switch.
Step 3	**no system default zone default-zone permit** **Example:** `switch(config)# no system default zone default-zone permit`	Configures deny (default) as the default zoning policy for new VSANs on the switch.
Step 4	**system default zone distribute full** **Example:** `switch(config)# system default zone distribute full`	Enables full zone database distribution as the default for new VSANs on the switch.
Step 5	**no system default zone distribute full** **Example:** `switch(config)# no system default zone distribute full`	Disables (default) full zone database distribution as the default for new VSANs on the switch. Only the active zone database is distributed.

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 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 22: 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:

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:

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.

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

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

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 these steps:

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:

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:

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.

Default Settings for Zones

The following table lists the default settings for basic zone parameters.

Table 23: Default Basic Zone Parameters

Parameters	Default
Default zone policy	Denied to all members.
Full zone set distribute	The full zone set(s) is not distributed.
Enhanced zoning	Disabled.

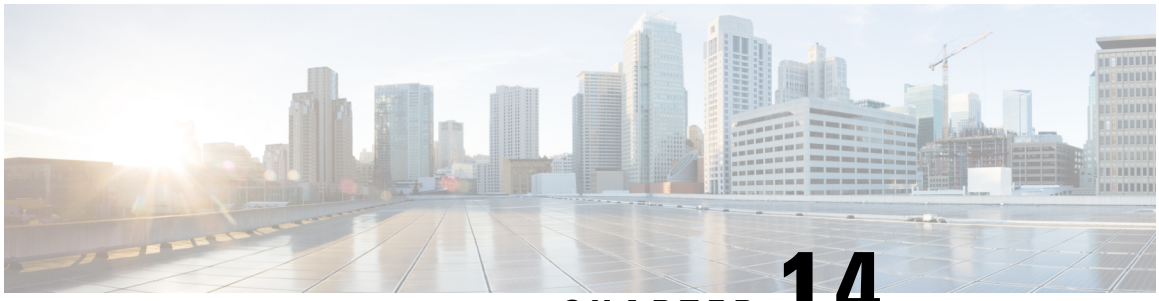

CHAPTER 14

Advanced Fibre Channel Features

This chapter describes how to configure advanced Fibre Channel features.

This chapter includes the following sections:

- Advanced Fibre Channel Features and Concepts, on page 185

Advanced Fibre Channel Features and Concepts

Fibre Channel Timeout Values

You can modify Fibre Channel protocol-related timer values for the switch by configuring the following timeout values (TOVs):

- Distributed services TOV (D_S_TOV)—The valid range is from 5,000 to 10,000 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.

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.

You can configure Fibre Channel timers across all VSANs.

SUMMARY STEPS

1. **configure terminal**
2. **fctimer R_A_TOV** *timeout*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**fctimer R_A_TOV** *timeout* **Example:** `switch(config)# fctimer R_A_TOV 8008000`	Configures the R_A_TOV timeout value for all VSANs. The unit is milliseconds. 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 such as Fibre Channel. 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.

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.

You can configure per-VSAN Fibre Channel timers.

SUMMARY STEPS

1. **configure terminal**
2. **fctimer D_S_TOV** *timeout* **vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	**configure terminal** **Example:**	Enters global configuration mode.

	Command or Action	Purpose
	`switch# configure terminal` `switch(config)#`	
Step 2	**fctimer D_S_TOV** *timeout* **vsan** *vsan-id* **Example:** `switch(config#)# fctimer D_S_TOV 9009000 vsan 15`	Configures the D_S_TOV timeout value (in milliseconds) for the specified VSAN. Suspends the VSAN temporarily. You have the option to end this command, if required.

EXAMPLES

This example shows how to configure the timer value for VSAN 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
```

fctimer Distribution

You can enable per-VSAN fctimer fabric distribution for all Cisco SAN 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 enter 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.

Note CFS is enabled by default. All devices in the fabric must have CFS enabled, or they do not receive distributions.If CFS is disabled for an application, that application does not distribute any configuration, and it does not accept a distribution from other devices in the fabric. You can enable CFS using **cfs distribute** command.

Enabling or Disabling fctimer Distribution

You can enable or disable fctimer fabric distribution.

SUMMARY STEPS

1. **configure terminal**
2. **fctimer distribute**
3. **no fctimer distribute**

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**fctimer distribute** **Example:** `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	**no fctimer distribute** **Example:** `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.

SUMMARY STEPS

1. **configure terminal**
2. **fctimer commit**

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**fctimer commit** **Example:** `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.

SUMMARY STEPS

1. **configure terminal**

2. **fctimer abort**

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**fctimer abort** **Example:** `switch(config)# fctimer abort`	Discards the fctimer configuration changes in the pending database and releases the fabric lock.

Overriding the Fabric Lock

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.

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
```

Fabric Database Merge Guidelines

When merging two fabrics, follow these guidelines:

- Be aware of the following merge conditions:
  - The merge protocol is not implemented for distribution of the fctimer values. You must manually merge the fctimer values when a fabric is merged.
  - The per-VSAN fctimer configuration is distributed in the physical fabric.
  - The fctimer configuration is only applied to those switches containing the VSAN with a modified fctimer value.
  - The global fctimer values are not distributed.
- Do not configure global timer values when distribution is enabled.

Note The number of pending fctimer configuration operations cannot be more than 15. After 15 operations, you must commit or abort the pending configurations before performing any more operations.

Verifying Configured fctimer Values

Use the **show fctimer** command to display the configured fctimer values. The following example displays the 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.

The following example displays the configured TOV for VSAN 10:

```
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
```

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.

Cisco SAN switches support three network address authority (NAA) address formats. (see the following table).

Table 24: Standardized NAA WWN Formats

NAA Address	NAA Type	WWN Format	
IEEE 48-bit address	Type 1 = 0001b	000 0000 0000b	48-bit MAC address
IEEE extended	Type 2 = 0010b	Locally assigned	48-bit MAC address
IEEE registered	Type 5 = 0101b	IEEE company ID: 24 bits	VSID: 36 bits

Caution Changes to the world-wide names should be made by an administrator or individual who is completely familiar with switch operations.

Verifying the WWN Configuration

Use the **show wwn** commands to display the status of the WWN configuration. This example shows how to display the status of all WWNs:

```
switch# show wwn status
Type    Configured     Available    Resvd.  Alarm State
----    ----------  --------------  ------  -----------
   1           64        48 ( 75%)      16    NONE
 2,5       524288    442368 ( 84%)   73728    NONE
```

This example shows how to display the information for block ID 51:

```
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
```

This example shows how to display 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. ELPs and EFPs both 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.

Configuring a Secondary MAC Address

You can allocate secondary MAC addresses.

SUMMARY STEPS

1. **configure terminal**
2. **wwn secondary-mac** *wwn-id* **range** *value*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**wwn secondary-mac** *wwn-id* **range** *value* **Example:** `switch(config)# wwn secondary-mac 33:e8:00:05:30:00:16:df range 55`	Configures the secondary MAC address. This command cannot be undone.

EXAMPLES

This example shows how to configure the secondary MAC address:

```
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
```

FC ID Allocation for HBAs

Fibre Channel standards require a unique FC ID to be allocated to an N port attached to an F port in any switch. To conserve the number of FC IDs used, Cisco SAN switches use a special allocation scheme.

Some HBAs do not discover targets that have FC IDs with the same domain and area. The switch software maintains a list of tested company IDs that do not exhibit this behavior. These HBAs are allocated with single FC IDs. If the HBA can discover targets within the same domain and area, a full area is allocated.

To allow further scalability for switches with numerous ports, the switch software maintains a list of HBAs that can discover targets within the same domain and area. Each HBA is identified by its company ID (also known as Organizational Unique Identifier, or OUI) used in the pWWN during a fabric log in. 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 type (whole area or single) of FC ID allocated, the FC ID entries remain persistent.

Default Company ID List

All Cisco SAN switches 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 enter a **write erase**, the list inherits the default list of company IDs shipped with a relevant release.

Verifying the Company ID Configuration

You can view the configured company IDs by entering the **show fcid-allocation area** command. 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.

This example shows how to display 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 entering the **show fcid-allocation company-id-from-wwn** command. Some WWN formats do not support company IDs. In these cases, you many need to configure the FC ID persistent entry.

This example shows how to display 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
```

Switch Interoperability

Interoperability enables the products of multiple vendors to interwork with each other. Fibre Channel standards guide vendors towards common external Fibre Channel interfaces.

Not all vendors follow the standards in the same way, which results in the need for interoperability modes. This section briefly explains the basic concepts of these modes.

Each vendor has a regular mode and an equivalent interoperability mode, which specifically turns off advanced or proprietary features and provides the product with a standards-compliant implementation.

About Interop Mode

The software supports only one interop mode (Mode 3—Brocade native mode (Core PID 1)). In mode 3 of the interop mode, you can seamlessly add Brocade switches with Core PID 1 (Brocade native mode) without altering their native modes. All the existing functionalities remains same.

The following table lists the changes in switch operation when you enable interoperability mode.

Table 25: Changes in Switch Operation When Interoperability Is Enabled

Switch Feature	Changes if Interoperability Is Enabled
Domain IDs	Domain IDs can either be static or preferred, which operate as follows: • Static: Cisco switches accept only one domain ID; if a switch does not get that domain ID it isolates itself from the fabric. • Preferred: If the switch does not get its requested domain ID, it accepts any assigned domain ID.
Timers	All Fibre Channel timers must be the same on all switches as these values are exchanged by E ports when establishing an ISL. The timers are F_S_TOV, D_S_TOV, E_D_TOV, and R_A_TOV.
F_S_TOV	Verify that the Fabric Stability Time Out Value timers match exactly.
D_S_TOV	Verify that the Distributed Services Time Out Value timers match exactly.
E_D_TOV	Verify that the Error Detect Time Out Value timers match exactly.
R_A_TOV	Verify that the Resource Allocation Time Out Value timers match exactly.
Trunking	Trunking is not supported between two different vendor's switches. This feature may be disabled per port or per switch.
Default zone	The default zone operation of permit (all nodes can see all other nodes) or deny (all nodes are isolated when not explicitly placed in a zone) may change.
Zoning attributes	Zones may be limited to the pWWN and other proprietary zoning methods (physical port number) may be eliminated. **Note** On a Brocade switch, use the **cfgsave** command to save fabric-wide zoning configuration. This command does not have any effect on Cisco SAN switches if they are part of the same fabric. You must explicitly save the configuration on each Cisco SAN switch.

Switch Feature	Changes if Interoperability Is Enabled
Zone propagation	Some vendors do not pass the full zone configuration to other switches, only the active zone set gets passed. Verify that the active zone set or zone configuration has correctly propagated to the other switches in the fabric.
VSAN	Interop mode only affects the specified VSAN.
TE ports and SAN port channels	TE ports and SAN port channels cannot be used to connect Cisco switches to non-Cisco SAN switches. Only E ports can be used to connect to non-Cisco SAN switches. TE ports and SAN port channels can still be used to connect a Cisco switch to other Cisco SAN switches even when in interop mode.
FSPF	The routing of frames within the fabric is not changed by the introduction of interop mode. The switch continues to use src-id, dst-id, and ox-id to load balance across multiple ISL links.
Domain reconfiguration disruptive	This is a switch-wide impacting event. Brocade and McData require the entire switch to be placed in offline mode and/or rebooted when changing domain IDs.
Domain reconfiguration nondisruptive	This event is limited to the affected VSAN. Cisco SAN switches have the capability to restart only the domain manager process for the affected VSAN and not the entire switch.
Name server	Verify that all vendors have the correct values in their respective name server database.

Configuring Interop Mode 3

You can configure interop mode3 in Cisco SAN switches disruptively or nondisruptively.

Note Brocade's **msplmgmtdeactivate** command must explicitly be run prior to connecting from a Brocade switch to Cisco SAN switches. This command uses Brocade proprietary frames to exchange platform information, which Cisco SAN switches do not recognize. Rejecting these frames causes the common E ports to become isolated.

Procedure

	Command or Action	Purpose

Step 1	Place the VSAN of the E ports that connect to the OEM switch in interoperability mode.	`switch# configuration terminal`
		`switch(config)# vsan database`
		`switch(config-vsan-db)# vsan 10 interop 3`
		`switch(config-vsan-db)# exit`

	Command or Action	Purpose

Step 2	Change the Fibre Channel timers (if they have been changed from the system defaults).	**Note** The Cisco SAN switches, Brocade, and McData FC Error Detect (ED_TOV) and Resource Allocation (RA_TOV) timers default to the same values. They can be changed if needed. The RA_TOV default is 10 seconds, and the ED_TOV default is 2 seconds. Per the FC-SW2 standard, these values must be the same on each switch within the fabric.
		`switch(config)# fctimer e_d_tov ?`
		`<1000-100000> E_D_TOV in milliseconds(1000-100000)`
		`switch(config)# fctimer r_a_tov ?`
		`<1000-4000> E_D_TOV in milliseconds(1000-4000)`

Step 3	When making changes to the domain, you may or may not need to restart the Domain Manager function for the altered VSAN.	• Force a fabric reconfiguration with the **disruptive** option.
		`switch(config)# fcdomain restart disruptive vsan 1`
		or
		• Do not force a fabric reconfiguration.
		`switch(config# fcdomain restart vsan 10`

Default Settings for Advanced Fibre Channel Features

The following table lists the default settings for the features included in this chapter.

Table 26: 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

Parameters	Default
Remote capture connection protocol	TCP
Remote capture connection mode	Passive
Local capture frame limits	10 frames
FC ID allocation mode	Auto mode
Loop monitoring	Disabled
Interop mode	Disabled

INDEX

* (asterisk) 89
  first operational port[asterisk (asterisk) 89
    first operational port] 89

A

active zone sets 156, 165
  considerations 156
  enabling distribution 165
address allocation cache 109
  description 109
administrative speeds 52
  configuring 52
administrative states 39
  description 39
auto mode 49
  configuring 49
auto port mode 39
  description 39
autosensing speed 54

B

BB_credits 42, 60
  description 42
  displaying information 60
  reason codes 42
bit error thresholds 55
  configuring 55
  description 55
bit errors 55
  reasons 55
Brocade 194
  native interop mode 194
buffer-to-buffer credits 42
build fabric frames 92
  description 92

C

changed information 1
  description 1
company IDs 192
  FC ID allocations 192
configuring 158
  zones example 158
Contiguous Domain ID Assignments 103
  About 103
creating 116, 118
  virtual fibre channel interfaces 116, 118

D

default zones 161, 194
  description 161
  interoperability 194
  policies 161
destination IDs 77
  exchange based 77
  flow based 77
device alias databases 148–151
  disabling distribution 150
  discarding changes 149
  enabling distribution 150
  locking the fabric 148
  merging 151
device aliases 143–145, 151–152
  creating 144
  default settings 152
  description 143
  displaying information 151
  displaying zone set information 151
  enhanced mode 145
  features 143
  modifying databases 144
  requirements 144
disabling 10
  FCoE 10
domain IDs 40, 91, 96, 99–100, 103–104, 162, 194
  allowed lists 99
  assignment failures 40
  configuring allowed lists 99
  configuring CFS distribution 100
  configuring fcalias members 162
  contiguous assignments 103
  description 96
  distributing 91
  enabling contiguous assignments 103–104
  interoperability 194

domain IDs *(continued)*
preferred 96
static 96
domain manager 40
isolation 40

E

E port mode 38
classes of service 38
description 38
E ports 40, 49, 167
configuring 49
isolation 40
recovering from link isolations 167
EISLs 75
SAN port channel links 75
ELP 40
enhanced zones 171–173, 176–177
advantages over basic zones 171
changing from basic zones 172
configuring default full database distribution 177
configuring default policies 176
configuring default switch-wide zone policies 177
description 171
modifying database 173
exchange link parameter 40
expansion port mode 38

F

F port mode 38
classes of service 38
description 38
F ports 38, 49
configuring 49
description 38
fabric login 127
fabric port mode 38
fabric pWWNs 153
zone membership 153
fabric reconfiguration 91
fcdomain phase 91
Fabric-Device Management Interface 132
fabrics 92
FC IDs 91, 103–104, 162, 192
allocating 91
allocating default company ID lists 192
configuring fcalias members 162
description 103
persistent 104
fcaliases 162, 168–169
cloning 169
configuring for zones 162
creating 162
fcaliases *(continued)*
renaming 168
fcdomains 40, 91, 93, 95–96, 100, 109–110
autoreconfigured merged fabrics 95
configuring CFS distribution 100
default settings 110
description 91
displaying information 109
displaying statistics 109
domain IDs 96
enabling autoreconfiguration 96
incoming RCFs 95
overlap isolation 40
restarts 91
switch priorities 93
FCoE 10–11
disable LAN traffic 11
disabling 10
fctimers 190
displaying configured values 190
FDMI 132
description 132
displaying database information 132
Fibre Channel 185
timeout values 185
TOV 185
Fibre Channel domains 91
Fibre Channel interfaces 39–40, 42, 48–50, 52, 55, 61
administrative states 39
BB_credits 42
configuring 48
configuring auto port mode 49
configuring bit error thresholds 55
configuring descriptions 50
configuring port modes 49
configuring range 48
configuring speeds 52
default settings 61
operational states 40
reason codes 40
states 39
FLOGI 127
description 127
FSPF 194
interoperability 194
full zone sets 156, 165
considerations 156
enabling distribution 165
fWWNs 162
configuring fcalias members 162
Fx ports 38, 66
VSAN membership 66

H

hard zoning 165
  description 165
HBA ports 106
  configuring area FCIDs 106

I

interfaces 50, 58, 68–69, 84–85, 162
  adding to SAN port channels 84–85
  assigning to VSANs 69
  configuring descriptions 50
  configuring fcalias members 162
  displaying SFP information 58
  isolated states 85
  SFP types 58
  suspended states 85
  VSAN membership 68
Interfaces 39
interop modes 194, 196
  configuring mode 1 194
  default settings 196
  description 194
interoperability 73, 194
  configuring interop mode 1 194
  description 194
  VSANs 73
ISLs 75
  SAN port channel links 75
isolated VSANs 70–71
  description 70
  displaying membership 71

L

layer 2 interfaces 51
  unified ports 51
load balancing 67, 75, 77
  attributes for VSANs 67
  description 77
  SAN port channels 75

M

MAC addresses 191
  configuring secondary 191
McData 194
  native interop mode 194
merged fabrics 95
  autoreconfigured 95

N

N port identifier virtualization 56
N ports 153, 165
  hard zoning 165
  zone enforcement 165
  zone membership 153
N5K-M1008 expansion module 47
N5K-M1404 expansion module 47
name servers 128–129, 194
  displaying database entries 129
  interoperability 194
  proxy feature 128
  registering proxies 128
new information 1
  description 1
NPIV 56
  description 56
  enabling 56

O

operational states 40, 49
  configuring on Fibre Channel interfaces 49
  description 40

P

persistent FC IDs 104–105, 107, 109
  configuring 105
  description 104
  displaying 109
  enabling 105
  purging 107
PLOGI 129
  name server 129
port channels 40, 194
  administratively down 40
  interoperability 194
port modes 39
  auto 39
port speeds 52
  configuring 52
port world wide names 153
ports 68
  VSAN membership 68
principal switches 96, 99
  assigning domain ID 96
  configuring 99
proxies 128
  registering for name servers 128
pWWNs 153, 162
  configuring fcalias members 162
  zone membership 153

R

RCFs 92, 95
  description 92
  incoming 95
  rejecting incoming 95
reason codes 40
  description 40
reconfigure fabric frames 92
redundancy 66
  VSANs 66
Registered State Change Notifications 134
RSCN 134, 136, 141
  default settings 141
  description 134
  displaying information 134
  multiple port IDs 134
  suppressing domain format SW-RSCNs 136
RSCN timers 136, 138
  configuration distribution using CFS 138
  configuring 136

S

SAN port channel 89
  verifying configurations 89
SAN port Channel 90
  default settings 90
SAN port channel Protocol 87
  creating channel group 87
SAN port channels 75–77, 80, 84–85
  adding interfaces 84–85
  comparison with trunking 76
  compatibility checks 85
  configuration guidelines 80
  description 75
  interface states 85
  load balancing 77
  misconfiguration error detection 80
scalability 66
  VSANs 66
SCR 134
  request 134
SD ports 49
  configuring 49
secondary MAC addresses 191
  configuring 191
SFPs 58
  displaying transmitter types 58
  transmitter types 58
smart zoning 179
soft zoning 165
  description 165
source IDs 77
  exchange based 77
  flow based 77
storage devices 153
  access control 153
switch ports 56
  configuring attribute default values 56
switch priorities 93
  default 93
  description 93

T

TE port mode 38
  classes of service 38
  description 38
TE ports 167, 194
  interoperability 194
  recovering from link isolations 167
timeout values 185
TOV 185–186, 194, 196
  configuring across all VSANs 185
  configuring for a VSAN 186
  default settings 196
  interoperability 194
  ranges 185
traffic isolation 66
  VSANs 66
trunk mode 56
  administrative default 56
trunking 76, 194
  comparison with port channels 76
  interoperability 194
trunking E port mode 38
trunking ports 69
  associated with VSANs 69

U

unified ports 51
  configuring 51
unique area FC IDs 106
  configuring 106
  description 106

V

Virtual Fibre Channel interfaces 61
  default settings 61
VSAN IDs 38, 66–67
  description 67
  multiplexing traffic 38
  range 66
  VSAN membership 66
VSANs 38, 40, 63, 66–71, 73, 96, 109, 128, 156, 185, 194
  advantages 63
  cache contents 109
  comparison with zones (table) 66

VSANs *(continued)*
configuring 68
default settings 73
deleting 71
description 63
displaying configuration 73
displaying membership 69
displaying usage 73
domain ID automatic reconfiguration 96
FC IDs 63
features 63
interop mode 194
isolated 70
load balancing attributes 67
mismatches 40
multiple zones 156
name server 128
names 67
operational states 71
port membership 68
states 67
TE port mode 38
timer configuration 185
TOV 185
traffic isolation 63
trunking ports 69

W

world wide names 190
WWNs 40, 190–191
description 190
displaying information 190
link initialization 191
secondary MAC addresses 191
suspended connections 40

Z

zone attribute groups 169
cloning 169
zone databases 170, 174
migrating a non-Cisco SAN database 170
zone databases *(continued)*
release locks 174
zone members 161
displaying information 161
zone server databases 170
clearing 170
zone sets 153, 156, 159–160, 165–170
activating 160
cloning 169
considerations 156
creating 159
displaying information 170
distributing configuration 165
enabling distribution 165
exporting 167
exporting databases 167
features 153
importing 167
importing databases 167
one-time distribution 166
recovering from link isolations 167
renaming 168
viewing information 170
zones 40, 66, 153, 155, 159, 162, 167–170
access control 159
backing up (procedure) 168
cloning 169
comparison with VSANs (table) 66
configuring aliases 162
configuring fcaliases 162
default policies 153
displaying information 170
exporting databases 167
features 153, 155
importing databases 167
membership using pWWNs 66
merge failures 40
renaming 168
restoring (procedure) 168
viewing information 170
zoning 153, 155
description 153
example 155
implementation 155