Cisco Nexus 5000 Series NX-OS SAN Switching Configuration Guide, Release 4.2(1)N1(1)

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

Text Part Number: OL-20923-02

CONTENTS

Preface

This preface describes the audience, organization, and conventions of the *Cisco Nexus 5000 Series NX-OS SAN Switching Configuration Guide*. It also provides information on how to obtain related documentation.

Audience

This preface describes the audience, organization, and conventions of the . It also provides information on how to obtain related documentation.

Document Organization

This document is organized into the following chapters:

Chapter	Description
New and Changed Information	Describes the new and changed information for the new Cisco NX-OS software releases.
Overview	Provides an overview of all the features in this guide.
Configuring Fibre Channel Interfaces	Describes the licensing requirements, configuration information, interface modes, and interface states for Fibre Channel interfaces.
Configuring Domain Parameters	Describes configuration information for domain parameters, domain IDs, and FC IDs.

Chapter	Description
Configuring N Port Virtualization	Provides an overview of N Port Virtualization, information on how to configure NPV interfaces, and includes quidelines and requirements for configuring and verifying NPV.
Configuring VSAN Trunking	Explains TE ports and trunking concepts and describes how to configure VSAN trunking how to enable trunking protocols.
Configuring SAN Port Channel	Explains SAN PortChannels and load balancing concepts and provides details on configuring SAN PortChannels, and adding or deleting ports to SAN PortChannels.
Configuring and Managing VSANs	Describes how virtual SANs (VSANs) work, explains the concept of default VSANs, isolated VSANs, VSAN IDs, and attributes, and provides details on how to create, delete, and view VSANs.
Configuring and Managing Zones	Defines various zoning concepts and provides details on configuring a zone set and zone management features.
Distributing Device Alias Services	Describes the use of the Distributed Device Alias Services (device alias) to distribute device alias names on a fabric-wide basis.
Configuring Fibre Channel Routing Services and Protocols	Provides details and configuration information on Fibre Channel routing services and protocols.
Managing FLOGI, Name Server, FDMI, an RSCN Databases	Provides name server and fabric login details required to manage storage devices and display registered state change notification (RSCN) databases.
Discovering SCSI Targets	Describes how the SCSI LUN discovery feature is started and displayed.
Advanced Fibre Channel Features and Concepts	Describes the advanced configuration features-time out values, fctrace, fabric analyzer, world wide names, flat FC IDs, loop monitoring, and interoperating switches.
Configuring FC-SP and DHCHAP	Describes the DHCHAP protocol, an FC-SP protocol, that provides authentication between Cisco MDS 9000 Family switches and other devices.
Configuring Port Security	Provides details on port security features that can prevent unauthorized access to a switch port in the Cisco MDS 9000 Family.
Configuring Fabric Binding	Describes the fabric binding security feature for VSANs, which ensures that ISLs are only enabled between specific switches.
Configuring Fabric Configuration Servers	Describes how the fabric configuration server (FCS) feature is configured and displayed.
Configuring Port Tracking	Describes the port tracking feature and provides information to enable port tracking and to configure linked ports.

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 the user supplies the values.
[x]	Square brackets enclose an optional element(keyword or argument).
[x | y]	Square brackets enclosing keywords or arguments separated by a vertical bar indicate an optional choice.
{x | y}	Braces enclosing keywords or arguments separated by a vertical bar indicate a required choice.
[x {y | z}]	Nested set of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element.
`variable`	Indicates a variable for which you supply values, in context where italics cannot be used.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.

Screen examples use the following conventions:

Convention	Description
`screen font`	Terminal sessions and information the switch displays are in screen font.
`boldface screen font`	Information you must enter is in boldface screen font.
`italic screen font`	Arguments for which you supply values are in italic screen font.
< >	Nonprinting characters, such as passwords, are in angle brackets.
[ ]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

This document uses the following conventions:

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

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

Related Documentation for Nexus 5000 Series NX-OS Software

Cisco NX-OS documentation is available at the following URL:

http://www.cisco.com/en/US/products/ps9670/tsd_products_support_series_home.htmll

The documentation set for the Cisco Nexus 5000 Series NX-OS software includes the following documents:

Release Notes

- *Cisco Nexus 5000 Series and Cisco Nexus 2000 Series Release Notes*
- *Cisco Nexus 5000 Series Switch Release Notes*

Cisco Nexus 5000 Series NX-OS Configuration Guides

- *Cisco Nexus 5000 Series NX-OS Fibre Channel over Ethernet Configuration Guide*
- *Cisco Nexus 5000 Series NX-OS Layer 2 Switching Configuration Guide*
- *Cisco Nexus 5000 Series NX-OS Quality of Service Configuration Guide*
- *Cisco Nexus 5000 Series NX-OS SAN Switching Configuration Guide*
- *Cisco Nexus 5000 Series NX-OS Security Configuration Guide*
- *Cisco Nexus 5000 Series NX-OS System Management Configuration Guide*
- *Cisco Nexus 5000 Series Switch CLI Software Configuration Guide*
- *Cisco Nexus 5000 Series Fabric Manager Configuration Guide, Release 3.4(1a)*

Installation and Upgrade Guides

- *Cisco Nexus 5000 Series Hardware Installation Guide*
- *Regulatory Compliance and Safety Information for the Cisco Nexus 5000 Series*

Cisco NX-OS Command References

- *Cisco Nexus 5000 Series Command Reference*

Cisco NX-OS Technical References

- *Cisco Nexus 5000 MIBs Reference*

Cisco NX-OS Error and System Messages

- *Cisco NX-OS System Messages Reference*

Obtaining Documentation and Submitting a Service Request

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

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

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

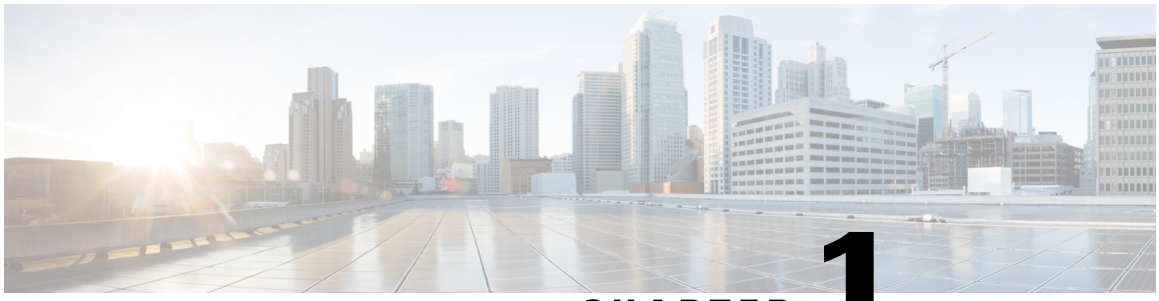

CHAPTER 1

New and Changed Information

This chapter provides release-specific information for each new and changed feature in the *Cisco Nexus 5000 Series NX-OS SAN Switching Configuration Guide*.

New and Changed Information

This chapter provides release-specific information for each new and changed feature in the *Cisco Nexus 5000 Series NX-OS SAN Switching Configuration Guide*.

The latest version of this document is available at the following Cisco website:

http://www.cisco.com/en/US/products/ps9670/products_installation_and_configuration_guides_list.html

To check for additional information about Cisco NX-OS, see the *Cisco Nexus 5000 Series NX-OS Release Notes* available at the following Cisco website:

http://www.cisco.com/en/US/products/ps9670/prod_release_notes_list.html

This table summarizes the new and changed features documented in the *Cisco Nexus 5000 Series NX-OS SAN Switching Configuration Guide, Release 5.0(2)N2(1)*, and tells you where they are documented.

Table 1: New and Changed SAN Switching Features for Cisco NX-OS Release 5.0(2)N2(1)

Feature	Description	Changed in Release	Where to find it documented...
VSAN Trunking	Added information abut configuring virtual Fibre Channel interfaces in trunking mode.	5.0(2)N2(1)	Configuring VSAN Trunking
VE Port	Added information about configuring multi-hop FCoE fabric using VE ports.	5.0(2)N2(1)	Configuring Fibre Channel Interfaces

This table summarizes the new and changed features documented in the *Cisco Nexus 5000 Series NX-OS SAN Switching Configuration Guide, Release 4.2(1)N1(1)*, and tells you where they are documented.

Table 2: New and Changed SAN Switching Features for Cisco NX-OS Release 4.2(1)N1(1)

Feature	Description	Changed in Release	Where to find it documented...
F Port Trunking	Added information about TF Ports and TNP Ports.	4.2(1)N1(1)	Configuring Fibre Channel Interfaces
F Port Channels	Added information about F and TF Port Channels.	4.2(1)N1(1)	Configuring SAN Port Channels

Documentation Organization

As of Cisco NX-OS Release 4.2(1)N1(1), the Nexus 5000 Series configuration information is available in new feature-specific configuration guides for the following information:

- System Management
- Layer 2 Switching
- SAN Switching
- Fibre Channel over Ethernet
- Security
- Quality of Service

The information in these new guides previously existed in the *Cisco Nexus 5000 Series CLI Configuration Guide* which remains available on Cisco.com and should be used for all software releases prior to Cisco Nexus 5000 NX-OS Software Rel 4.1(3). Each new configuration guide addresses the features that are introduced in or are available in a particular release. Select and view the configuration guide that pertains to the software installed in your switch.

The information in the new *Cisco Nexus 5000 Series NX-OS SAN Switching Configuration Guide* previously existed in Part 7: SAN Switching of the *Cisco Nexus 5000 Series CLI Configuration Guide*.

For a complete list of Nexus 5000 Series document titles, see the list of Related Documentation in the "Preface."

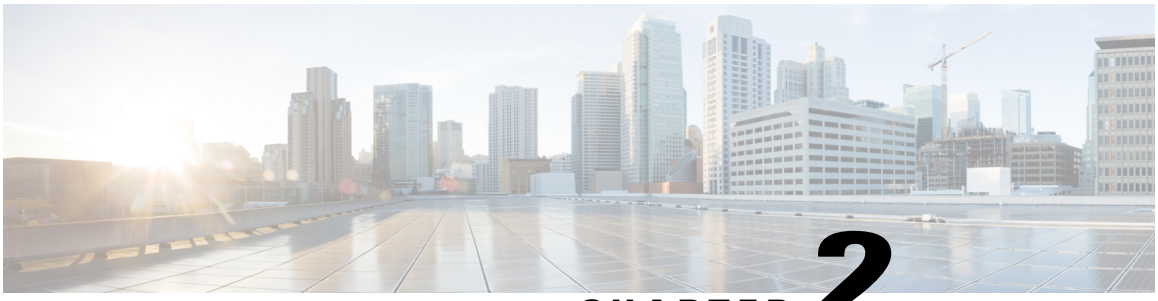

CHAPTER 2

Overview

The Cisco Nexus 5000 Series NX-OS software can configure and manage features such as VSANs, SAN device virtualization, dynamic VSANs, zones, distributed device alias services, Fibre Channel routing services and protocols, FLOGI, name server, FDMI, RSCN database, SCSI targets, FICON, and other advanced features described in this guide.

SAN Switching Overview

The SAN switching features documented in this guide are described below.

Fibre Channel Interfaces

Fibre Channel ports are optional on the Cisco Nexus 5000 Series switch. When you use expansion modules up to 8 Fibre Channel ports are available on the Cisco Nexus 5010 switch and up to 16 Fibre Channel ports are available on the Cisco Nexus 5020 switch.

Each Fibre Channel port can be used as a downlink (connected to a server) or as an uplink (to the data center SAN fabric).

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.

N Port Virtualization

Cisco NX-OS software supports industry-standard N port identifier virtualization (NPIV), which allows multiple N port fabric logins concurrently on a single physical Fibre Channel link. HBAs that support NPIV can help improve SAN security by enabling zoning and port security to be configured independently for each virtual machine (OS partition) on a host. In addition to being useful for server connections, NPIV is beneficial for connectivity between core and edge SAN switches.

N port virtualizer (NPV) is a complementary feature that reduces the number of Fibre Channel domain IDs in core-edge SANs. Cisco MDS 9000 family 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. This feature is available only for Cisco MDS Blade Switch Series, the Cisco MDS 9124 Multilayer Fabric Switch, and the Cisco MDS 9134 Multilayer Fabric Switch.

VSAN Trunking

Trunking, also known as VSAN trunking, is a feature specific to switches in the Cisco MDS 9000 Family. 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.

SAN Port Channels

PortChannels aggregate multiple physical ISLs into one logical link with higher bandwidth and port resiliency for both Fibre Channel and FICON traffic. With this feature, up to 16 expansion ports (E-ports) or trunking E-ports (TE-ports) can be bundled into a PortChannel. ISL ports can reside on any switching module, and they do not need a designated master port. If a port or a switching module fails, the PortChannel continues to function properly without requiring fabric reconfiguration.

Cisco NX-OS software uses a protocol to exchange PortChannel configuration information between adjacent switches to simplify PortChannel management, including misconfiguration detection and autocreation of PortChannels among compatible ISLs. In the autoconfigure mode, ISLs with compatible parameters automatically form channel groups; no manual intervention is required.

PortChannels load balance Fibre Channel traffic using a hash of source FC-ID and destination FC-ID, and optionally the exchange ID. Load balancing using PortChannels is performed over both Fibre Channel and FCIP links. Cisco NX-OS software also can be configured to load balance across multiple same-cost FSPF routes.

Virtual SANs

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

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

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

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

Zoning

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

- N port zoning-Defines zone members based on the end-device (host and storage) port.
    - WWN

- Fibre Channel identifier (FC-ID)

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

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

Device Alias Services

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

Fibre Channel Routing

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

Dynamically compute routes throughout a fabric by establishing the shortest and quickest path between any two switches.

Select an alternative path in the event of the failure of a given path. FSPF supports multiple paths and automatically computes an alternative path around a failed link. FSPF provides a preferred route when two equal paths are available.

SCSI Targets

Small Computer System Interface (SCSI) targets include disks, tapes, and other storage devices. These targets do not register logical unit numbers (LUNs) with the name server. The SCSI LUN discovery feature is initiated on demand, through CLI or SNMP. This information is also synchronized with neighboring switches, if those switches belong to the Cisco Nexus 5000 Series.

Advanced Fibre Channel Features

Fibre Channel protocol-related timer values can be configured 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. Cisco Nexus 5000 Series switches support three network address authority (NAA) address formats

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 Nexus 5000 Series switches use a special allocation scheme.

FC-SP and DHCHAP

Fibre Channel Security Protocol (FC-SP) capabilities provide switch to switch and hosts to switch authentication to overcome security challenges for enterprise-wide fabrics. Diffie-Hellman Challenge Handshake Authentication Protocol (DHCHAP) is an FC-SP protocol that provides authentication between Cisco MDS 9000 Family switches and other devices. DHCHAP consists of the CHAP protocol combined with the Diffie-Hellman exchange.

With FC-SP, switches, storage devices, and hosts are able to prove their identity through a reliable and manageable authentication mechanism. With FC-SP, Fibre Channel traffic can be secured on a frame-by-frame basis to prevent snooping and hijacking, even over untrusted links. A consistent set of policies and management actions are propagated through the fabric to provide a uniform level of security across the entire fabric.

Port Security

The port security feature prevents unauthorized access to a switch port by binding specific world-wide names (WWNs) that have access to one or more given switch ports.

When port security is enabled on a switch port, all devices connecting to that port must be in the port security database and must be listed in the database as bound to a given port. If both of these criteria are not met, the port will not achieve an operationally active state and the devices connected to the port will be denied access to the SAN.

Fabric Binding

The fabric binding feature ensures Inter-Switch Links (ISLs) are enabled only between specified switches in the fabric binding configuration. This feature helps prevent unauthorized switches from joining the fabric or disrupting the current fabric operations. This feature uses the Exchange Fabric Membership Data (EEMD) protocol to ensure that the list of authorized switches is identical in all of the switches in a fabric.

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. In the Cisco MDS 9000 Family switch environment, multiple VSANs constitute a fabric, where one instance of the FCS is present per VSAN.

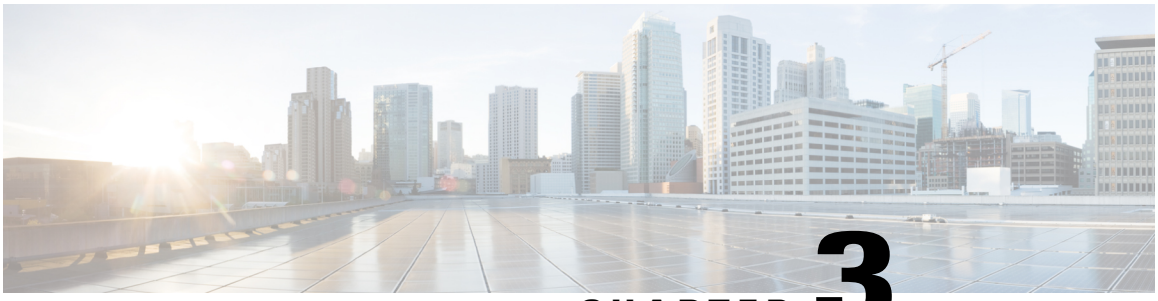

CHAPTER 3

Configuring Fibre Channel Interfaces

This chapter contains the following sections:

Configuring Fibre Channel Interfaces

Information About Fibre Channel Interfaces

Licensing Requirements for Fibre Channel

On Cisco Nexus 5000 Series switches, Fibre Channel capability is included in the Storage Protocol Services license.

Ensure that you have the correct license installed (N5010SS or N5020SS) before using Fibre Channel interfaces and capabilities.

Note You can configure virtual Fibre Channel interfaces without a Storage Protocol Services license, but these interfaces will not become operational until the license is activated.

QOS Requirements for Fibre Channel

The FCoE QoS must be configured if the following types of interfaces are in use:

- Native FC - for FC
- FCoE - for vFC
- FC and FCoE - for FC and vFC

The FCoE QoS must be added even if Ethernet is not configured on the switch.

The following commands will enable the default QoS configuration which must be configured for native FC or FCoE or FC and FCoE:

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

Physical Fibre Channel Interfaces

Cisco Nexus 5000 Series switches support up to sixteen physical Fibre Channel (FC) uplinks through the use of two, optional explansion modules. The first module contains eight FC interfaces. The second module includes four Fibre Channel ports and four Ethernet ports.

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, NP, TE, TF, TNP, SD, and Auto.

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 5000 Series switches, 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.

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.

Interface Modes

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

In NPV mode, Fibre Channel interfaces may operate in NP mode, F mode, or SD 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.

NP 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.

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 5000 Series switch 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 5000 Series switch. 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 5000 Series switches. TF ports support class 3 and class F service.

TNP Port

In trunking NP port (TNP port) mode, an interface functions as a trunking expansion port. A TNP Port may be connected to a trunked F port (TF port) to create a link to a core NPIV switch from an NPV switch.

SD Port

In SPAN destination port (SD port) mode, an interface functions as a switched port analyzer (SPAN). The SPAN feature monitors network traffic that passes though a Fibre Channel interface. This monitoring is done using a standard Fibre Channel analyzer (or a similar switch probe) that is attached to an SD port. SD ports do not receive frames, instead they transmit a copy of the source traffic. The SPAN feature is nonintrusive and does not affect switching of network traffic for any SPAN source ports.

Auto Mode

Interfaces configured in auto mode can operate in one of the following modes: E, F, NP, TE, TF, and TNP port. 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 or Cisco MDS 9000 Family, it may become operational in TE port mode.

SD ports are not determined during initialization and are administratively configured.

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.
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 or 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

Reason Code (long version)	Description	Applicable Modes
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.	
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 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.

In Cisco Nexus 5000 Series switches, the BB_credit mechanism is used on Fibre Channel interfaces but not on virtual Fibre Channel interfaces. Virtual Fibre Channel interfaces provide flow control based on capabilities of the underlying physical Ethernet interface.

The receive BB_credit value (fcrxbbcredit) may be configured for each Fibre Channel interface. In most cases, you do not need to modify the default configuration.

Note The receive BB_credit values depend on the port mode. For physical Fibre Channel interfaces, the default value is 16 for F mode and E mode interfaces. This value can be changed as required. The maximum value is 64.

For virtual Fibre Channel interfaces, BB_credits are not used.

Configuring Fibre Channel Interfaces

Configuring a Fibre Channel Interface

To configure a Fibre Channel interface, perform this task:

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 enable 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

To configure the interface mode, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface** {**fc** *slot*/*port*}|{**vfc** *vfc-id*}
3. switch(config-if)# **switchport mode E** | **F** | **NP** | **TE** | **TF** | **TNP** | **SD** | **auto**

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)# **switchport mode E** | **F** | **NP** | **TE** | **TF** | **TNP** | **SD** | **auto**	For a Fibre Channel interface, you can set the mode to E, F, NP, TE, TF, TNP, or SD port mode. Set the mode to auto to auto-negotiate an E, F, NP, TE, TF, or TNP port mode. **Note** SD ports cannot be configured automatically. They must be administratively configured.

This example shows how to configure VE port 20 and bind it to Ethernet slot 1, port 3:

```
switch# config t
switch(config)# interface vfc 20
switch(config-if)# bind interface ethernet 1/3
switch(config-if)# switchport mode E
switch(config-if)# exit
switch#
```

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

```
switch# show running-config
  interface vfc20
  bind interface Ethernet1/3
  switchport mode E
  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.
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 Port Speeds

Port speed can be configured on a physical Fibre Channel interface but not on a virtual Fibre Channel interface. By default, the port speed for an interface is automatically calculated by the switch.

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 1000**
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 1000**	Configures the port speed of the interface to 1000 Mbps. The number indicates the speed in megabits per second (Mbps). You can set the speed to 1000 (for 1-Gbps interfaces), 2000 (for 2-Gbps interfaces), 4000 (for 4-Gbps interfaces), or auto (default).

	Command or Action	Purpose
Step 4	switch(config-if)# **no switchport speed**	Reverts to the factory default (auto) administrative speed of the interface.

Autosensing

Autosensing speed is enabled on all 4-Gbps interfaces by default. This configuration enables the interfaces to operate at speeds of 1 Gbps, 2 Gbps, or 4 Gbps on the 4-Gbps ports. When autosensing is enabled for an interface operating in dedicated rate mode, 4-Gbps of bandwidth is reserved, even if the port negotiates at an operating speed of 1-Gbps or 2-Gbps.

Configuring SD Port Frame Encapsulation

The **switchport encap eisl** command only applies to SD port interfaces. This command determines the frame format for all frames transmitted by the interface in SD port mode. If the encapsulation is set to EISL, all outgoing frames are transmitted in the EISL frame format, for all SPAN sources.

The **switchport encap eisl** command is disabled by default. If you enable encapsulation, all outgoing frames are encapsulated, and you will see a new line (Encapsulation is eisl) in the **show interface** *SD_port_interface* command output.

Configuring Receive Data Field Size

You can configure the receive data field size for native Fibre Channel interfaces (but not for virtual Fibre Channel interfaces). If the default data field size is 2112 bytes, the frame length will be 2148 bytes.

To configure the receive data field size, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface fc** *slot*/*port*
3. switch(config-if)# **switchport fcrxbufsize 2000**

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.

	Command or Action	Purpose
Step 3	switch(config-if)# **switchport fcrxbufsize 2000**	Reduces the data field size for the selected interface to 2000 bytes. The default is 2112 bytes and the range is from 256 to 2112 bytes.

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.
- GBIC or SFP is specified to operate at 1 Gbps but is used at 2 Gbps.
- GBIC or SFP is specified to operate at 2 Gbps but is used at 4 Gbps.
- 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 Buffer-to-Buffer Credits

To configure BB_credits for a Fibre Channel interface, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface fc** *slot*/*port*
3. switch(config-if)# **switchport fcrxbbcredit default**
4. switch(config-if)# **switchport fcrxbbcredit 5**
5. switch(config-if)# **switchport fcrxbbcredit 5 mode E**
6. switch(config-if)# **switchport fcrxbbcredit 5 mode Fx**
7. switch(config-if# **do show int fc** *slot*/*port*

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 fcrxbbcredit default**	Applies the default operational value to the selected interface. The operational value depends on the port mode. The default values are assigned based on the port capabilities.
Step 4	switch(config-if)# **switchport fcrxbbcredit 5**	Assigns a BB_credit of 5 to the selected interface. The range to assign BB_credits is between 1 and 64.

	Command or Action	Purpose
Step 5	switch(config-if)# **switchport fcrxbbcredit 5 mode E**	Assigns this value if the port is operating in E or TE mode. The range to assign BB_credits is between 1 and 64.
Step 6	switch(config-if)# **switchport fcrxbbcredit 5 mode Fx**	Assigns this value if the port is operating in F mode. The range to assign BB_credits is between 1 and 64.
Step 7	switch(config-if# **do show int fc** *slot*/*port*	Displays the receive and transmit BB_credit along with other pertinent interface information for this interface. **Note** The BB_credit values are correct at the time the registers are read. They are useful to verify situations when the data traffic is slow.

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.

	Command or Action	Purpose
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.

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 1: NPIV Example

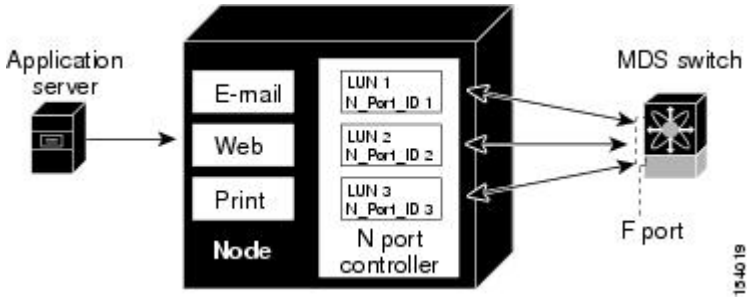

Enabling N Port Identifier Virtualization

To enable or disable NPIV on the switch, perform this task:

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. switch# **configuration terminal**
2. switch(config)# **feature npiv**
3. switch(config)# **no npiv enable**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **feature npiv**	Enables NPIV for all VSANs on the switch.
Step 3	switch(config)# **no npiv enable**	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.

This example shows how to create the port channel:

```
switch(config)# interface port-channel 2
switch(config-if)# switchport mode F
switch(config-if)# switchport  dedicated
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 in dedicated mode:

```
switch(config)# interface fc1/4-6
switch(config-if)# shut
switch(config-if)# switchport mode F
switch(config-if)# switchport speed 4000
switch(config-if)# switchport rate-mode dedicated
switch(config-if)# switchport trunk mode on
switch(config-if)# channel-group 2
switch(config-if)# no shut
switch(config-if)# exit
```

This example shows how to create the port channel in dedicated mode on the NPV switch:

```
switch(config)# interface san-port-channel 2
switch(config-if)# switchport mode NP
switch(config-if)# no shut
switch(config-if)# exit
```

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

```
switch(config)# interface fc2/1-2
switch(config-if)# shut
switch(config-if)# switchport mode NP
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 SPF 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-configuration** 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 configuration
...
interface fc3/5
  switchport speed 2000
...
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 fc3/5
interface fc3/5
  switchport speed 2000
  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 bbcredit
...
fc2/3 is trunking
    Transmit B2B Credit is 255
    Receive B2B Credit is 12
    Receive B2B Credit performance buffers is 375
      12 receive B2B credit remaining
      255 transmit B2B credit remaining
```

Default Fibre Channel Interface Settings

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

Table 7: 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 8: 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

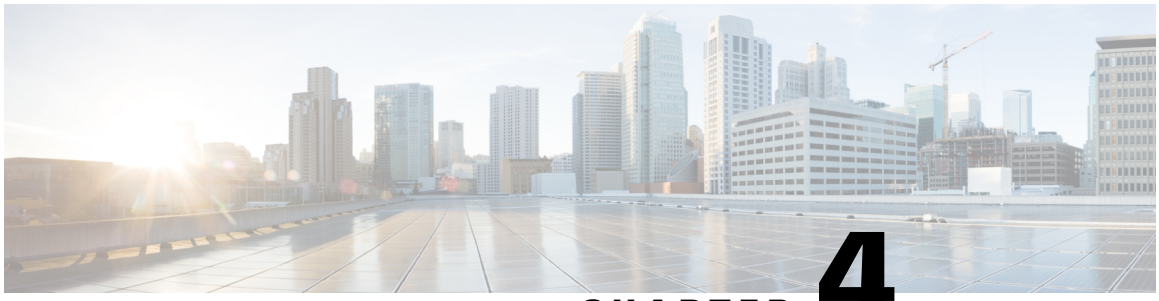

CHAPTER 4

Configuring Domain Parameters

This chapter contains the following sections:

Configuring 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.

Information About Fibre Channel Domains

This section describes each fcdomain phase:

- Principal switch selection—This phase guarantees the selection of a unique principal switch across the fabric.
- Domain ID distribution—This phase guarantees 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 illustrates an example fcdomain configuration.

Figure 2: Sample fcdomain Configuration

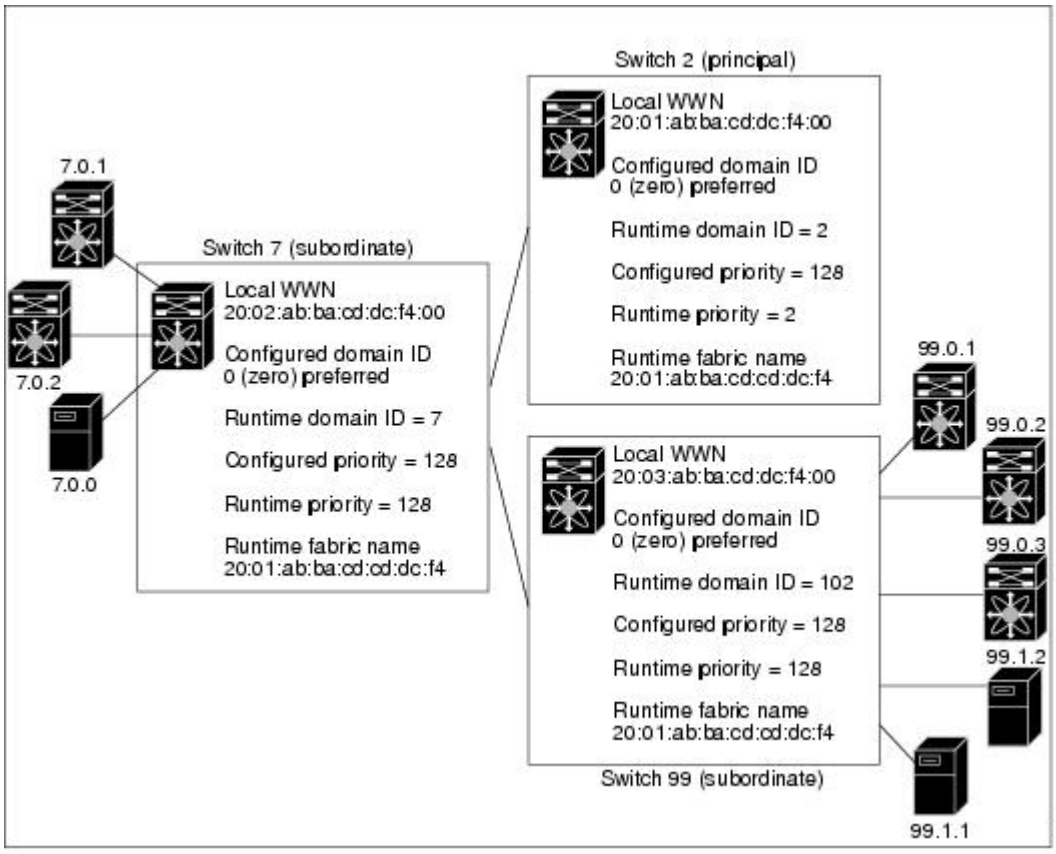

About Domain Restart

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, either disruptive or nondisruptive.

If a VSAN is in interop mode, you cannot disruptively restart the fcdomain for that VSAN.

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. Use the disruptive option to apply most of the configurations to their corresponding runtime values, including preferred domain IDs.

Restarting a Domain

To restart the fabric disruptively or nondisruptively, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcdomain restart vsan** *vsan-id*
3. switch(config)# **fcdomain restart disruptive vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fcdomain restart vsan** *vsan-id*	Forces the VSAN to reconfigure without traffic disruption.
Step 3	switch(config)# **fcdomain restart disruptive vsan** *vsan-id*	Forces the VSAN to reconfigure with data traffic disruption.

About Domain Manager Fast Restart

When a principal link fails, the domain manager must select a new principal link. By default, the domain manager starts a build fabric (BF) phase, followed by a principal switch selection phase. Both of these phases involve all the switches in the VSAN, and together take at least 15 seconds to complete. To reduce the time required for the domain manager to select a new principal link, you can enable the domain manager fast restart feature.

When fast restart is enabled and a backup link is available, the domain manager needs only a few milliseconds to select a new principal link to replace the one that failed. Also, the reconfiguration required to select the new principal link only affects the two switches that are directly attached to the failed link, not the entire VSAN. When a backup link is not available, the domain manager reverts to the default behavior and starts a BF phase, followed by a principal switch selection phase. The fast restart feature can be used in any interoperability mode.

Enabling Domain Manager Fast Restart

To enable the domain manager fast restart feature, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcdomain optimize fast-restart vsan** *vsan-id*
3. switch(config)# **fcdomain optimize fast-restart vsan** *vsan-id* **-** *vsan-id*
4. switch(config)# **no fcdomain optimize fast-restart vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fcdomain optimize fast-restart vsan** *vsan-id*	Enables domain manager fast restart in the specified VSAN.
Step 3	switch(config)# **fcdomain optimize fast-restart vsan** *vsan-id* **-** *vsan-id*	Enables domain manager fast restart in the specified range of VSANs.
Step 4	switch(config)# **no fcdomain optimize fast-restart vsan** *vsan-id*	Disables (default) domain manager fast restart in the specified VSAN.

About 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

To configure the priority for the principal switch, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcdomain priority** *number* **VSAN** *vsan-id*
3. switch(config)# **no fcdomain priority** *number* **VSAN** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fcdomain priority** *number* **VSAN** *vsan-id*	Configures the specified priority for the local switch in the specified VSAN.
Step 3	switch(config)# **no fcdomain priority** *number* **VSAN** *vsan-id*	Reverts the priority to the factory default (128) in the specified VSAN.

About fcdomain Initiation

By default, the fcdomain feature is enabled on each switch. If you disable the fcdomain feature in a switch, that switch can no longer participate with other switches in the fabric. The fcdomain configuration is applied to runtime through a disruptive restart.

Disabling or Reenabling fcdomains

To disable or reenable fcdomains in a single VSAN or a range of VSANs, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **no fcdomain vsan** *vsan-id* **-** *vsan-id*
3. switch(config)# **fcdomain vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **no fcdomain vsan** *vsan-id* **-** *vsan-id*	Disables the fcdomain configuration in the specified VSAN range.
Step 3	switch(config)# **fcdomain vsan** *vsan-id*	Enables the fcdomain configuration in the specified VSAN.

Configuring Fabric Names

To set the fabric name value for a disabled fcdomain, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcdomain fabric-name 20:1:ac:16:5e:0:21:01 vsan** *vsan-id*
3. switch(config)# **no fcdomain fabric-name 20:1:ac:16:5e:0:21:01 vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fcdomain fabric-name 20:1:ac:16:5e:0:21:01 vsan** *vsan-id*	Assigns the configured fabric name value in the specified VSAN.
Step 3	switch(config)# **no fcdomain fabric-name 20:1:ac:16:5e:0:21:01 vsan** *vsan-id*	Changes the fabric name value to the factory default (20:01:00:05:30:00:28:df) in VSAN 3010.

About 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 RFC reject option on virtual Fibre Channel interfaces, because these interfaces operate only in F port mode.

Rejecting Incoming RCFs

To reject incoming RCF request frames, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface fc** *slot/port*
3. switch(config-if)# **fcdomain rcf-reject vsan** *vsan-id*
4. switch(config-if)# **no fcdomain rcf-reject vsan** *vsan-id*

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.
Step 3	switch(config-if)# **fcdomain rcf-reject vsan** *vsan-id*	Enables the RCF filter on the specified interface in the specified VSAN.
Step 4	switch(config-if)# **no fcdomain rcf-reject vsan** *vsan-id*	Disables (default) the RCF filter on the specified interface in the specified VSAN.

About 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 may 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

To enable automatic reconfiguration in a specific VSAN (or range of VSANs), perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcdomain auto-reconfigure vsan** *vsan-id*
3. switch(config)# **no fcdomain auto-reconfigure vsan** *vsan-id*

DETAILED STEPS

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

	Command or Action	Purpose
Step 2	switch(config)# **fcdomain auto-reconfigure vsan** *vsan-id*	Enables the automatic reconfiguration option in the specified VSAN.
Step 3	switch(config)# **no fcdomain auto-reconfigure vsan** *vsan-id*	Disables the automatic reconfiguration option and reverts it to the factory default in the specified VSAN.

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.

About Domain IDs

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 3: Configuration Process Using the Preferred Option

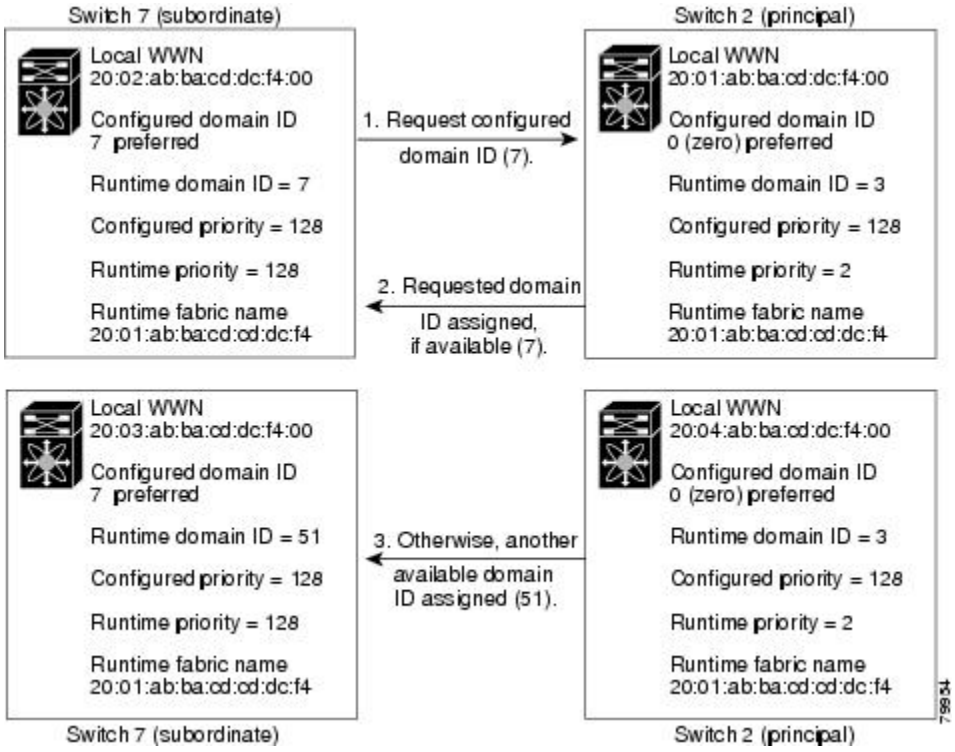

The operation of a subordinate switch changes based on three factors:

- 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

About Allowed Domain ID Lists, on page 37

Specifying Static or Preferred Domain IDs

When you assign a static domain ID type, you are requesting a particular domain ID. If the switch does not obtain the requested address, it will isolate itself from the fabric. When you specify a preferred domain ID, you are also requesting a particular domain ID; however, if the requested domain ID is unavailable, then the switch will accept another domain ID.

While the static option can be applied at runtime after a disruptive or nondisruptive restart, the preferred option is applied at runtime only after a disruptive restart.

Note Within a VSAN all switches should have the same domain ID type (either static or preferred). If a configuration is mixed (some switches with static domain types and others with preferred), you may experience link isolation.

To specify a static or preferred domain ID, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcdomain domain** *domain-id* **static vsan** *vsan-id*
3. switch(config)# **no fcdomain domain** *domain-id* **static vsan** *vsan-id*
4. switch(config)# **fcdomain domain** *domain-id* **preferred vsan** *vsan-id*
5. switch(config)# **no fcdomain domain** *domain-id* **preferred vsan** *vsan-id*

DETAILED STEPS

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

	Command or Action	Purpose
Step 2	switch(config)# **fcdomain domain** *domain-id* **static vsan** *vsan-id*	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.
Step 3	switch(config)# **no fcdomain domain** *domain-id* **static vsan** *vsan-id*	Resets the configured domain ID to factory defaults in the specified VSAN. The configured domain ID becomes 0 preferred.
Step 4	switch(config)# **fcdomain domain** *domain-id* **preferred vsan** *vsan-id*	Configures the switch in the specified VSAN to request a preferred domain ID 3 and accepts any value assigned by the principal switch. The domain is range is 1 to 239.
Step 5	switch(config)# **no fcdomain domain** *domain-id* **preferred vsan** *vsan-id*	Resets the configured domain ID to 0 (default) in the specified VSAN. The configured domain ID becomes 0 preferred.

About 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.

An allowed domain ID list must satisfy the following conditions:

- If this switch is a principal switch, all the currently assigned domain IDs must be in the allowed list.
- If this switch is a subordinate switch, the local runtime domain ID must be in the allowed list.
- The locally configured domain ID of the switch must be in the allowed list.
- The intersection of the assigned domain IDs with other already configured domain ID lists must not be empty.

Configuring Allowed Domain ID Lists

To configure the allowed domain ID list, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcdomain allowed** *domain-id range* **vsan** *vsan-id*
3. switch(config)# **no fcdomain allowed** *domain-id range* **vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fcdomain allowed** *domain-id range* **vsan** *vsan-id*	Configures the list to allow switches with the domain ID range in the specified VSAN.
Step 3	switch(config)# **no fcdomain allowed** *domain-id range* **vsan** *vsan-id*	Reverts to the factory default of allowing domain IDs from 1 through 239 in the specified VSAN.

About 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 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.

For additional information, refer to Using Cisco Fabric Services in the *Cisco Nexus 5000 Series System Management Configuration Guide*.

Enabling 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.

To enable (or disable) allowed domain ID list configuration distribution, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcdomain distribute**
3. switch(config)# **no fcdomain distribute**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fcdomain distribute**	Enables domain configuration distribution.
Step 3	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

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.

To commit pending domain configuration changes and release the lock, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcdomain commit vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fcdomain commit vsan** *vsan-id*	Commits the pending domain configuration changes.

Discarding Changes

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.

To discard pending domain configuration changes and release the lock, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcdomain abort vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fcdomain abort vsan** *vsan-id*	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 using the **show fcdomain status** command.

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

Displaying Pending Changes

You can display the pending configuration changes 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 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 using the **show fcdomain session-status vsan** command.

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

About 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

To enable contiguous domains in a specific VSAN (or a range of VSANs), perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcdomain contiguous-allocation vsan** *vsan-id* **-** *vsan-id*
3. switch(config)# **no fcdomain contiguous-allocation vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fcdomain contiguous-allocation vsan** *vsan-id* **-** *vsan-id*	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	switch(config)# **no fcdomain contiguous-allocation vsan** *vsan-id*	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 Cisco Nexus 5000 Series 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 Cisco Nexus 5000 Series 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).

About 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

To enable the persistent FC ID feature, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcdomain fcid persistent vsan** *vsan-id*
3. switch(config)# **no fcdomain fcid persistent vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fcdomain fcid persistent vsan** *vsan-id*	Activates (default) persistency of FC IDs in the specified VSAN.
Step 3	switch(config)# **no fcdomain fcid persistent vsan** *vsan-id*	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

To configure persistent FC IDs, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcdomain fcid database**
3. switch(config-fcid-db)# **vsan** *vsan-id* **wwn 33:e8:00:05:30:00:16:df fcid** *fcid*
4. switch(config-fcid-db)# **vsan** *vsan-id* **wwn 11:22:11:22:33:44:33:44 fcid** *fcid* **dynamic**
5. switch(config-fcid-db)# **vsan** *vsan-id* **wwn 11:22:11:22:33:44:33:44 fcid** *fcid* **area**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fcdomain fcid database**	Enters FC ID database configuration submode.
Step 3	switch(config-fcid-db)# **vsan** *vsan-id* **wwn 33:e8:00:05:30:00:16:df fcid** *fcid*	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	switch(config-fcid-db)# **vsan** *vsan-id* **wwn 11:22:11:22:33:44:33:44 fcid** *fcid* **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	switch(config-fcid-db)# **vsan** *vsan-id* **wwn 11:22:11:22:33:44:33:44 fcid** *fcid* **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.

About Unique Area FC IDs for HBAs

Note Only read this section 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 Nexus 5000 Series 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

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 vfc20/1 and the storage port connects to interface fc2/3 on the same switch.

To configure a different area ID for the HBA port, perform this task:

SUMMARY STEPS

1. Obtain the port WWN (Port Name field) ID of the HBA using the **show flogi database** command.
2. Shut down the HBA interface in the Cisco Nexus 5000 Series switch.
3. Verify that the FC ID feature is enabled using the **show fcdomain vsan** command.
4. Enable the persistent FC ID feature in the Cisco Nexus 5000 Series switch.
5. Assign a new FC ID with a different area allocation. In this example, replace *77* with *ee*.
6. Enable the HBA interface in the Cisco Nexus 5000 Series switch.
7. Verify the pWWN ID of the HBA by using the **show flogi database** command.

DETAILED STEPS

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
--------------------------------------------------------------------------------------------
 vfc10/1    3       0x6f7703    50:05:08:b2:00:71:c8:c2     50:05:08:b2:00:71:c8:c0
 fc2/3      3       0x6f7704    50:06:0e:80:03:29:61:0f     50:06:0e:80:03:29:61:0f
```

Note Both FC IDs in this setup have the same area 77 assignment.

Step 2 Shut down the HBA interface in the Cisco Nexus 5000 Series switch.

```
switch# configuration terminal
switch(config)# interface vfc20/1

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 1
...
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 Cisco Nexus 5000 Series switch.

```
switch# configuration terminal
switch(config)# fcdomain fcid persistent vsan 1
switch(config)# end
```

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

```
switch# configuration 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 Cisco Nexus 5000 Series switch.

```
switch# configuration terminal
switch(config)# interface vfc20/1
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
--------------------------------------------------------------------------------
 vfc20/1    3       0x6fee00    50:05:08:b2:00:71:c8:c2   50:05:08:b2:00:71:c8:c0
 fc2/3      3       0x6f7704    50:06:0e:80:03:29:61:0f   50:06:0e:80:03:29:61:0f
```

Note Both FC IDs now have different area assignments.

About 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 9: 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

To purge persistent FC IDs, perform this task:

SUMMARY STEPS

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

DETAILED STEPS

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

Verifying fcdomain Information

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 Fibre Channel Domain Settings

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

Table 10: 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
Persistent FC ID	Enabled
Allowed domain ID list configuration distribution	Disabled

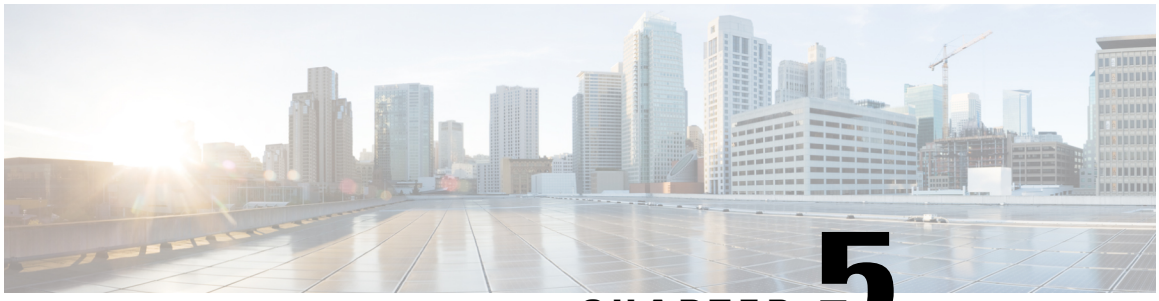

CHAPTER 5

Configuring N Port Virtualization

This chapter contains the following sections:

Configuring N Port Virtualization

Information About NPV

NPV Overview

By default, Cisco Nexus 5000 Series switches operate in fabric mode. In this mode, the switch provides standard Fibre Channel switching capability and features.

In fabric mode, each switch that joins a SAN is assigned a domain ID. Each SAN (or VSAN) supports a maximum of 239 domain IDs, so the SAN has a limit of 239 switches. In a SAN topology with a large number of edge switches, the SAN may need to grow beyond this limit. NPV alleviates the domain ID limit by sharing the domain ID of the core switch among multiple edge switches.

In NPV mode, the edge switch relays all traffic from server-side ports to the core switch. The core switch provides F port functionality (such as login and port security) and all the Fibre Channel switching capabilities.

The edge switch appears as a Fibre Channel host to the core switch and as a regular Fibre Channel switch to its connected devices.

The figure below shows an interface-level view of an NPV configuration.

Figure 4: NPV Interface Configuration

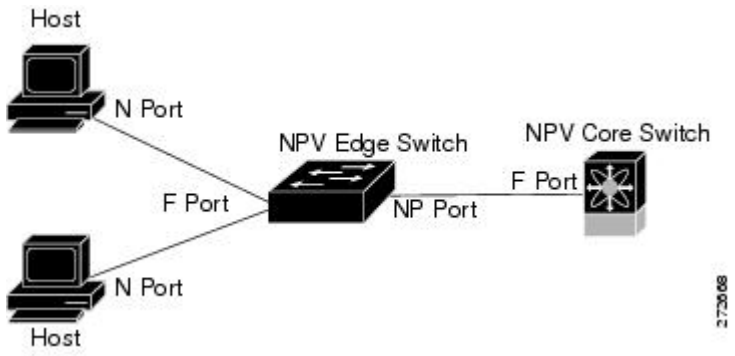

NPV Mode

In NPV mode, the edge switch relays all traffic to the core switch, which provides the Fibre Channel switching capabilities. The edge switch shares the domain ID of the core switch.

To convert a switch into NPV mode, you set the NPV feature to enabled. This configuration command automatically triggers a switch reboot. You cannot configure NPV mode on a per-interface basis. NPV mode applies to the entire switch.

In NPV mode, a subset of fabric mode CLI commands and functionality is supported. For example, commands related to fabric login and name server registration are not required on the edge switch, because these functions are provided in the core switch. To display the fabric login and name server registration databases, you must enter the **show flogi database** and **show fcns database** commands on the core switch.

Server Interfaces

Server interfaces are F ports on the edge switch that connect to the servers. A server interface may support multiple end devices by enabling the N port identifier virtualization (NPIV) feature. NPIV provides a means to assign multiple FC IDs to a single N port, which allows the server to assign unique FC IDs to different applications.

Note To use NPIV, enable the NPIV feature and reinitialize the server interfaces that will support multiple devices.

Server interfaces are automatically distributed among the NP uplinks to the core switch. All of the end devices connected to a server interface are mapped to the same NP uplink.

In Cisco Nexus 5000 Series switches, server interfaces can be physical or virtual Fibre Channel interfaces.

NP Uplinks

All interfaces from the edge switch to the core switch are configured as proxy N ports (NP ports).

An NP uplink is a connection from an NP port on the edge switch to an F port on the core switch. When an NP uplink is established, the edge switch sends a fabric login message (FLOGI) to the core switch, and then (if the FLOGI is successful) it registers itself with the name server on the core switch. Subsequent FLOGIs from end devices connected to this NP uplink are converted to fabric discovery messages (FDISCs).

Note In the switch CLI configuration commands and output displays, NP uplinks are called External Interfaces.

In Cisco Nexus 5000 Series switches, NP uplink interfaces must be native Fibre Channel interfaces.

FLOGI Operation

When an NP port becomes operational, the switch first logs itself in to the core switch by sending a FLOGI request (using the port WWN of the NP port).

After completing the FLOGI request, the switch registers itself with the fabric name server on the core switch (using the symbolic port name of the NP port and the IP address of the edge switch).

The following table identifies port and node names in the edge switch used in NPV mode.

Table 11: Edge Switch FLOGI Parameters

Parameter	Derived From
pWWN	The fWWN of the NP port on the edge switch.
nWWN	The VSAN-based sWWN of the edge switch.
symbolic port name	The edge switch name and NP port interface string. **Note** If no switch name is available, the output will read "switch." For example, switch: fc 1/5.
IP address	The IP address of the edge switch.
symbolic node name	The edge switch name.

Note The buffer-to-buffer state change number (BB_SCN) of internal FLOGIs on an NP port is always set to zero. The BB_SCN is supported by the F port on the edge switch.

We do not recommend using fWWN-based zoning on the edge switch for the following reasons:

- Zoning is not enforced at the edge switch (rather, it is enforced on the core switch).
- Multiple devices attached to an edge switch log in through the same F port on the core, so they cannot be separated into different zones.
- The same device might log in using different fWWNs on the core switch (depending on the NPV link it uses) and may need to be zoned using different fWWNs.

NPV Traffic Management

Automatic Uplink Selection

NPV supports automatic selection of NP uplinks. When a server interface is brought up, the NP uplink interface with the minimum load is selected from the available NP uplinks in the same VSAN as the server interface.

When a new NP uplink interface becomes operational, the existing load is not redistributed automatically to include the newly available uplink. Server interfaces that become operational after the NP uplink can select the new NP uplink.

Traffic Maps

In Release 4.0(1a)N2(1) and later software releases, NPV supports traffic maps. A traffic map allows you to specify the NP uplinks that a server interface can use to connect to the core switches.

Note When an NPV traffic map is configured for a server interface, the server interface must select only from the NP uplinks in its traffic map. If none of the specified NP uplinks are operational, the server remains in a non-operational state.

The NPV traffic map feature provides the following benefits:

- Facilitates traffic engineering by allowing configuration of a fixed set of NP uplinks for a specific server interface (or range of server interfaces).
- Ensures correct operation of the persistent FC ID feature, because a server interface will always connect to the same NP uplink (or one of a specified set of NP uplinks) after an interface reinitialization or switch reboot.

Disruptive Load Balancing

In Release 4.0(0)N1(2a) and later software releases, NPV supports disruptive load balancing. When disruptive load balancing is enabled, NPV redistributes the server interfaces across all available NP uplinks when a new NP uplink becomes operational. To move a server interface from one NP uplink to another NP uplink, NPV forces reinitialization of the server interface so that the server performs a new login to the core switch.

Only server interfaces that are moved to a different uplink are reinitialized. A system message is generated for each server interface that is moved.

Note Redistributing a server interface causes traffic disruption to the attached end devices.

To avoid disruption of server traffic, you should enable this feature only after adding a new NP uplink, and then disable it again after the server interfaces have been redistributed.

If disruptive load balancing is not enabled, you can manually reinitialize some or all of the server interfaces to distribute server traffic to new NP uplink interfaces.

NPV Traffic Management Guidelines

When deploying NPV traffic management, follow these guidelines:

- Use NPV traffic management only when automatic traffic engineering does not meet your network requirements.
- You do not need to configure traffic maps for all server interfaces. By default, NPV will use automatic traffic management.
- Server interfaces configured to use a set of NP uplink interfaces cannot use any other available NP uplink interfaces, even if none of the configured interfaces are available.
- When disruptive load balancing is enabled, a server interface may be moved from one NP uplink to another NP uplink. Moving between NP uplink interfaces requires NPV to relogin to the core switch, causing traffic disruption.
- To link a set of servers to a specific core switch, associate the server interfaces with a set of NP uplink interfaces that all connect to that core switch.
- Configure Persistent FC IDs on the core switch and use the Traffic Map feature to direct server interface traffic onto NP uplinks that all connect to the associated core switch.

NPV Guidelines and Limitations

When configuring NPV, note the following guidelines and limitations:

- In-order data delivery is not required in NPV mode because the exchange between two end devices always takes the same uplink from the edge switch to the core. Upstream of the edge switch, core switches will enforce in-order delivery if configured.
- You can configure zoning for end devices that are connected to edge switches using all available member types on the core switch. For fWWN, sWWN, domain, or port-based zoning, use the fWWN, sWWN, domain, or port of the core switch in the configuration commands.
- Port tracking is not supported in NPV mode.
- Port security is supported on the core switch for devices logged in through the NPV switch. Port security is enabled on the core switch on a per-interface basis. To enable port security on the core switch for devices that log in through an NPV switch, you must adhere to the following requirements:
  - The internal FLOGI must be in the port security database; in this way, the port on the core switch will allow communications and links.
  - All the end device pWWNs must also be in the port security database.
- Edge switches can connect to multiple core switches. In other words, different NP ports can be connected to different core switches.
- NPV uses a load-balancing algorithm to automatically assign end devices in a VSAN to one of the NP uplinks (in the same VSAN) upon initial login. If there are multiple NP uplinks in the same VSAN, you cannot assign an end device to a specific NP uplink.
- If a server interface goes down and then returns to service, the interface is not guaranteed to be assigned to the same NP uplink.

- The server interface is only operational when its assigned NP uplink is operational.
- Servers can be connected to the switch when in NPV mode.
- Targets can not be connected to the switch when in NPV mode.
- Fibre Channel switching is not performed in the edge switch; all traffic is switched in the core switch.
- NPV supports NPIV-capable module servers. This capability is called nested NPIV.
- Only F, NP, and SD ports are supported in NPV mode.

Configuring NPV

Enabling NPV

When you enable NPV, the system configuration is erased and the switch reboots.

Note We recommend that you save your current configuration either in boot flash memory or to a TFTP server before you enable NPV.

To enable NPV, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **npv enable**
3. switch(config-npv)# **no npv enable**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters configuration mode.
Step 2	switch(config)# **npv enable**	Enables NPV mode. The switch reboots, and it comes back up in NPV mode. **Note** A write-erase is performed during the initialization.
Step 3	switch(config-npv)# **no npv enable**	Disables NPV mode, which results in a reload of the switch.

Configuring NPV Interfaces

After you enable NPV, you should configure the NP uplink interfaces and the server interfaces.

Configuring an NP Interface

After you enable NPV, you should configure the NP uplink interfaces and the server interfaces. To configure an NP uplink interface, perform this task:

To configure a server interface, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface fc** *slot*/*port*
3. switch(config-if)# **switchport mode NP**
4. switch(config-if)# **no shutdown**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters configuration mode.
Step 2	switch(config)# **interface fc** *slot*/*port*	Selects an interface that will be connected to the core NPV switch.
Step 3	switch(config-if)# **switchport mode NP**	Configures the interface as an NP port.
Step 4	switch(config-if)# **no shutdown**	Brings up the interface.

Configuring a Server Interface

To configure a server interface, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **interface** {**fc** *slot*/*port* | **vfc** *vfc-id*}
3. switch(config-if)# **switchport mode F**
4. switch(config-if)# **no shutdown**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters configuration mode.
Step 2	switch(config)# **interface** {**fc** *slot*/*port*	**vfc** *vfc-id*}	Selects a server interface.

	Command or Action	Purpose
Step 3	switch(config-if)# **switchport mode F**	Configures the interface as an F port.
Step 4	switch(config-if)# **no shutdown**	Brings up the interface.

Configuring NPV Traffic Management

Configuring NPV Traffic Maps

An NPV traffic map associates one or more NP uplink interfaces with a server interface. The switch associates the server interface with one of these NP uplinks.

Note If a server interface is already mapped to an NP uplink, you should include this mapping in the traffic map configuration.

To configure a traffic map, perform this task:

SUMMARY STEPS

1. switch# **config t**
2. switch(config)# **npv traffic-map server-interface** {**fc** *slot/port* | **vfc** *vfc-id*} e**xternal-interface fc** *slot/port*
3. switch(config)# **no npv traffic-map server-interface** {**fc** *slot/port* | **vfc** *vfc-id*} **external-interface fc** *slot/port*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **config t**	Enters configuration mode on the NPV.
Step 2	switch(config)# **npv traffic-map server-interface** {**fc** *slot/port* | **vfc** *vfc-id*} e**xternal-interface fc** *slot/port*	Configures a mapping between a server interface (or range of server interfaces) and an NP uplink interface (or range of NP uplink interfaces).
Step 3	switch(config)# **no npv traffic-map server-interface** {**fc** *slot/port* | **vfc** *vfc-id*} **external-interface fc** *slot/port*	Removes the mapping between the specified server interfaces and NP uplink interfaces.

Enabling Disruptive Load Balancing

If you configure additional NP uplinks, you can enable the disruptive load-balancing feature to distribute the server traffic load evenly among all the NP uplinks.

To enable disruptive load balancing, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **npv auto-load-balance disruptive**
3. switch (config)# **no npv auto-load-balance disruptive**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters configuration mode on the NPV.
Step 2	switch(config)# **npv auto-load-balance disruptive**	Enables disruptive load balancing on the switch.
Step 3	switch (config)# **no npv auto-load-balance disruptive**	Disables disruptive load balancing on the switch.

Verifying NPV

To display information about NPV, perform the following task:

SUMMARY STEPS

1. switch# **show npv flogi-table** [**all**]

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **show npv flogi-table** [**all**]	Displays the NPV configuration.

Verifying NPV Examples

To display a list of devices on a server interface and their assigned NP uplinks, enter the **show npv flogi-table** command on the Cisco Nexus 5000 Series switch:

```
switch# show npv flogi-table
--------------------------------------------------------------------------------
SERVER                                                                  EXTERNAL
INTERFACE VSAN FCID            PORT NAME               NODE NAME       INTERFACE
--------------------------------------------------------------------------------
vfc3/1    1    0xee0008 10:00:00:00:c9:60:e4:9a 20:00:00:00:c9:60:e4:9a fc2/1

vfc3/1    1    0xee0009 20:00:00:00:0a:00:00:01 20:00:00:00:c9:60:e4:9a fc2/2
vfc3/1    1    0xee000a 20:00:00:00:0a:00:00:02 20:00:00:00:c9:60:e4:9a fc2/3
```

```
vfc3/1    1    0xee000b 33:33:33:33:33:33:33:33 20:00:00:00:c9:60:e4:9a fc2/4
Total number of flogi = 4
```

Note For each server interface, the External Interface value displays the assigned NP uplink.

To display the status of the server interfaces and the NP uplink interfaces, enter the **show npv status** command:

```
switch# show npv status
npiv is enabled

External Interfaces:
====================
  Interface: fc2/1, VSAN: 1, FCID: 0x1c0000, State: Up
  Interface: fc2/2, VSAN: 1, FCID: 0x040000, State: Up
  Interface: fc2/3, VSAN: 1, FCID: 0x260000, State: Up
  Interface: fc2/4, VSAN: 1, FCID: 0x1a0000, State: Up
  Number of External Interfaces: 4

Server Interfaces:
==================
  Interface: vfc3/1, VSAN: 1, NPIV: No, State: Up
  Number of Server Interfaces: 1
```

Note To view fcns database entries for NPV edge switches, you must enter the **show fcns database** command on the core switch.

To view all the NPV edge switches, enter the **show fcns database** command on the core switch:

```
core-switch# show fcns database
```

For additional details (such as IP addresses, switch names, interface names) about the NPV edge switches that you see in the **show fcns database** output, enter the **show fcns database detail** command on the core switch:

```
core-switch# show fcns database detail
```

Verifying NPV Traffic Management

To display the NPV traffic map, enter the **show npv traffic-map** command.

```
switch# show npv traffic-map
NPV Traffic Map Information:
----------------------------------------
Server-If       External-If(s)
----------------------------------------
fc1/3           fc1/10,fc1/11
fc1/5           fc1/1,fc1/2
----------------------------------------
```

To display the NPV internal traffic details, enter the **show npv internal info traffic-map** command.

To display the disruptive load-balancing status, enter the **show npv status** command:

```
switch# show npv status
npiv is enabled
disruptive load balancing is enabled
External Interfaces:
====================
  Interface: fc2/1, VSAN: 2, FCID: 0x1c0000, State: Up
...
```

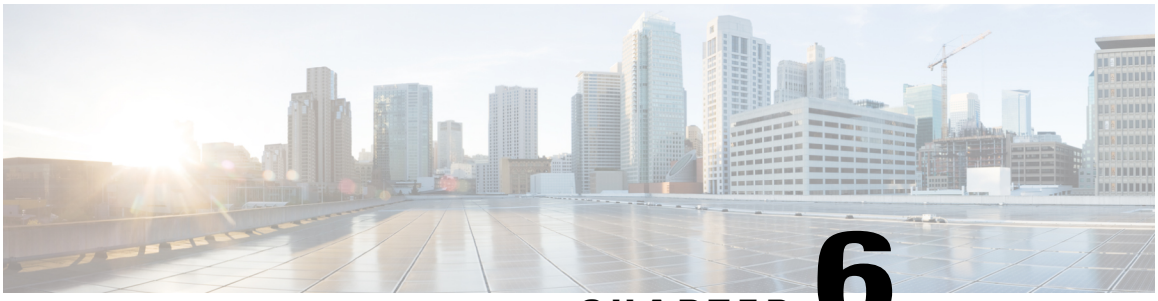

CHAPTER 6

Configuring VSAN Trunking

This chapter contains the following sections:

Configuring VSAN Trunking

Information About VSAN Trunking

VSAN trunking enables interconnected ports to transmit and receive frames in more than one VSAN. Trunking is supported on E ports and F ports.

Beginning in Cisco NX-OS Release 5.0(2)N1(1), VSAN trunking is supported on native Fibre Channel interfaces and virtual Fibre Channel interfaces.

The VSAN trunking feature includes the following restrictions:

- Trunking configurations are only applicable to E ports. If trunk mode is enabled in an E port and that port becomes operational as a trunking E port, it is referred to as a TE port.
- The trunk-allowed VSANs configured for TE ports are used by the trunking protocol to determine the allowed-active VSANs in which frames can be received or transmitted.
- If a trunking-enabled E port is connected to a third-party switch, the trunking protocol ensures seamless operation as an E port.

Trunking E Ports

Trunking E ports enables interconnected ports to transmit and receive frames in more than one VSAN, over the same physical link, using enhanced ISL (EISL) frame format.

Figure 5: Trunking E Ports

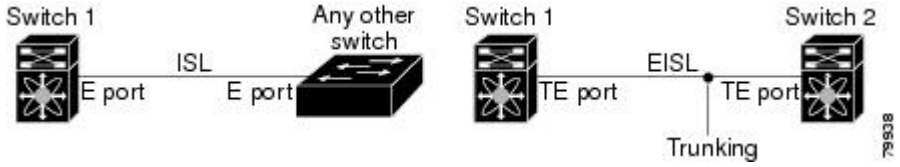

Trunking F Ports

Trunking F ports allows interconnected ports to transmit and receive tagged frames in more than one VSAN, over the same physical link.

VSAN Trunking Mismatches

If you misconfigure VSAN configurations across E ports, issues can occur such as the merging of traffic in two VSANs (causing both VSANs to mismatch). The VSAN trunking protocol validates the VSAN interfaces at both ends of an ISL to avoid merging VSANs (see the following figure).

Figure 6: VSAN Mismatch

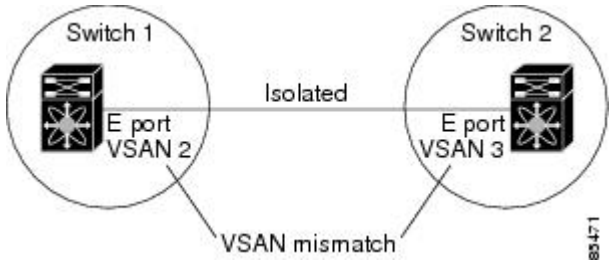

In this example, the trunking protocol detects potential VSAN merging and isolates the ports involved.

The trunking protocol cannot detect merging of VSANs when a third-party switch is placed in between two Cisco Nexus 5000 Series switches (see the following figure).

Figure 7: Third-Party Switch VSAN Mismatch

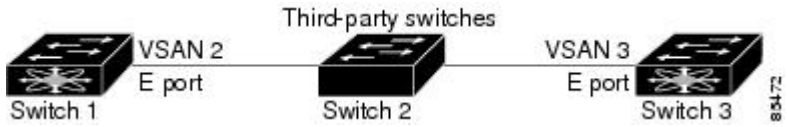

VSAN 2 and VSAN 3 are effectively merged with overlapping entries in the name server and the zone applications. The Cisco MDS 9000 Fabric Manager helps detect such topologies.

VSAN Trunking Protocol

The trunking protocol is important for E-port and TE-port operations. It supports the following capabilities:

- Dynamic negotiation of operational trunk mode.
- Selection of a common set of trunk-allowed VSANs.
- Detection of a VSAN mismatch across an ISL.

By default, the VSAN trunking protocol is enabled. If the trunking protocol is disabled on a switch, no port on that switch can apply new trunk configurations. Existing trunk configurations are not affected: the TE port continues to function in trunk mode, but only supports traffic in VSANs that it negotiated with previously (when the trunking protocol was enabled). Other switches that are directly connected to this switch are similarly affected on the connected interfaces. If you need to merge traffic from different port VSANs across a nontrunking ISL, disable the trunking protocol.

Configuring VSAN Trunking

Guidelines and Restrictions

When configuring VSAN trunking, note the following guidelines:

- We recommend that both ends of a VSAN trunking ISL belong to the same port VSAN. On platforms or fabric switches where the port VSANs are different, one end returns an error, and the other is not connected.
- To avoid inconsistent configurations, disable all E ports with a **shutdown** command before enabling or disabling the VSAN trunking protocol.

Difference Between TE Ports and TF-TNP Ports

In case of TE ports, the VSAN will in be initializing state when VSAN is coming up on that interface and when peers are in negotiating phase. Once the handshake is done, VSAN will be moved to up state in the successful case, and isolated state in the case of failure. Device Manager will show the port status as amber during initializing state and it will be green once VSANs are up.

In case of TF ports, after the handshake, one of the allowed VSAN will be moved to up state. And all other VSAN will be in initializing state even though the handshake with the peer is completed and successful. Each VSAN will be moved from initializing state to up state when a server or target logs in through the trunked F or NP ports in the corresponding VSAN.

Enabling or Disabling the VSAN Trunking Protocol

To enable or disable the VSAN trunking protocol, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **no trunk protocol enable**
3. switch(config)# **trunk protocol enable**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **no trunk protocol enable**	Disables the trunking protocol.
Step 3	switch(config)# **trunk protocol enable**	Enables trunking protocol (default).

About Trunk Mode

By default, trunk mode is enabled in all Fibre Channel interfaces. However, trunk mode configuration takes effect only in E-port mode. You can configure trunk mode as on (enabled), off (disabled), or auto (automatic). The default trunk mode is on. The trunk mode configurations at the two ends of the link determine the trunking state of the link and the port modes at both ends (see the following table).

Table 12: Trunk Mode Status Between Switches

Your Trunk Mode Configuration	Resulting State and Port Mode		
Switch 1	**Switch 2**	**Trunking State**	**Port Mode**
On	Auto or on	Trunking (EISL)	TE port
Off	Auto, on, or off	No trunking (ISL)	E port
Auto	Auto	No trunking (ISL)	E port

The preferred configuration on the Cisco Nexus 5000 Series switches is that one side of the trunk is set to auto and the other is set to on.

Note When connected to a third-party switch, the trunk mode configuration has no effect. The ISL is always in a trunking disabled state.

Configuring Trunk Mode

To configure trunk mode, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface** {**fc** *slot/port* | **vfc** *vfc-id*}
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* | **vfc** *vfc-id*}	Configures the specified Fibre Channel or virtual Fibre Channel interface.
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. **Note** Trunk mode cannot be turned off for virtual Fibre Channel interfaces.
Step 5	switch(config-if)# **switchport trunk mode auto**	Configures the trunk mode to **auto** mode, which provides automatic sensing for the interface.

The following example shows how to configure a vFC interface in trunk mode.

```
switch# config t
switch#(config)# vfc 200
switch(config-if)# switchport trunk mode on
```

The following example shows the output for the vFC interface 200 in trunk mode.

```
switch(config-if)# show interface vfc200
vfc200 is trunking (Not all VSANs UP on the trunk)
    Bound interface is Ethernet1/3
    Hardware is Virtual Fibre Channel
    Port WWN is 20:c7:00:0d:ec:f2:08:ff
    Peer port WWN is 00:00:00:00:00:00:00:00
    Admin port mode is E, trunk mode is on
    snmp link state traps are enabled
    Port mode is TE
    Port vsan is 1
    Trunk vsans (admin allowed and active) (1-6,10,22)
    Trunk vsans (up)                       ()
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             (1-6,10,22)
    5 minute input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minute output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
```

```
  0 frames input, 0 bytes
    0 discards, 0 errors
  0 frames output, 0 bytes
    0 discards, 0 errors
last clearing of "show interface" counters never
Interface last changed at Mon Jan 18 10:01:27 2010
```

About Trunk-Allowed VSAN Lists

Each Fibre Channel interface has an associated trunk-allowed VSAN list. In TE-port mode, frames are transmitted and received in one or more VSANs specified in this list. By default, the complete VSAN range (1 through 4093) is included in the trunk-allowed list.

The common set of VSANs that are configured and active in the switch are included in the trunk-allowed VSAN list for an interface, and they are called *allowed-active VSANs*. The trunking protocol uses the list of allowed-active VSANs at the two ends of an ISL to determine the list of operational VSANs in which traffic is allowed.

In the following figure, switch 1 has VSANs 1 through 5, switch 2 has VSANs 1 through 3, and switch 3 has VSANs 1, 2, 4, and 5 with a default configuration of trunk-allowed VSANs. All VSANs configured in all three switches are allowed-active. However, only the common set of allowed-active VSANs at the ends of the ISL become operational as shown in below.

Figure 8: Default Allowed-Active VSAN Configuration

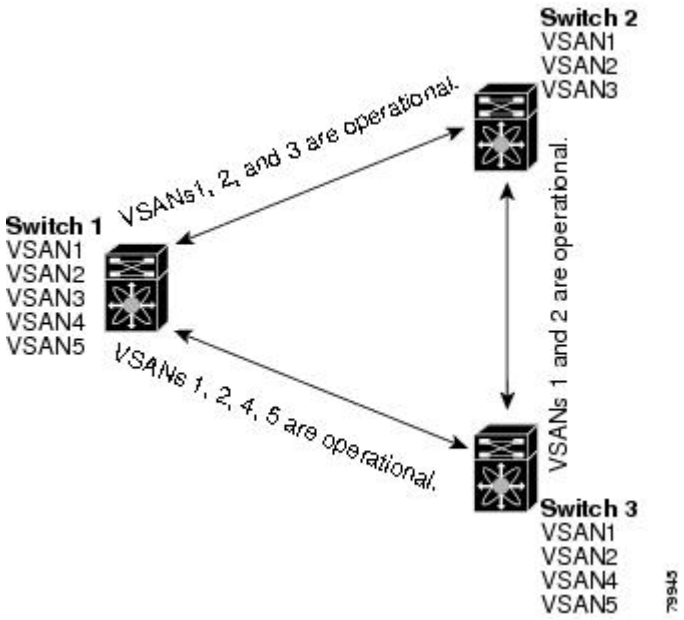

You can configure a selected set of VSANs (from the allowed-active list) to control access to the VSANs specified in a trunking ISL.

Using the figure above as an example, you can configure the list of allowed VSANs on a per-interface basis (see the following figure). For example, if VSANs 2 and 4 are removed from the allowed VSAN list of ISLs connecting to switch 1, the operational allowed list of VSANs for each ISL would be as follows:

- The ISL between switch 1 and switch 2 includes VSAN 1 and VSAN 3.

- The ISL between switch 2 and switch 3 includes VSAN 1 and VSAN 2.
- The ISL between switch 3 and switch 1 includes VSAN 1, 2, and 5.

Consequently, VSAN 2 can only be routed from switch 1 through switch 3 to switch 2.

Figure 9: Operational and Allowed VSAN Configuration

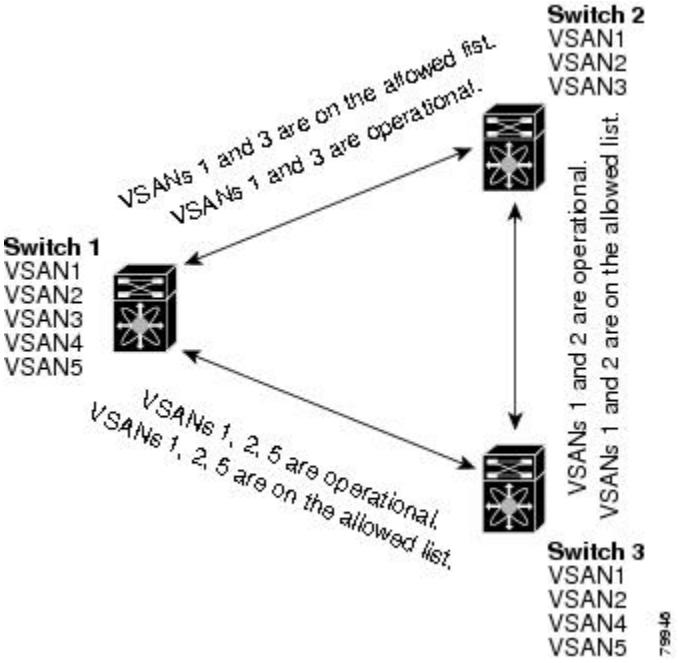

Configuring an Allowed-Active List of VSANs

To configure an allowed-active list of VSANs for an interface, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface fc** *slot/port*
3. switch(config-if)# **switchport trunk allowed vsan** *vsan-id* **-** *vsan-id*
4. switch(config-if)# **switchport trunk allowed vsan add** *vsan-id*
5. switch(config-if)# **no switchport trunk allowed vsan** *vsan-id* **-** *vsan-id*
6. switch(config-if)# **no switchport trunk allowed vsan add** *vsan-id*

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*	Configures the specified interface.
Step 3	switch(config-if)# **switchport trunk allowed vsan** *vsan-id* **-** *vsan-id*	Changes the allowed list for the specified VSAN range.
Step 4	switch(config-if)# **switchport trunk allowed vsan add** *vsan-id*	Expands the specified VSAN to the new allowed list.
Step 5	switch(config-if)# **no switchport trunk allowed vsan** *vsan-id* **-** *vsan-id*	Deletes the specified VSAN range.
Step 6	switch(config-if)# **no switchport trunk allowed vsan add** *vsan-id*	Deletes the expanded allowed list.

Displaying VSAN Trunking Information

The **show interface** command is invoked from the EXEC mode and displays VSAN trunking configurations for a TE port. Without any arguments, this command displays the information for all of the configured interfaces in the switch.

The following example shows how to display the trunk mode of a Fibre Channel interface:

```
switch# show interface fc3/3
fc3/3 is up
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:83:00:0d:ec:6d:78:40
    Peer port WWN is 20:0c:00:0d:ec:0d:d0:00
    Admin port mode is auto, trunk mode is on
...
```

The following example shows how to display the trunk protocol of a Fibre Channel interface:

```
switch# show trunk protocol
Trunk protocol is enabled
```

The following example shows how to display the VSAN information for all trunk interfaces:

```
switch# show interface trunk vsan 1-1000
fc3/1 is not trunking
...
fc3/11 is trunking
    Belongs to san-port-channel 6
    Vsan 1 is up, FCID is 0xef0000
    Vsan 2 is up, FCID is 0xef0000
...
san-port-channel 6 is trunking
    Vsan 1 is up, FCID is 0xef0000
    Vsan 2 is up, FCID is 0xef0000
```

Default Trunk Configuration Settings

The following table lists the default settings for trunking parameters.

Table 13: Default Trunk Configuration Parameters

Parameters	Default
Switch port trunk mode	On
Allowed VSAN list	1 to 4093 user-defined VSAN IDs
Trunking protocol	Enabled

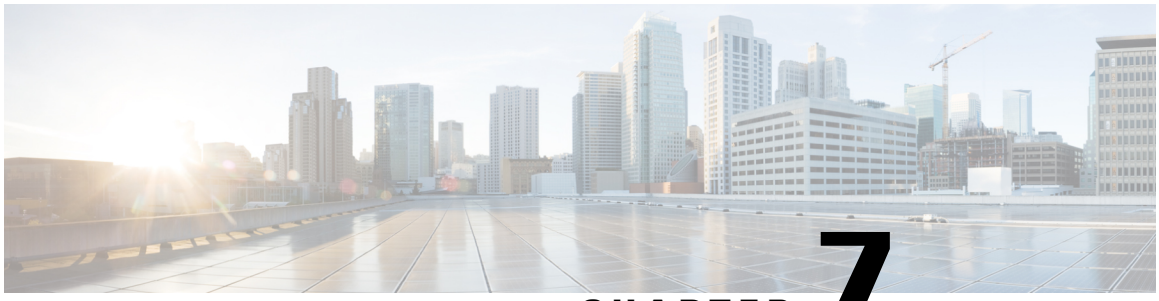

CHAPTER 7

Configuring SAN Port Channel

This chapter contains the following sections:

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 5000 Series switches, SAN port channels can include physical Fibre Channel interfaces, but not virtual Fibre Channel interfaces. A SAN port channel can include up to eight Fibre Channel interfaces.

Information About SAN Port Channels

About E 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.

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.

Cisco Nexus 5000 Series switches support a maximum of four SAN port channels (with eight interfaces per port channel). A port channel number refers to the unique (within each switch) identifier associated with each channel group. This number ranges from 1 to 256.

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 MDS 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).

Understanding Port Channels and VSAN Trunking

Switches in the Cisco Nexus 5000 Series 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 10: VSAN Trunking Only

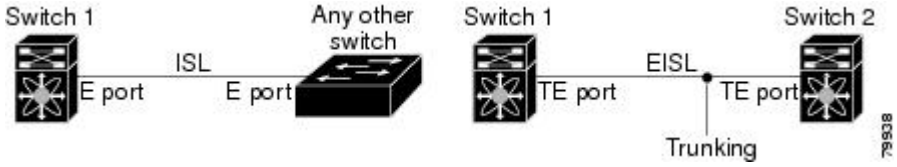

- You can create a SAN port channel with members that are E ports, as shown on the left side of the figure below. In this configuration, the port channel implements a logical ISL (carrying traffic for one VSAN).
- You can create a SAN port channel with members that are TE-ports, as shown on the right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

Figure 11: Port Channels and VSAN Trunking

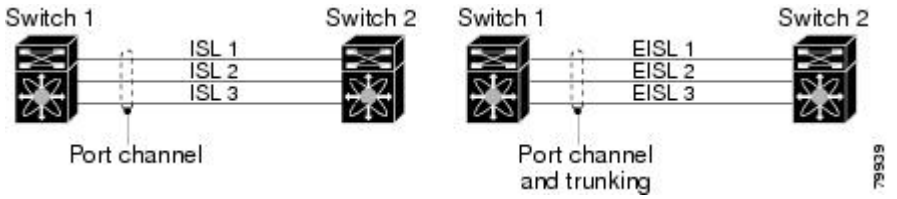

- Port channel interfaces can be channeled between the following port sets:

- 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 12: SID1, DID1, and Flow-Based Load Balancing

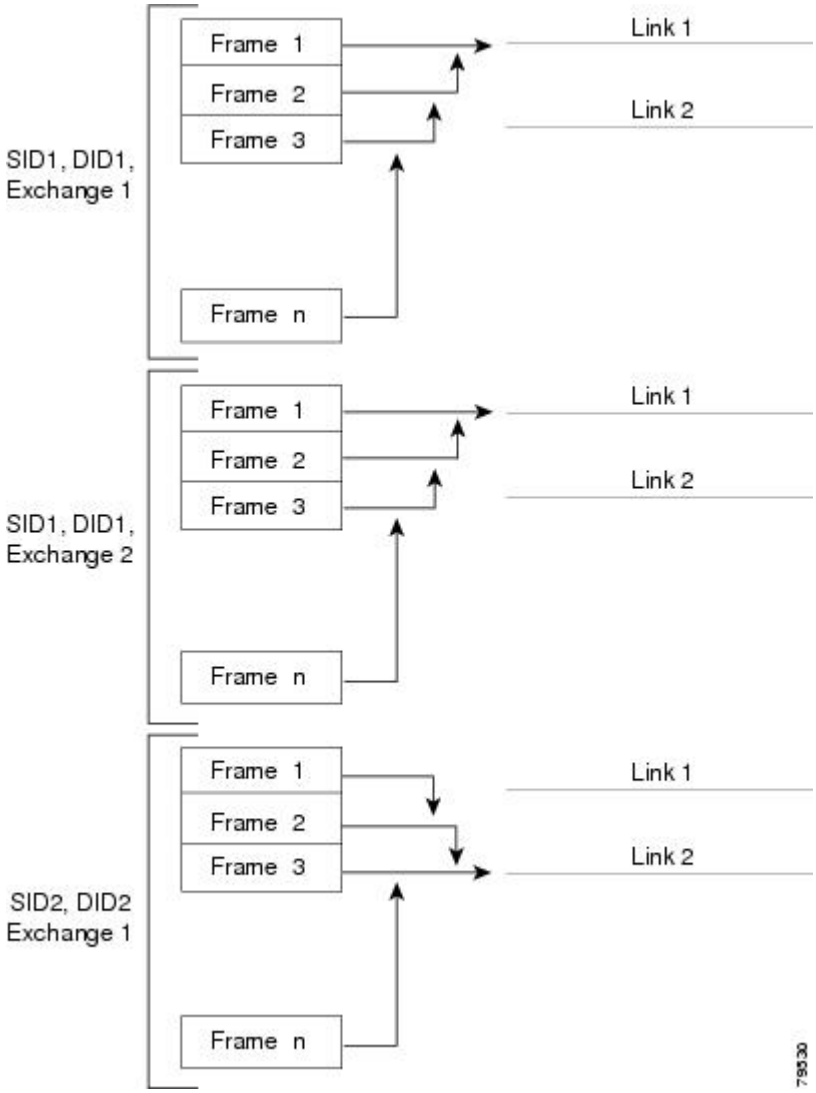

The following figure illustrates how exchange-based load balancing works. When the first frame in an exchange is received for forwarding on an interface, link 1 is chosen by a hash algorithm. All remaining frames in that

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 13: SID1, DID1, and Exchange-Based Load Balancing

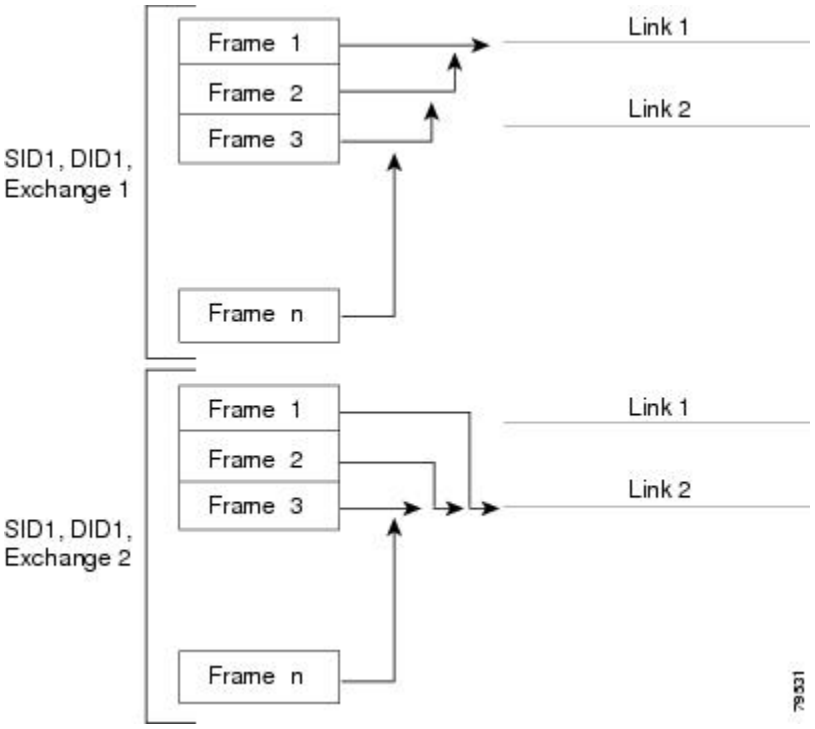

Configuring SAN Port Channels

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 14: Valid SAN Port Channel Configurations

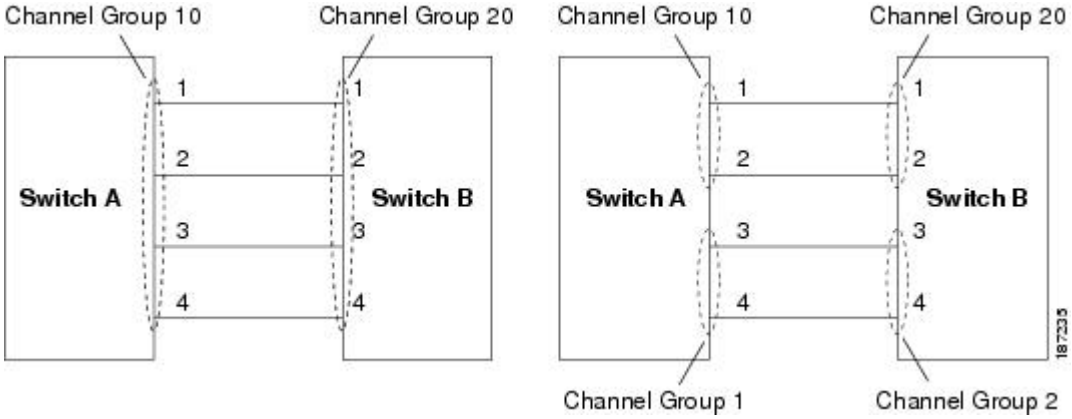

The following figure shows examples of invalid configurations. Assuming that the links are brought up in the 1, 2, 3, 4 sequence, links 3 and 4 will be operationally down as the fabric is misconfigured.

Figure 15: Misconfigured Configurations

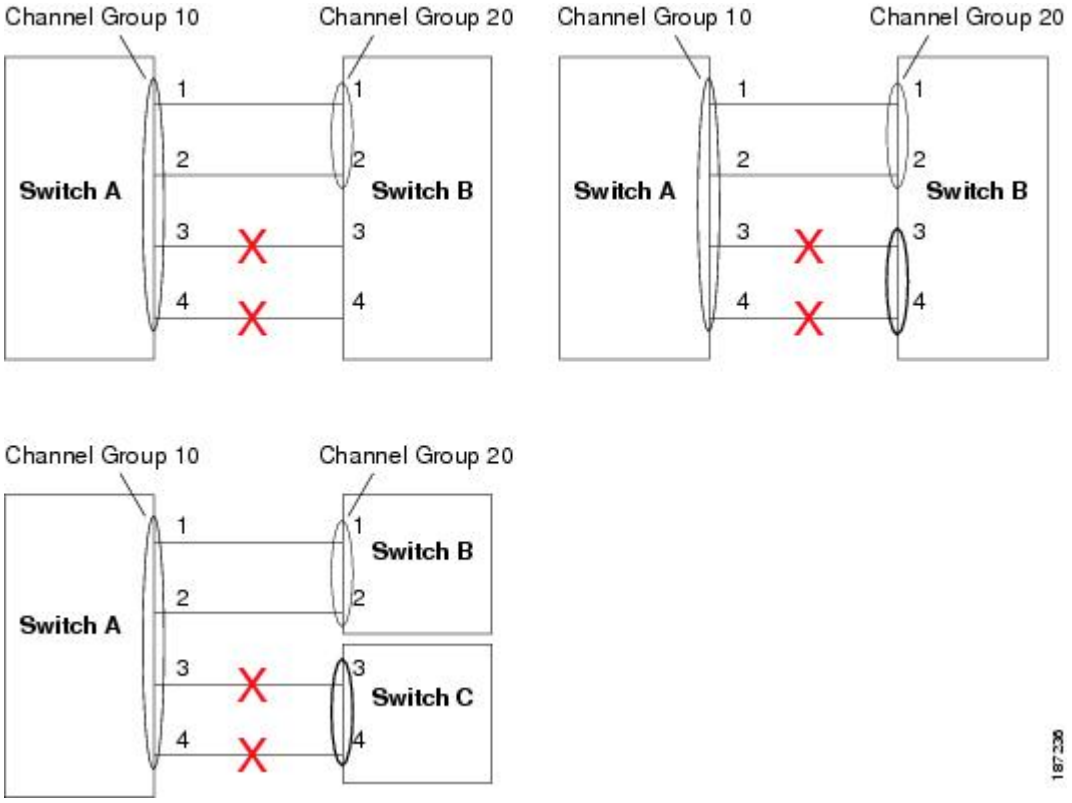

SAN Port Channel Configuration Guidelines

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

- Configure the SAN port channel using Fibre Channel ports from both expansion modules to provide increased availability (if one of the expansion modules failed).
- 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.
- Devices logged in through the F port channel on an MDS switch are not supported in IVR non-NAT configuration. The devices are supported only in IVR NAT configuration.
- Port security rules are enforced only on physical PWWNs at the single link level.
- The name server registration of the N ports logging in through an F port channel will use the FWWN of the port channel interface.
- DPVM configuration is not supported.
- The port channel port VSAN cannot be configured using Dynamic Port VSAN Membership (DPVM).

Creating a SAN Port Channel

To create a SAN port channel, perform this task:

SUMMARY STEPS

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

DETAILED STEPS

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

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 14: Channel Group Configuration Differences

On Mode	Active Mode
No protocol is exchanged.	A port channel protocol negotiation is performed with the peer ports.

On Mode	Active Mode
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.
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# **configuration 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# **configuration terminal**	Enters 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.

	Command or Action	Purpose
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# **configuration terminal**
2. switch(config)# **no interface san-port-channel** *channel-number*

DETAILED STEPS

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

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, allowed VSAN, and port security).
- 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.

Beginning with Cisco NX-OS Release 5.0(2)N2(1), 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:
  - QoS
  - Bandwidth
  - Delay
  - STP

- Service policy
- ACLs

When an interface joins or leaves a port channel, the following parameters remain unaffected:

- Beacon
- Description
- CDP
- LACP port priority
- Debounce
- UDLD
- 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# **configuration 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# **configuration terminal**	Enters 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# **configuration 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# **configuration terminal**	Enters 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(config)# **interface** *type slot*/*port*
2. switch(config-if)# **no channel-group** *channel-number*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch(config)# **interface** *type slot*/*port*	Enters configuration mode for the specified interface.
Step 2	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, or they can be automatically created. In both cases, 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. This helps make applications such as write acceleration work for SAN port channels over FCIP links.

• Autocreation protocol—Automatically aggregates compatible ports into a SAN port channel.

About Channel Group Creation

If channel group autocreation is enabled, ISLs can be configured automatically into channel groups without manual intervention. The following figure shows an example of channel group autocreation.

The first ISL comes up as an individual link. In the example shown in the following figure, this is link A1-B1. When the next link comes up (A2-B2 in the example), the port channel protocol determines if this link is compatible with link A1-B1 and automatically creates channel groups 10 and 20 in the respective switches. Link A3-B3 can join the channel groups (and the port channels) if the respective ports have compatible configurations. Link A4-B4 operates as an individual link, because it is not compatible with the existing member ports in the channel group.

Figure 16: Autocreating Channel Groups

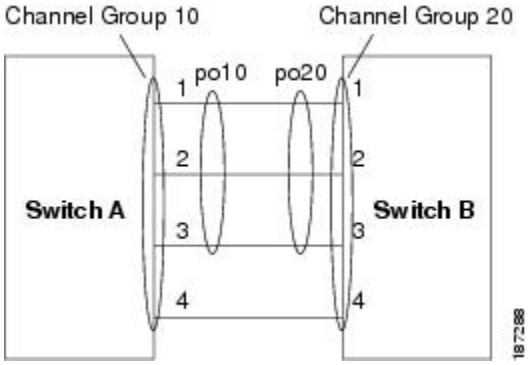

The channel group numbers are assigned dynamically (when the channel group is formed).

The channel group number may change across reboots for the same set of port channels depending on the initialization order of the ports.

The following table identifies the differences between user-configured and auto-configured channel groups.

Table 15: Channel Group Configuration Differences

User-Configured Channel Group	Autocreated Channel Group
Manually configured by the user.	Created automatically when compatible links come up between two compatible switches, if channel group autocreation is enabled in all ports at both ends.
Member ports cannot participate in autocreation of channel groups. The autocreation feature cannot be configured.	None of these ports are members of a user-configured channel group.
You can form the SAN port channel with a subset of the ports in the channel group. Incompatible ports remain in a suspended or isolated state depending on the On or Active mode configuration.	All ports included in the channel group participate in the SAN port channel. No member port becomes isolated or suspended; instead, the member port is removed from the channel group when the link is found to be incompatible.

User-Configured Channel Group	Autocreated Channel Group
Any administrative configuration made to the SAN port channel is applied to all ports in the channel group, and you can save the configuration for the port channel interface.	Any administrative configuration made to the SAN port channel is applied to all ports in the channel group, but the configurations are saved for the member ports; no configuration is saved for the port channel interface. You can explicitly convert this channel group, if required.
You can remove any channel group and add members to a channel group.	You cannot remove a channel group. You cannot add members to the channel group or remove members. The channel group is removed when no member ports exist.

Autocreation Guidelines

When using the autocreation protocol, follow these guidelines:

- A port is not allowed to be configured as part of a SAN port channel when the autocreation feature is enabled. These two configurations are mutually exclusive.
- Autocreation must be enabled in both the local and peer ports to negotiate a SAN port channel.
- Aggregation occurs in one of two ways:
  - A port is aggregated into a compatible autocreated SAN port channel.
  - A port is aggregated with another compatible port to form a new SAN port channel.
- Newly created SAN port channels are allocated from the maximum possible port channel number in a decreasing order based on availability. If all port channel numbers are used up, aggregation is not allowed.
- You cannot change the membership or delete an autocreated SAN port channel.
- When you disable autocreation, all member ports are removed from the autocreated SAN port channel.
- Once the last member is removed from an autocreated SAN port channel, the channel is automatically deleted and the number is released for reuse.
- An autocreated SAN port channel is not persistent through a reboot. An autocreated SAN port channel can be manually configured to appear the same as a persistent SAN port channel. Once the SAN port channel is made persistent, the autocreation feature is disabled in all member ports.
- You can enable or disable the autocreation feature on a per-port basis or for all ports in the switch. When this configuration is enabled, the channel group mode is assumed to be active. The default for this task is disabled.
- If autocreation of channel groups is enabled for an interface, you must first disable autocreation before downgrading to earlier software versions or before configuring the interface in a manually configured channel group.

Tip When enabling autocreation in any switch in the Cisco Nexus 5000 Series, we recommend that you retain at least one interconnected port between the switches without any autocreation configuration. If all ports between two switches are configured with the autocreation feature at the same time, a possible traffic disruption may occur between these two switches as ports are automatically disabled and reenabled when they are added to an autocreated SAN port channel.

Enabling and Configuring Autocreation

To configure automatic channel groups, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface** *type slot*/*port*
3. switch(config- if)# **channel-group auto**
4. switch(config- if)# **no channel-group auto**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **interface** *type slot*/*port*	Enters configuration mode for the specified interface.
Step 3	switch(config- if)# **channel-group auto**	Automatically creates the channel group for the selected interface(s).
Step 4	switch(config- if)# **no channel-group auto**	Disables the autocreation of channel groups for this interface, even if the system default configuration may have autocreation enabled.

Example of Configuring Autocreation

The following example configures an automatic channel group:

```
switch(config)# interface fc2/3
switch(config-if)# channel-group auto
```

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.

Converting to Manually Configured Channel Groups

You can convert autocreated channel group to a user-configured channel group using the **san-port-channel** *channel-group-number* persistent EXEC command. If the SAN port channel does not exist, this command is not executed.

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.

This example shows how to create the port channel:

```
switch(config)# interface port-channel 2
switch(config-if)# switchport mode F
switch(config-if)# switchport  dedicated
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 in dedicated mode:

```
switch(config)# interface fc1/4-6
switch(config-if)# shut
switch(config-if)# switchport mode F
switch(config-if)# switchport speed 4000
switch(config-if)# switchport rate-mode dedicated
switch(config-if)# switchport trunk mode on
switch(config-if)# channel-group 2
switch(config-if)# no shut
switch(config-if)# exit
```

This example shows how to create the port channel in dedicated mode on the NPV switch:

```
switch(config)# interface san-port-channel 2
switch(config-if)# switchport mode NP
switch(config-if)# no shut
switch(config-if)# exit
```

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

```
switch(config)# interface fc2/1-2
switch(config-if)# shut
switch(config-if)# switchport mode NP
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 (*).

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
```

Autocreated SAN port channels are indicated explicitly to help differentiate them from the manually created SAN port channels. The following example shows how to display an autocreated port channel:

```
switch# show interface fc2/1
fc2/1 is trunking
    Hardware is Fibre Channel, FCOT is short wave laser
    Port WWN is 20:0a:00:0b:5f:3b:fe:80
    ...
    Receive data field Size is 2112
    Beacon is turned off
    Port-channel auto creation is enabled
Belongs to port-channel 123
...
```

Default Settings for SAN Port Channels

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

Table 16: 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.

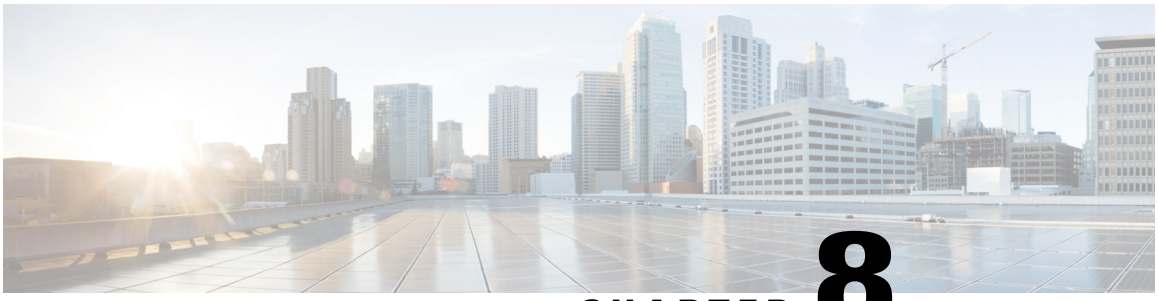

CHAPTER 8

Configuring and Managing VSANs

This chapter contains the following sections:

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

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

- Multiple VSANs can share the same physical topology.
- The same Fibre Channel IDs (FC IDs) can be assigned to a host in another VSAN, 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 17: Logical VSAN Segmentation

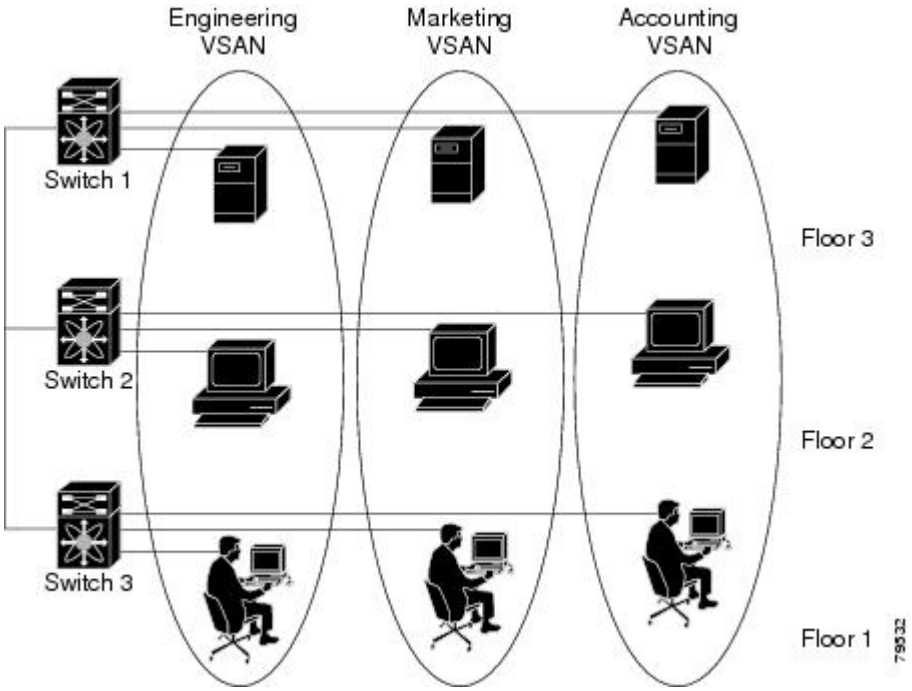

The application servers or storage arrays can be connected to the switch using Fibre Channel or virtual Fibre Channel interfaces. A VSAN can include a mixture of Fibre Channel and virtual Fibre Channel interfaces.

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 18: Example of Two VSANs

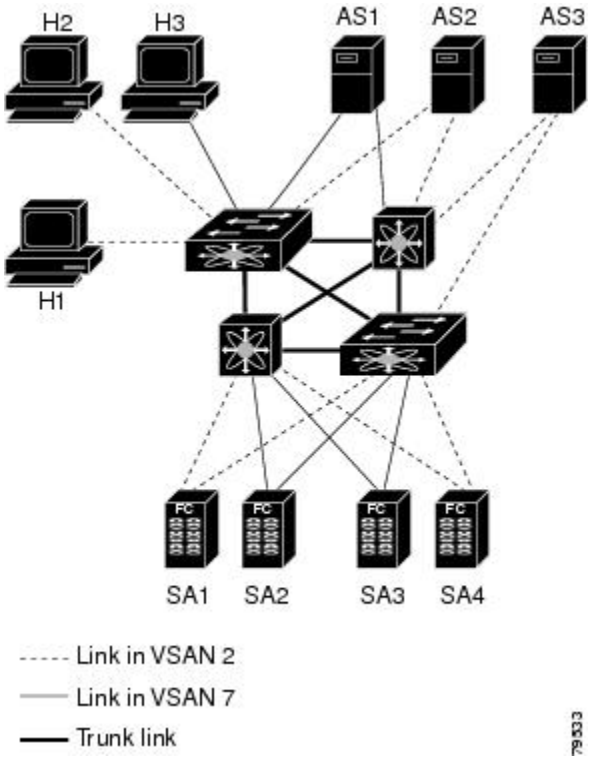

The four switches in this network are interconnected by VSAN trunk links that carry both VSAN 2 and

VSAN 7 traffic. You can configure a different inter-switch topology for each VSAN. In the preceeding 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 may be shared by multiple VSANs. VSANs allow SANs to be built on port granularity instead of switch granularity. The preceeding 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 256 VSANs can be configured in a switch. Of these, one is a default VSAN (VSAN 1), and another is an isolated VSAN (VSAN 4094). User-specified VSAN IDs range from 2 to 4093.

VSANs Versus Zones

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 17: 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.

VSAN Characteristic	Zone Characteristic
VSANs are defined for larger environments (storage service providers).	Zones are defined for a set of initiators and targets not visible outside the zone.
VSANs encompass the entire fabric.	Zones are configured at the fabric edge.

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 19: VSANS with Zoning

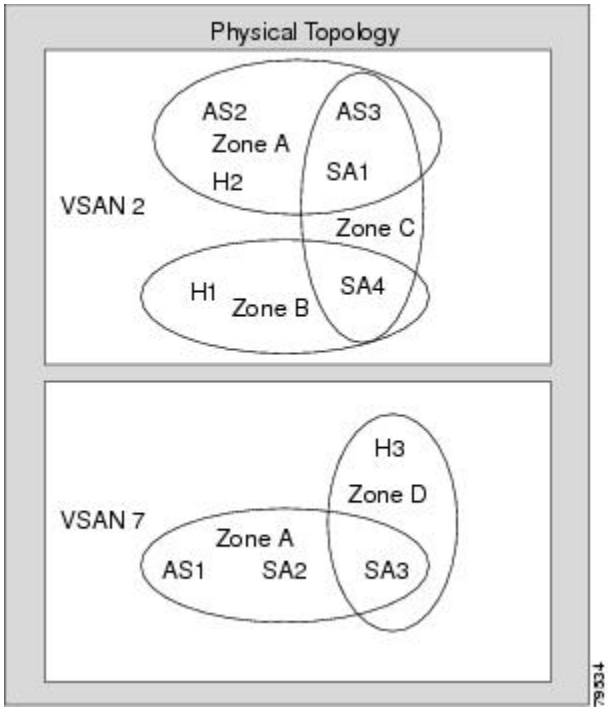

Configuring VSANs

VSANs have the following attributes:

- VSAN ID—The VSAN ID identifies the VSAN as the default VSAN (VSAN 1), user-defined VSANs (VSAN 2 to 4093), and the isolated VSAN (VSAN 4094).
- State—The administrative state of a VSAN can be configured to an active (default) or suspended state. Once VSANs are created, they may exist in various conditions or states.
  - The active state of a VSAN indicates that the VSAN is configured and enabled. By enabling a VSAN, you activate the services for that VSAN.

◦ The suspended state of a VSAN indicates that the VSAN is configured but not enabled. If a port is configured in this VSAN, it is disabled. Use this state to deactivate a VSAN without losing the VSAN's configuration. All ports in a suspended VSAN are disabled. By suspending a VSAN, you can preconfigure all the VSAN parameters for the whole fabric and activate the VSAN immediately.

- VSAN name—This text string identifies the VSAN for management purposes. The name can be from 1 to 32 characters long and it must be unique across all VSANs. By default, the VSAN name is a concatenation of VSAN and a four-digit string representing the VSAN ID. For example, the default name for VSAN 3 is VSAN0003.

Note A VSAN name must be unique.

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

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.

To create VSANs, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **vsan database**
3. switch(config-vsan-db)# **vsan** *vsan-id*
4. switch(config-vsan-db)# **vsan** *vsan-id* **name** *name*
5. switch(config-vsan-db)# **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	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **vsan database**	Configures the database for a VSAN. Application specific VSAN parameters cannot be configured from this prompt.

	Command or Action	Purpose
Step 3	switch(config-vsan-db)# **vsan** *vsan-id*	Creates a VSAN with the specified ID if that VSAN does not exist already.
Step 4	switch(config-vsan-db)# **vsan** *vsan-id* **name** *name*	Updates the VSAN with the assigned name.
Step 5	switch(config-vsan-db)# **vsan** *vsan-id* **suspend**	Suspends the selected VSAN.
Step 6	switch(config-vsan-db)# **no vsan** *vsan-id* **suspend**	Negates the **suspend** command issued in the previous step.
Step 7	switch(config-vsan-db)# **end**	Returns you to EXEC mode.

About Port VSAN Membership

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

- Statically—Assigning VSANs to ports.
- Dynamically—Assigning VSANs based on the device WWN. This method is referred to as dynamic port VSAN membership (DPVM). Cisco Nexus 5000 Series switches do not support DPVM.

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

Related Topics

Assigning Static Port VSAN Membership, on page 95

Assigning Static Port VSAN Membership

To statically assign VSAN membership for an interface port, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **vsan database**
3. switch(config-vsan-db)# **vsan** *vsan-id*
4. switch(config-vsan-db)# **vsan** *vsan-id* **interface** {*fc slot*/*port* | **vfc** *vfc-id*}
5. switch(config-vsan-db)# **vsan** *vsan-id* {*fc slot*/*port* | **vfc** *vfc-id*}

DETAILED STEPS

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

	Command or Action	Purpose
Step 2	switch(config)# **vsan database**	Configures the database for a VSAN.
Step 3	switch(config-vsan-db)# **vsan** *vsan-id*	Creates a VSAN with the specified ID if that VSAN does not exist already.
Step 4	switch(config-vsan-db)# **vsan** *vsan-id* **interface** {*fc slot*/*port* | **vfc** *vfc-id*}	Assigns the membership of the specified interface to the VSAN.
Step 5	switch(config-vsan-db)# **vsan** *vsan-id* {*fc slot*/*port* | **vfc** *vfc-id*}	Updates the membership information of the interface to reflect the changed VSAN. **Note** To remove the VSAN membership of a FC or vFC interface, assign the VSAN membership of that interface to another VSAN. Cisco recommends that you assign it to VSAN 1.

Displaying VSAN Static Membership

To display the VSAN static membership information, use the **show vsan membership** command.

The following example displays membership information for the specified VSAN:

```
switch # show vsan 1 membership
vsan 1 interfaces:
        fc2/1   fc2/2   fc2/3   fc2/4

        san-port-channel 3  vfc1/1
```

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

The following example displays membership information for all VSANs:

```
switch # show vsan membership
vsan 1 interfaces:
        fc2/1   fc2/2   fc2/3   fc2/4

        san-port-channel 3  vfc3/1
vsan 2 interfaces:
        fc2/3 vfc4/1
vsan 7 interfaces:
vsan 100 interfaces:
vsan 4094(isolated vsan) interfaces:
```

The following example displays static membership information for the specified interface:

```
switch # show vsan membership interface fc2/1
fc2/1
        vsan:1
        allowed list:1-4093
```

About the Default VSAN

The factory settings for switches in the Cisco Nexus 5000 Series 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 256 VSANs can be configured in a switch. Of these, one is a default VSAN (VSAN 1), and another is an isolated VSAN (VSAN 4094). User-specified VSAN IDs range from 2 to 4093.

About the Isolated VSAN

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

Displaying Isolated VSAN Membership

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

Operational State of a VSAN

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

About Static VSAN Deletion

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

- VSAN attributes and port membership details are maintained by the VSAN manager. This feature is affected when you delete a VSAN from the configuration. When a VSAN is deleted, all the ports in that VSAN are made inactive and the ports are moved to the isolated VSAN. If the same VSAN is recreated,

the ports do not automatically get assigned to that VSAN. You must explicitly reconfigure the port VSAN membership (see the figure below).

Figure 20: VSAN Port Membership Details

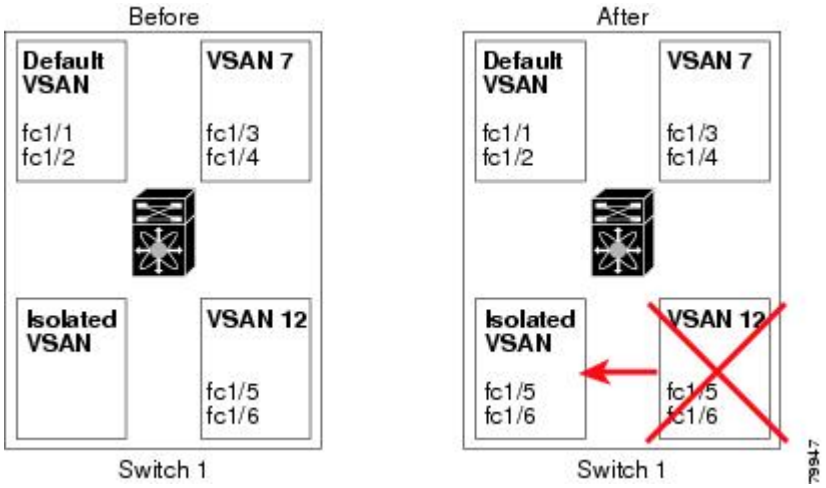

- VSAN-based runtime (name server), zoning, and configuration (static routes) information is removed when the VSAN is deleted.
- Configured VSAN interface information is removed when the VSAN is deleted.

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, then a command request to move a port to VSAN 10 is rejected.

Deleting Static VSANs

To delete a VSAN and its various attributes, perform this task:

SUMMARY STEPS

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

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **vsan database**	Configures the VSAN database.
Step 3	switch-config-db# **vsan 2**	Places you in VSAN configuration mode.
Step 4	switch(config-vsan-db)# **no vsan 5**	Deletes VSAN 5 from the database and switch.
Step 5	switch(config-vsan-db)# **end**	Places you in EXEC mode.

About Load Balancing

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

Configuring Load Balancing

To configure load balancing on an existing VSAN, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **vsan database**
3. switch(config-vsan-db)# **vsan** *vsan-id*
4. switch(config-vsan-db)# **vsan** *vsan-id* **loadbalancing src-dst-id**
5. switch(config-vsan-db)# **no vsan** *vsan-id* **loadbalancing src-dst-id**
6. switch(config-vsan-db)# **vsan** *vsan-id* **loadbalancing src-dst-ox-id**
7. switch(config-vsan-db)# **vsan** *vsan-id* **suspend**
8. switch(config-vsan-db)# **no vsan** *vsan-id* **suspend**
9. switch(config-vsan-db)# **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **vsan database**	Enters VSAN database configuration submode
Step 3	switch(config-vsan-db)# **vsan** *vsan-id*	Specifies an existing VSAN.

	Command or Action	Purpose
Step 4	switch(config-vsan-db)# **vsan** *vsan-id* **loadbalancing src-dst-id**	Enables the load-balancing guarantee for the selected VSAN and directs the switch to use the source and destination ID for its path selection process.
Step 5	switch(config-vsan-db)# **no vsan** *vsan-id* **loadbalancing src-dst-id**	Negates the command entered in the previous step and reverts to the default values of the load-balancing parameters.
Step 6	switch(config-vsan-db)# **vsan** *vsan-id* **loadbalancing src-dst-ox-id**	Changes the path selection setting to use the source ID, the destination ID, and the OX ID (default).
Step 7	switch(config-vsan-db)# **vsan** *vsan-id* **suspend**	Suspends the selected VSAN.
Step 8	switch(config-vsan-db)# **no vsan** *vsan-id* **suspend**	Negates the **suspend** command entered in the previous step.
Step 9	switch(config-vsan-db)# **end**	Returns you to EXEC mode.

About 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 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 VSAN Settings

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

Table 18: Default VSAN Parameters

Parameters	Default
Default VSAN	VSAN 1.
State	Active state.

Parameters	Default
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).

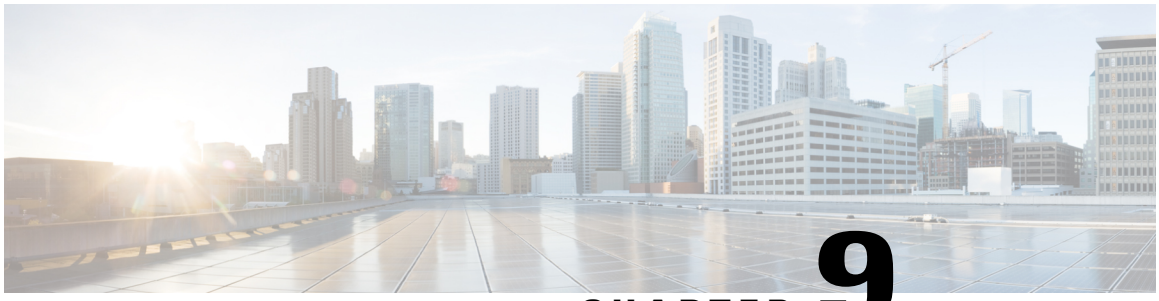

CHAPTER 9

Configuring and Managing Zones

This chapter contains the following sections:

Configuring and Managing Zones

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

Advanced zoning capabilities specified in the FC-GS-4 and FC-SW-3 standards are 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 500 zone sets.

- Zoning can be administered from any switch in the fabric.
  - When you activate a zone (from any switch), all switches in the fabric receive the active zone set. Additionally, full zone sets are distributed to all switches in the fabric, if this feature is enabled in the source switch.
  - If a new switch is added to an existing fabric, zone sets are acquired by the new switch.

- Zone changes can be configured 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 identifiers:
  - Port world wide name (pWWN)—Specifies the pWWN of an N port attached to the switch as a member of the zone.
  - Fabric pWWN—Specifies the WWN of the fabric port (switch port's WWN). This membership is also referred to as port-based zoning.
  - FC ID—Specifies the FC ID of an N port attached to the switch as a member of the zone.
  - Interface and switch WWN (sWWN)—Specifies the interface of a switch identified by the sWWN. This membership is also referred to as interface-based zoning.
  - Interface and domain ID—Specifies the interface of a switch identified by the domain ID.
  - Domain ID and port number—Specifies the domain ID of a Cisco switch domain and additionally specifies a port belonging to a non-Cisco switch.

Note For N ports attached to the switch over a virtual Fibre Channel interface, you can specify zone membership using the 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.

Note Interface-based zoning only works with Cisco SAN switches. Interface-based zoning does not work for VSANs configured in interop 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 21: Fabric with Two Zones

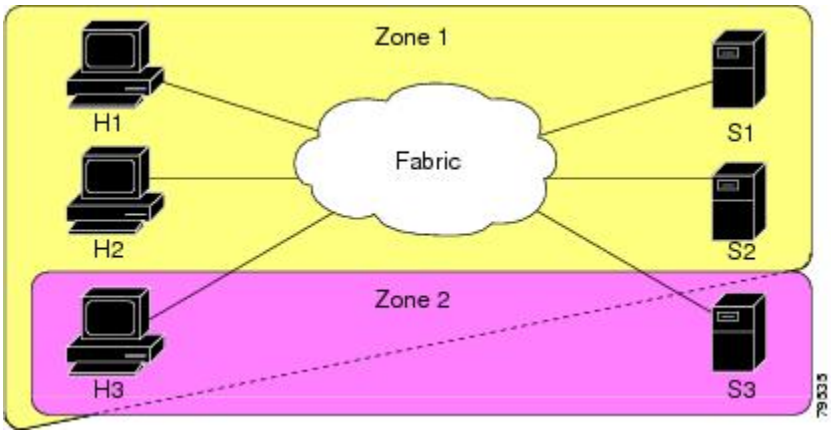

You can use other ways to partition this fabric into zones. The following figure shows another possibility. Assume that there is a need to isolate storage system S2 for the purpose of testing new software. To achieve this, zone 3 is configured, which contains only host H2 and storage S2. You can restrict access to only H2 and S2 in zone 3, and to H1 and S1 in zone 1.

Figure 22: Fabric with Three Zones

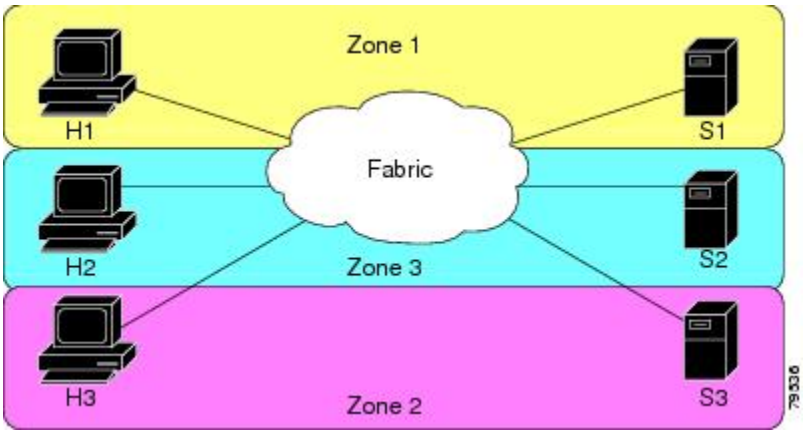

Zone Implementation

Cisco Nexus 5000 Series switches automatically support the following basic zone features (no additional configuration is required):

- Zones are contained in a VSAN.

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

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

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

Active and Full Zone Set Configuration Guidelines

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 23: Active and Full Zone Sets

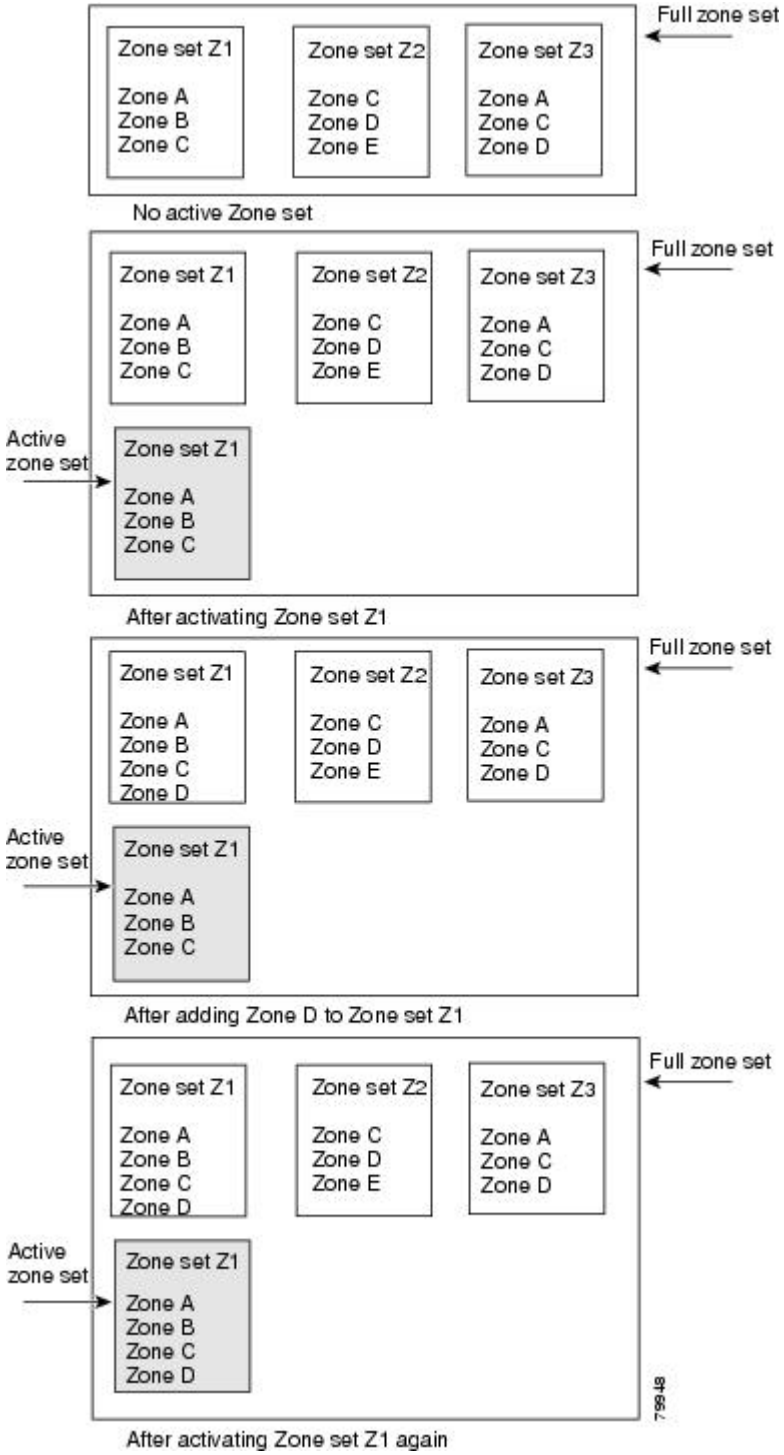

Configuring Zones

To configure a zone and assign a zone name, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **zone name** *zone-name* **vsan** *vsan-id*
3. switch(config-zone)# **member** *type value*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **zone name** *zone-name* **vsan** *vsan-id*	Configures a zone in the specified VSAN. **Note** All alphanumeric characters or one of the following symbols ($, -, ^, _) are supported.
Step 3	switch(config-zone)# **member** *type value*	Configures a member for the specified zone based on the type (pWWN, fabric pWWN, FC ID, fcalias, domain ID, or interface) and value specified. **Caution** You must only configure pWWN-type zoning on all SAN switches running Cisco NX-OS if there is a Cisco MDS 9020 switch running FabricWare in the same fabric. **Tip** Use a relevant display command (for example, **show interface** or **show flogi database**) to obtain the required value in hex format.

Configuring Zones Example

Table 19: Type and Value Syntax for the member Command

Domain ID	**member domain-id** *domain-id* **portnumber** *number*
FC alias	**member fcalias** *fc-alias-name*
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*

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
```

Domain ID example:

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

Local sWWN interface example:

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

Remote sWWN interface example:

```
switch(config-zone)# member interface fc 2/1 swwn 20:00:00:05:30:00:4a:de
```

Domain ID interface example:

```
switch(config-zone)# member interface fc 2/1 domain-id 25
```

Zone Sets

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

Figure 24: Hierarchy of Zone Sets, Zones, and Zone Members

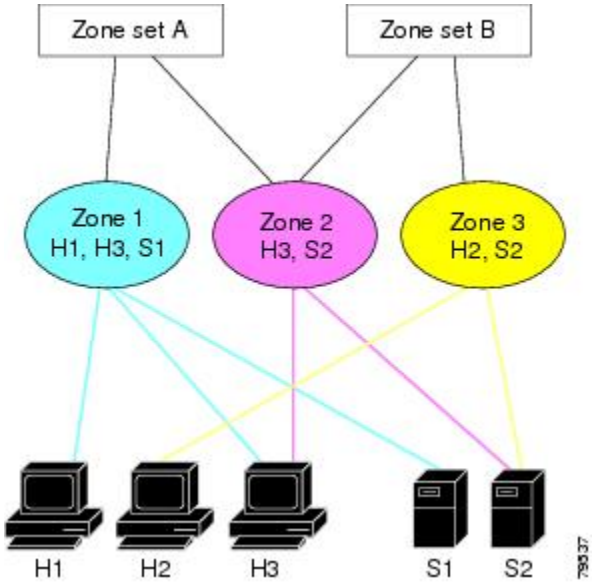

Zones provide a method for specifying access control, while zone sets are a grouping of zones to enforce access control in the fabric. Either zone set A or zone set B can be activated (but not together).

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

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

To activate or deactivate an existing zone set, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **zoneset activate name** *zoneset-name* **vsan** *vsan-id*
3. switch(config)# **no zoneset activate name** *zoneset-name* **vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **zoneset activate name** *zoneset-name* **vsan** *vsan-id*	Activates the specified zone set.
Step 3	switch(config)# **no zoneset activate name** *zoneset-name* **vsan** *vsan-id*	Deactivates the specified zone set.

About the Default Zone

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

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

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

Note When the switch is initialized for the first time, no zones are configured and all members are considered to be part of the default zone. Members are not permitted to 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. switch# **configuration terminal**
2. switch(config)# **zone default-zone permit vsan** *vsan-id*
3. switch(config)# **no zone default-zone permit vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **zone default-zone permit vsan** *vsan-id*	Permits traffic flow to default zone members.
Step 3	switch(config)# **no zone default-zone permit vsan** *vsan-id*	Denies (default) traffic flow to default zone members.

About FC Alias Creation

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

- pWWN—The WWN of the N 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.

Tip The switch supports a maximum of 2048 aliases per VSAN.

Creating FC Aliases

To create an alias, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcalias name AliasSample vsan** *vsan-id*
3. switch(config-fcalias)# **member** *type value*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fcalias name AliasSample vsan** *vsan-id*	Configures an alias name (AliasSample).
Step 3	switch(config-fcalias)# **member** *type value*	Configures a member for the specified fcalias (AliasSample) based on the type (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 20: 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
```

Local sWWN interface example:

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

Remote sWWN interface example:

```
switch(config-fcalias)# member interface fc 2/1 swwn 20:00:00:05:30:00:4a:de
```

Domain ID interface example:

```
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

To create a zone set to include several zones, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **zone set name** *zoneset-name* **vsan** *vsan-id*
3. switch(config-zoneset)# **member** *name*
4. switch(config-zoneset)# **zone name** *zone-name*
5. switch(config-zoneset-zone)# **member fcid** *fcid*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **zone set name** *zoneset-name* **vsan** *vsan-id*	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	switch(config-zoneset)# **member** *name*	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	switch(config-zoneset)# **zone name** *zone-name*	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	switch(config-zoneset-zone)# **member fcid** *fcid*	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 Nexus 5000 Series 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 21: 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 switches in the Cisco Nexus 5000 Series 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.

To enable full zone set and active zone set distribution to all switches on a per VSAN basis, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **zoneset distribute full vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **zoneset distribute full vsan** *vsan-id*	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 2004
```

About Recovering from Link Isolation

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

- Import the neighboring switch's active zone set database and replace the current active zone set (see 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 25: Importing and Exporting the Database

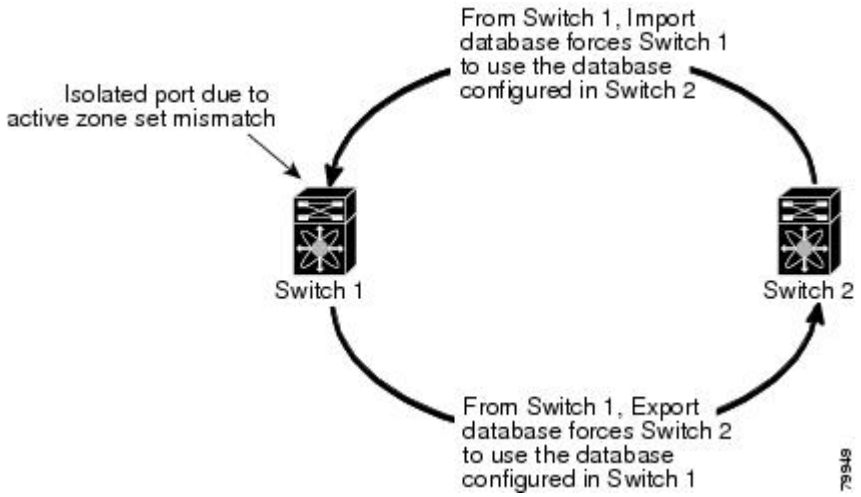

Importing and Exporting Zone Sets

To import or export the zone set information from or to an adjacent switch, perform this task:

SUMMARY STEPS

1. switch# **zoneset import interface fc** *slot*/*port* **vsan** *vsan-id*
2. switch# **zoneset import interface fc** *slot*/*port* **vsan** *vsan-id*
3. switch# **zoneset export vsan** *vsan-id*
4. switch# **zoneset export vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **zoneset import interface fc** *slot*/*port* **vsan** *vsan-id*	Imports the zone set from the adjacent switch connected through the specified interface for the VSAN .
Step 2	switch# **zoneset import interface fc** *slot*/*port* **vsan** *vsan-id*	Imports the zone set from the adjacent switch connected through the specified interface for the VSAN range.
Step 3	switch# **zoneset export vsan** *vsan-id*	Exports the zone set to the adjacent switch connected through the specified VSAN.
Step 4	switch# **zoneset export vsan** *vsan-id*	Exports the zone set to the adjacent switch connected through the specified 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 Nexus 5000 Series 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.

To make a copy of a zone set, perform this task:

SUMMARY STEPS

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

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **zone copy active-zoneset full-zoneset vsan** *vsan-id*	Makes a copy of the active zone set in the specified VSAN to the full zone set.
Step 2	switch# **zone copy vsan** *vsan-id* **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

To rename a zone, zone set, fcalias, or zone-attribute-group, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **zoneset rename** oldname newname **vsan** *vsan-id*
3. switch(config)# **zone rename** oldname newname **vsan** *vsan-id*
4. switch(config)# **fcalias rename** oldname newname **vsan** *vsan-id*
5. switch(config)# **zone-attribute-group rename** oldname newname **vsan** *vsan-id*
6. switch(config)# **zoneset activate name** newname **vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **zoneset rename** oldname newname **vsan** *vsan-id*	Renames a zone set in the specified VSAN.
Step 3	switch(config)# **zone rename** oldname newname **vsan** *vsan-id*	Renames a zone in the specified VSAN.
Step 4	switch(config)# **fcalias rename** oldname newname **vsan** *vsan-id*	Renames a fcalias in the specified VSAN.
Step 5	switch(config)# **zone-attribute-group rename** oldname newname **vsan** *vsan-id*	Renames a zone attribute group in the specified VSAN.
Step 6	switch(config)# **zoneset activate name** newname **vsan** *vsan-id*	Activates the zone set and updates the new zone name in the active zone set.

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

To clone a zone, zone set, fcalias, or zone-attribute-group, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **zoneset clone** *oldname newname* **vsan** *vsan-id*
3. switch(config)# **zone clone** *oldname newname* **vsan** *number*
4. switch(config)# **fcalias clone** *oldname newname* **vsan** *vsan-id*
5. switch(config)# **zone-attribute-group clone** *oldname newname* **vsan** *vsan-id*
6. switch(config)# **zoneset activate name** *newname* **vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **zoneset clone** *oldname newname* **vsan** *vsan-id*	Clones a zone set in the specified VSAN.
Step 3	switch(config)# **zone clone** *oldname newname* **vsan** *number*	Clones a zone in the specified VSAN.
Step 4	switch(config)# **fcalias clone** *oldname newname* **vsan** *vsan-id*	Clones a fcalias in the specified VSAN.
Step 5	switch(config)# **zone-attribute-group clone** *oldname newname* **vsan** *vsan-id*	Clones a zone attribute group in the specified VSAN.
Step 6	switch(config)# **zoneset activate name** *newname* **vsan** *vsan-id*	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 Zone Information

You can view any zone information by using the **show** command. If you request information for a specific object (for example, a specific zone, zone set, VSAN, or alias, or keywords such as brief or active), only information for the specified object is displayed.

The following example shows how to display zone information for all VSANs:

```
switch# show zone
```

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

```
switch# show zone vsan 1
```

The following example shows how to display the configured zone sets for a range of VSANs:

```
switch# show zoneset vsan 2-3
```

The following example shows how to display the members of a specific zone:

```
switch# show zone name Zone1
```

The following example shows how to display fcalias configuration:

```
switch# show fcalias vsan 1
```

The following example shows how to display all zones to which a member belongs:

```
switch# show zone member pwwn 21:00:00:20:37:9c:48:e5
```

The following example shows how to display the number of control frames exchanged with other switches:

```
switch# show zone statistics
```

The following example shows how to display the active zone set:

```
switch# show zoneset active
```

The following example shows how to display the active zones:

```
switch# show zone active
```

The following example shows how to display the zone status:

```
switch# show zone status
```

Enhanced Zoning

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

About Enhanced Zoning

The following table lists the advantages of the enhanced zoning feature in all switches in the Cisco Nexus 5000 Series.

Table 22: 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.

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

Changing from Basic Zoning to Enhanced Zoning

To change to the enhanced zoning mode from the basic mode, perform this task:

SUMMARY STEPS

1. Verify that all switches in the fabric are capable of working in the enhanced mode.
2. If one or more switches are not capable of working in enhanced mode, then your request to move to enhanced mode is rejected.
3. Set the operation mode to enhanced zoning mode.

DETAILED STEPS

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

Step 2 If one or more switches are not capable of working in 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 SAN switches allow you to change from enhanced zoning to basic zoning to enable you to downgrade and upgrade to other Cisco NX-OS releases.

To change to the basic zoning mode from the enhanced mode, perform this task:

SUMMARY STEPS

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

DETAILED STEPS

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

By default, the enhanced zoning feature is disabled in all switches in the Cisco Nexus 5000 Series.

To enable enhanced zoning in a VSAN, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **zone mode enhanced vsan** *vsan-id*
3. switch(config)# **no zone mode enhanced vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **zone mode enhanced vsan** *vsan-id*	Enables enhanced zoning in the specified VSAN.
Step 3	switch(config)# **no zone mode enhanced vsan** *vsan-id*	Disables enhanced zoning in the specified VSAN.

Modifying the Zone Database

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

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

To commit or discard changes to the zoning database in a VSAN, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **zone commit vsan** *vsan-id*
3. switch(config)# **zone commit vsan** *vsan-id* **force**
4. switch(config)# **no zone commit vsan** *vsan-id*
5. switch(config)# **no zone commit vsan** *vsan-id* **force**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **zone commit vsan** *vsan-id*	Applies the changes to the enhanced zone database and closes the session.
Step 3	switch(config)# **zone commit vsan** *vsan-id* **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*	Discards the changes to the enhanced zone database and closes the session.

	Command or Action	Purpose
Step 5	switch(config)# **no zone commit vsan** *vsan-id* **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# configuration 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.

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 23: Database Zone Merge Status

Local Database	Adjacent Database	Merge Status	Results of the Merge
The databases contain zone sets with the same nameIn the enhanced zoning mode, the active zone set does not have a name in interop mode 1. The zone set names are only present for full zone sets. but different zones, aliases, and attributes groups.		Successful.	ISLs are isolated.
The databases contains a zone, zone alias, or zone attribute group object with same name1 but different members.		Failed.	The adjacent database information populates the local database.
Empty.	Contains data.	Successful.	The union of the local and adjacent databases.

Local Database	Adjacent Database	Merge Status	Results of the Merge
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, then the ISL is isolated.
- If the protocol versions are the same, then the zone policies are compared. If the zone policies differ, then the ISL is isolated.
- If the zone merge options are the same, then 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, then the merge rules are used to perform the merge.

Configuring Zone Merge Control Policies

To configure merge control policies, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **zone merge-control restrict vsan** *vsan-id*
3. switch(config)# **no zone merge-control restrict vsan** *vsan-id*
4. switch(config)# **zone commit vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **zone merge-control restrict vsan** *vsan-id*	Configures a restricted merge control setting for this VSAN.
Step 3	switch(config)# **no zone merge-control restrict vsan** *vsan-id*	Defaults to using the allow merge control setting for this VSAN.
Step 4	switch(config)# **zone commit vsan** *vsan-id*	Commits the changes made to the specified VSAN.

Default Zone Policies

To permit or deny traffic in the default zone, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **zone default-zone permit vsan** *vsan-id*
3. switch(config)# **no zone default-zone permit vsan** *vsan-id*
4. switch(config)# **zone commit vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **zone default-zone permit vsan** *vsan-id*	Permits traffic flow to default zone members.
Step 3	switch(config)# **no zone default-zone permit vsan** *vsan-id*	Denies traffic flow to default zone members and reverts to factory default.
Step 4	switch(config)# **zone commit vsan** *vsan-id*	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. To configure switch-wide default settings, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **system default zone default-zone permit**
3. switch(config)# **no system default zone default-zone permit**
4. switch(config)# **system default zone distribute full**
5. switch(config)# **no system default zone distribute full**

DETAILED STEPS

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

	Command or Action	Purpose
Step 2	switch(config)# **system default zone default-zone permit**	Configures permit as the default zoning policy for new VSANs on the switch.
Step 3	switch(config)# **no system default zone default-zone permit**	Configures deny (default) as the default zoning policy for new VSANs on the switch.
Step 4	switch(config)# **system default zone distribute full**	Enables full zone database distribution as the default for new VSANs on the switch.
Step 5	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.

Verifying Enhanced Zone Information

The following example shows how to display the zone status for a specified VSAN:

```
switch# show zone status vsan 2
```

Compacting the Zone Database

You can delete excess zones and compact the zone database for the VSAN.

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

To delete zones and compact the zone database for a VSAN, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **no zone name** *zone-name* **vsan** *vsan-id*
3. switch(config)# **zone compact vsan** *vsan-id*

DETAILED STEPS

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

	Command or Action	Purpose
Step 2	switch(config)# **no zone name** *zone-name* **vsan** *vsan-id*	Deletes a zone to reduce the number of zones to 2000 or fewer.
Step 3	switch(config)# **zone compact vsan** *vsan-id*	Compacts the zone database for the specified VSAN to recover the zone ID released when a zone was deleted.

Zone and Zone Set Analysis

To better manage the zones and zone sets on your switch, you can display zone and zone set information using the **show zone analysis** command.

The following example shows how to display full zoning analysis:

```
switch# show zone analysis vsan 1
```

The following example shows how to display active zoning analysis:

```
switch# show zone analysis active vsan 1
```

See the Cisco Nexus 5000 Series *Switch Command Reference* for the description of the information displayed in the command output.

Default Basic Zone Settings

The following table lists the default settings for basic zone parameters.

Table 24: 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.

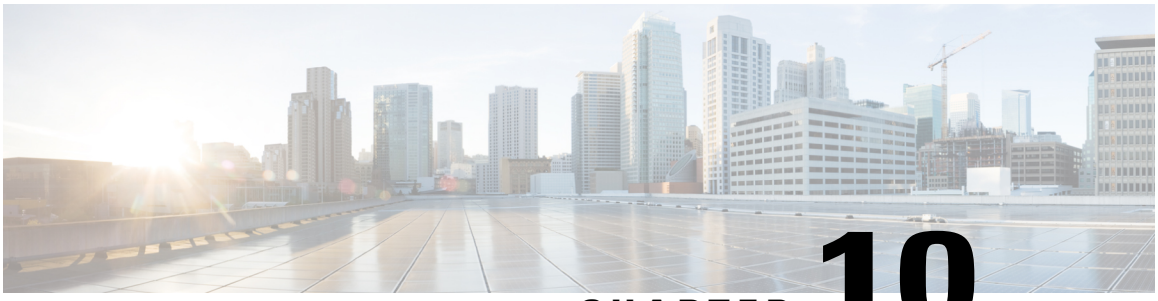

CHAPTER 10

Distributing Device Alias Services

This chapter contains the following sections:

Distributing Device Alias Services

Switches in the Cisco Nexus 5000 Series support Distributed Device Alias Services (device aliases) on a fabric-wide basis.

Information About Device Aliases

When the port WWN (pWWN) of a device must be specified to configure features (for example, zoning, DPVM, or port security) in a Cisco Nexus 5000 Series 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.
- You can import legacy zone alias configurations without losing data.
- The device alias application uses the Cisco Fabric Services (CFS) infrastructure to enable efficient database management and distribution. Device aliases use the coordinated distribution mode and the fabric-wide distribution scope.
- Basic and enhanced modes.
- Device aliases used to configure zones, IVR zones, or port security features are displayed automatically with their respective pWWNs in the **show** command output.

For additional information, refer to Using Cisco Fabric Services in the *Cisco Nexus 5000 Series System Management Configuration Guide*.

Related Topics

Device Alias Modes, on page 134

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)

Zone Aliases Versus Device Aliases

The following table compares the configuration differences between zone-based alias configuration and device alias configuration.

Table 25: Comparison Between Zone Aliases and Device Aliases

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

Device Alias Databases

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

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

To a create a device alias in the pending database, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **device-alias database**
3. switch(config-device-alias-db)# **device-alias name** *device-name* **pwwn** *pwwn-id*
4. switch(config-device-alias-db)# **no device-alias name** *device-name*
5. switch(config-device-alias-db)# **device-alias rename** *old-device-name new-device-name*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **device-alias database**	Enters the pending database configuration submode.
Step 3	switch(config-device-alias-db)# **device-alias name** *device-name* **pwwn** *pwwn-id*	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	switch(config-device-alias-db)# **no device-alias name** *device-name*	Removes the device name for the device that is identified by its pWWN.
Step 5	switch(config-device-alias-db)# **device-alias rename** *old-device-name new-device-name*	Renames an existing device alias with a new name.

Example of Creating a Device Alias

To display the device alias configuration, use the **show device-alias name** command:

```
switch# show device-alias name x
device-alias name x pwwn 21:01:00:e0:8b:2e:80:93
```

Device Alias Modes

You can specify that aliases operate in basic or enhanced 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.

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 applications such as zone server, PSM or DPVM 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 will fail in interop mode VSANs if the corresponding zones have native device alias-based members.

Changing Device Alias Mode Guidelines

When changing device alias modes, follow these guidelines:

- If two fabrics running in different device alias modes are joined together, the device alias merge will fail. There is no automatic conversion to one mode or the other during the merge process. In this situation, you must to 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 will automatically stop enforcing the corresponding device alias. If that corresponding device alias is part of an active zoneset, 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 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 zoneset.

- If a device alias name is mapped to a new HBA's pWWN, then 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

To configure device aliases to operate in enhanced mode, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **device-alias mode enhanced**
3. switch(config)# **no device-alias mode enhance**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **device-alias mode enhanced**	Assigns the device alias to operate in enhanced mode.
Step 3	switch(config)# **no device-alias mode enhance**	Assigns the device alias to operate in basic mode.

Viewing the Device Alias Mode Setting

To view the current device alias mode setting, enter the **show device-alias status** command.

```
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
```

About 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 not be any pending database and commit or abort operations either. If you have not committed the changes and you disable distribution, then a commit task will fail.

The following example displays 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

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.

To commit the changes, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **device-alias commit**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **device-alias commit**	Commits the changes made to the currently active session.

Discarding 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.

To discard the device alias session, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **device-alias abort**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **device-alias abort**	Discards the currently active session.

Displaying the Status of a Discard Operation

To display the status of the discard operation, use the show **device alias status** command.

```
switch# show device-alias status
Fabric Distribution: Enabled
Database:- Device Aliases 24
Status of the last CFS operation issued from this switch:
=========================================================
Operation: Abort
Status: Success
```

Fabric Lock Override

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
```

To display the status of the clear operation, use the **show device-alias status** command.

```
switch# show device-alias status
Fabric Distribution: Enabled
Database:- Device Aliases 24
Status of the last CFS operation issued from this switch:
==========================================================
Operation: Clear Session<--------------------Lock released by administrator
Status: Success<-----------------------------Successful status of the operation
```

Disabling and Enabling Device Alias Distribution

To disable or enable the device alias distribution, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **no device-alias distribute**
3. switch(config)# **device-alias distribute**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **no device-alias distribute**	Disables the distribution.
Step 3	switch(config)# **device-alias distribute**	Enables the distribution (default).

Viewing the Status of Device Alias Distribution

To display the status of device alias distribution, use the **show device-alias status** command. The following example shows the device alias display when distribution is enabled:

```
switch# show device-alias status
Fabric Distribution: Enabled <-------------------------------Distribution is enabled

Database:-Device Aliases 24

Locked By:-User "Test" SWWN 20:00:00:0c:cf:f4:02:83<-Lock holder's user name and switch ID

Pending Database:- Device Aliases 24

Status of the last CFS operation issued from this switch:

==========================================================

Operation: Enable Fabric Distribution

Status: Success
```

The following example shows the device alias display when distribution is disabled:

```
switch# show device-alias status
Fabric Distribution: Disabled

Database:- Device Aliases 24
```

```
Status of the last CFS operation issued from this switch:

=========================================================

Operation: Disable Fabric Distribution

Status: Success
```

About Legacy Zone Alias Configuration

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

- Each zone alias has only one member.
- The member type is pWWN.

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

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

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

Importing a Zone Alias

To import the zone alias for a specific VSAN, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **device-alias import fcalias vsan** *vlan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **device-alias import fcalias vsan** *vlan-id*	Imports the fcalias information for the specified VSAN.

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.
- Verify that the combined number of device aliases in both databases does not exceed 8K (8191 device aliases) in fabrics running Cisco MDS SAN-OS Release 3.0 (x) and earlier, and 20K in fabrics running Cisco MDS SAN-OS Release 3.1(x) and later.

If the combined number of device entries in both databases exceeds the supported configuration limit, then the merge will fail. For example, if database *N* has 6000 device aliases and database *M* has 2192 device aliases, and you are running SAN-OS Release 3.0(x) or earlier, then this merge operation will fail. Merge operations will also fail if there is a device alias mode mismatch.

For additional information, refer to CFS Merge Support in the *Cisco Nexus 5000 Series System Management Configuration Guide.*

Verifying Device Alias Configuration

To display device alias information, perform one of the following tasks:

SUMMARY STEPS

1. switch# **show zoneset** [**active**]
2. switch# **show device-alias database** [**pending** | **pending-diffs**]
3. switch# **show device-alias** {**pwwn** *pwwn-id* | **name** *device-name* } [**pending**]
4. switch# **show flogi database** [**pending**]
5. switch# **show fcns database** [**pending**]

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **show zoneset** [**active**]	Displays the device aliases in the zone set information.
Step 2	switch# **show device-alias database** [**pending** | **pending-diffs**]	Displays the device alias database.
Step 3	switch# **show device-alias** {**pwwn** *pwwn-id* | **name** *device-name* } [**pending**]	Displays the device alias information for the specified pwwn or alias.
Step 4	switch# **show flogi database** [**pending**]	Displays device alias information the the flogi database.
Step 5	switch# **show fcns database** [**pending**]	Displays device alias information the the fcns database.

Examples of Verifying Device Alias Configuration

The following example shows how to display device alias information in the zone set:

```
switch# show zoneset
zoneset name s1 vsan 1
  zone name z1 vsan 1
    pwwn 21:01:00:e0:8b:2e:80:93 [x] <---------------Device alias displayed for each pWWN.
    pwwn 21:00:00:20:37:39:ab:5f [y]
  zone name z2 vsan 1
    pwwn 21:00:00:e0:8b:0b:66:56 [SampleName]
    pwwn 21:00:00:20:37:39:ac:0d [z]
```

The following example shows how to display pending changes in the device alias database:

```
switch# show device-alias database pending
```

The following example shows how to display a specific pWWN in the device alias database:

```
switch# show device-alias pwwn 21:01:00:e0:8b:2e:80:93 pending
```

The following example shows how to display the difference between the pending and effective device alias databases:

```
switch# show device-alias database pending-diff
- device-alias name Doc pwwn 21:01:02:03:00:01:01:01
+ device-alias name SampleName pwwn 21:00:00:e0:8b:0b:66:56
```

Where available, device aliases are displayed regardless of a member being configured using a **device-alias** command or a zone-specific **member pwwn** command.

Default Device Alias Settings

The following table lists the default settings for device alias parameters.

Table 26: 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.

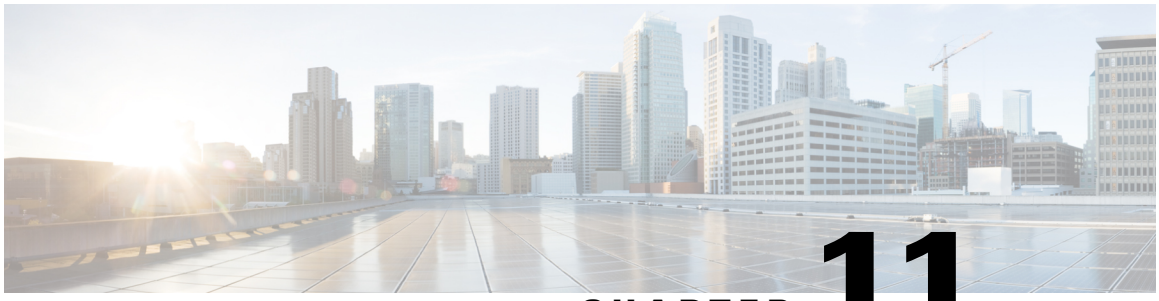

CHAPTER 11

Configuring Fibre Channel Routing Services and Protocols

This chapter contains the following sections:

Configuring Fibre Channel Routing Services and Protocols

Fabric Shortest Path First (FSPF) is the standard path selection protocol used by Fibre Channel fabrics. The FSPF feature is enabled by default on the E mode and TE mode Fibre Channel interfaces on Cisco Nexus 5000 Series switches. Except in configurations that require special consideration, you do not need to configure any FSPF services. FSPF automatically calculates the best path between any two switches in a fabric. FSPF provides the following capabilities:

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

Information About FSPF

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

- Supports multipath routing.
- Bases path status on a link state protocol.
- Routes hop by hop, based only on the domain ID.
- Runs only on E ports or TE ports and provides a loop free topology.
- Runs on a per VSAN basis. Connectivity in a given VSAN in a fabric is guaranteed only for the switches configured in that VSAN.

- Uses a topology database to keep track of the state of the links on all switches in the fabric and associates a cost with each link.
- Guarantees a fast reconvergence time in case of a topology change. Uses the standard Dijkstra algorithm, but there is a static dynamic option for a more robust, efficient, and incremental Dijkstra algorithm. The reconvergence time is fast and efficient as the route computation is done on a per VSAN basis.

Note The FSPF feature can be used on any topology.

FSPF Examples

Fault Tolerant Fabric Example

The following figure depicts a fault tolerant fabric using a partial mesh topology. If a link goes down anywhere in the fabric, any switch can still communicate with all others in the fabric. In the same way, if any switch goes down, the connectivity of the rest of the fabric is preserved.

Figure 26: Fault Tolerant Fabric

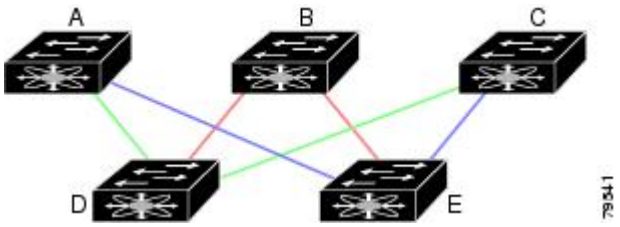

For example, if all links are of equal speed, the FSPF calculates two equal paths from A to C: A-D-C (green) and A-E-C (blue).

Redundant Link Example

To improve on the topology, each connection between any pair of switches can be replicated; two or more links can be present between a pair of switches. The following figure shows this arrangement. Because switches in the Cisco Nexus 5000 Series support SAN port channels, each pair of physical links can appear to the FSPF protocol as one single logical link.

By bundling pairs of physical links, FSPF efficiency is considerably improved by the reduced database size and the frequency of link updates. Once physical links are aggregated, failures are not attached to a single link but to the entire SAN port channel. This configuration also improves the resiliency of the network. The

failure of a link in a SAN port channel does not trigger a route change, which reduces the risks of routing loops, traffic loss, or fabric downtime for route reconfiguration.

Figure 27: Fault Tolerant Fabric with Redundant Links

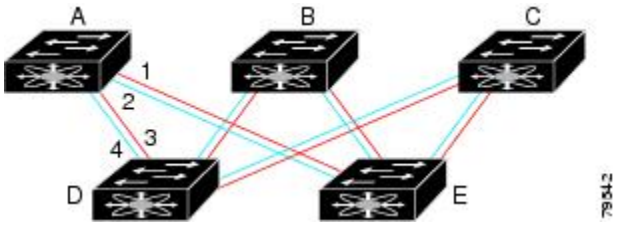

For example, if all links are of equal speed and no SAN port channels exist, the FSPF calculates four equal paths from A to C: A1-E-C, A2-E-C, A3-D-C, and A4-D-C. If SAN port channels exist, these paths are reduced to two.

FSPF Global Configuration

By default, FSPF is enabled on switches in the Cisco Nexus 5000 Series .

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

Note FSPF is enabled by default. Generally, you do not need to configure these advanced features.

Caution The default for the backbone region is 0 (zero). You do not need to change this setting unless your region is different from the default. If you are operating with other vendors using the backbone region, you can change this default to be compatible with those settings.

About SPF Computational Hold Times

The SPF computational hold time sets the minimum time between two consecutive SPF computations on the VSAN. Setting this to a small value means that FSPF reacts faster to any fabric changes by recomputing paths on the VSAN. A small SPF computational hold time uses more switch CPU time.

About Link State Records

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

The following table displays the default settings for switch responses.

Table 27: LSR Default Settings

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

The LSR minimum arrival time is the period between receiving LSR updates on this VSAN. Any LSR updates that arrive before the LSR minimum arrival time are discarded.

The LSR minimum interval time is the frequency at which this switch sends LSR updates on a VSAN.

Configuring FSPF on a VSAN

To configure an FSPF feature for the entire VSAN, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fspf config vsan** *vsan-id*
3. switch-config-(fspf-config)# **spf static**
4. switch-config-(fspf-config)# **spf hold-time** *value*
5. switch-config-(fspf-config)# **region** *region-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fspf config vsan** *vsan-id*	Enters FSPF global configuration mode for the specified VSAN.
Step 3	switch-config-(fspf-config)# **spf static**	Forces static SPF computation for the dynamic (default) incremental VSAN.
Step 4	switch-config-(fspf-config)# **spf hold-time** *value*	Configures the hold time between two route computations in milliseconds (msec) for the entire VSAN. The default value is 0. **Note** If the specified time is shorter, the routing is faster. However, the processor consumption increases accordingly.

	Command or Action	Purpose
Step 5	switch-config-(fspf-config)# **region** *region-id*	Configures the autonomous region for this VSAN and specifies the region ID.

Resetting FSPF to the Default Configuration

To return the FSPF VSAN global configuration to its factory default, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **no fspf config vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **no fspf config vsan** *vsan-id*	Deletes the FSPF configuration for the specified VSAN.

Enabling or Disabling FSPF

To enable or disable FSPF routing protocols, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fspf enable vsan** *vsan-id*
3. switch(config)# **no fspf enable vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fspf enable vsan** *vsan-id*	Enables the FSPF routing protocol in the specified VSAN.
Step 3	switch(config)# **no fspf enable vsan** *vsan-id*	Disables the FSPF routing protocol in the specified VSAN.

Clearing FSPF Counters for the VSAN

To clear the FSPF statistics counters for the entire VSAN, perform this task:

SUMMARY STEPS

1. switch# **clear fspf counters vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **clear fspf counters vsan** *vsan-id*	Clears the FSPF statistics counters for the specified VSAN. If an interface reference is not specified, all counters are cleared.

FSPF Interface Configuration

Several FSPF commands are available on a per-interface basis. These configuration procedures apply to an interface in a specific VSAN.

About FSPF Link Cost

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

Configuring FSPF Link Cost

To configure FSPF link cost, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface fc** *slot*/*port*
3. switch(config-if)# **fspf cost** *value* **vsan** *vsan-id*

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, or if already configured, enters configuration mode for the specified interface.
Step 3	switch(config-if)# **fspf cost** *value* **vsan** *vsan-id*	Configures the cost for the selected interface in the specified VSAN.

About Hello Time Intervals

You can set the FSPF Hello time interval to specify the interval between the periodic hello messages sent to verify the health of the link. The integer value can range from 1 to 65,535 seconds.

Note This value must be the same in the ports at both ends of the ISL.

Configuring Hello Time Intervals

To configure the FSPF Hello time interval, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface fc** *slot*/*port*
3. switch(config-if)# **fspf hello-interval** *value* **vsan** *vsan-id*

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, or if already configured, enters configuration mode for the specified interface.
Step 3	switch(config-if)# **fspf hello-interval** *value* **vsan** *vsan-id*	Specifies the hello message interval to verify the health of the link in VSAN 175. The default is 20 seconds.

About Dead Time Intervals

You can set the FSPF dead time interval to specify the maximum interval for which a hello message must be received before the neighbor is considered lost and removed from the database. The integer value can range from 1 to 65,535 seconds.

Note This value must be the same in the ports at both ends of the ISL.

Caution An error is reported at the command prompt if the configured dead time interval is less than the hello time interval.

Configuring Dead Time Intervals

To configure the FSPF dead time interval, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface fc** *slot*/*port*
3. switch(config-if)# **fspf dead-interval** *value* **vsan** *vsan-id*

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, or if already configured, enters configuration mode for the specified interface.
Step 3	switch(config-if)# **fspf dead-interval** *value* **vsan** *vsan-id*	Specifies the maximum interval for the specified VSAN before which a hello message must be received on the selected interface before the neighbor is considered lost. The default is 80 seconds.

About Retransmitting Intervals

You can specify the time after which an unacknowledged link state update should be transmitted on the interface. The integer value to specify retransmit intervals can range from 1 to 65,535 seconds.

Note This value must be the same on the switches on both ends of the interface.

Configuring Retransmitting Intervals

To configure the FSPF retransmit time interval, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface fc** *slot*/*port*
3. switch(config-if)# **fspf retransmit-interval** *value* **vsan** *vsan-id*

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, or if already configured, enters configuration mode for the specified interface.
Step 3	switch(config-if)# **fspf retransmit-interval** *value* **vsan** *vsan-id*	Specifies the retransmit time interval for unacknowledged link state updates in the specified VSAN. The default is 5 seconds.

About Disabling FSPF for Specific Interfaces

You can disable the FSPF protocol for selected interfaces. By default, FSPF is enabled on all E ports and TE ports. This default can be disabled by setting the interface as passive.

Note FSPF must be enabled at both ends of the interface for the protocol to work.

Disabling FSPF for Specific Interfaces

You can disable the FSPF protocol for selected interfaces. By default, FSPF is enabled on all E ports and TE ports. This default can be disabled by setting the interface as passive.

To disable FSPF for a specific interface, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface fc** *slot*/*port*
3. switch(config-if)# **fspf passive vsan** *vsan-id*
4. switch(config-if)# **no fspf passive vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **interface fc** *slot*/*port*	Configures a specified interface, or if already configured, enters configuration mode for the specified interface.
Step 3	switch(config-if)# **fspf passive vsan** *vsan-id*	Disables the FSPF protocol for the specified interface in the specified VSAN.
Step 4	switch(config-if)# **no fspf passive vsan** *vsan-id*	Reenables the FSPF protocol for the specified interface in the specified VSAN.

Clearing FSPF Counters for an Interface

To clear the FSPF statistics counters for an interface, perform this task:

SUMMARY STEPS

1. switch# **clear fspf counters vsan** *vsan-id* **interface fc** *slot*/*port*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **clear fspf counters vsan** *vsan-id* **interface fc** *slot*/*port*	Clears the FSPF statistics counters for the specified interface in the specified VSAN.

FSPF Routes

FSPF routes traffic across the fabric, based on entries in the FSPF database. These routes can be learned dynamically, or configured statically.

About Fibre Channel Routes

Each port implements forwarding logic, which forwards frames based on its FC ID. Using the FC ID for the specified interface and domain, you can configure the specified route (for example, FC ID 111211 and domain ID 3) in the switch with domain ID 1 (see the following figure).

Figure 28: Fibre Channel Routes

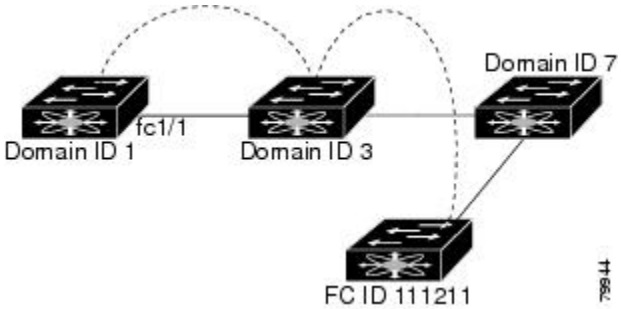

Configuring Fibre Channel Routes

To configure a Fibre Channel route, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)#**fcroute** *fcid* **interface fc** *slot*/*port* **domain** *domain-id* **vsan** *vsan-id*
3. switch(config)#**fcroute** *fcid* **interface san-port-channel** *port* **domain** *domain-id* **vsan** *vsan-id*
4. switch(config)# **fcroute** *fcid* **interface fc** *slot*/*port* **domain** *domain-id* **metric** *value* **vsan** *vsan-id*
5. switch(config)#**fcroute** *fcid* **interface fc** *slot*/*port* **domain** *domain-id* **metric** *value* **remote vsan** *vsan-id*
6. switch(config)#**fcroute** *fcid netmask* **interface fc** *slot*/*port* **domain** *domain-id* **vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)#**fcroute** *fcid* **interface fc** *slot*/*port* **domain** *domain-id* **vsan** *vsan-id*	Configures the route for the specified Fibre Channel interface and domain. In this example, the specified interface is assigned an FC ID and a domain ID to the next hop switch.
Step 3	switch(config)#**fcroute** *fcid* **interface san-port-channel** *port* **domain** *domain-id* **vsan** *vsan-id*	Configures the route for the specified SAN port channel interface and domain. In this example, interface san-port-channel 1 is assigned an FC ID (0x111211) and a domain ID to the next hop switch.
Step 4	switch(config)# **fcroute** *fcid* **interface fc** *slot*/*port* **domain** *domain-id* **metric** *value* **vsan** *vsan-id*	Configures the static route for a specific FC ID and next hop domain ID and also assigns the cost of the route.

	Command or Action	Purpose
		If the remote destination option is not specified, the default is direct.
Step 5	switch(config)#**fcroute** *fcid* **interface fc** *slot*/*port* **domain** *domain-id* **metric** *value* **remote vsan** *vsan-id*	Adds a static route to the RIB. If this is an active route and the FIBFIB = Forwarding Information Base records are free, it is also added to the FIB. If the cost (metric) of the route is not specified, the default is 10.
Step 6	switch(config)#**fcroute** *fcid netmask* **interface fc** *slot*/*port* **domain** *domain-id* **vsan** *vsan-id*	Configures the netmask for the specified route the in interface (or SAN port channel). You can specify one of three routes: 0xff0000 matches only the domain, 0xffff00 matches the domain and the area, 0xffffff matches the domain, area, and port.

In-Order Delivery

In-order delivery (IOD) of data frames guarantees frame delivery to a destination in the same order that they were sent by the originator.

Some Fibre Channel protocols or applications cannot handle out-of-order frame delivery. In these cases, switches in the Cisco Nexus 5000 Series preserve frame ordering in the frame flow. The source ID (SID), destination ID (DID), and optionally the originator exchange ID (OX ID) identify the flow of the frame.

On a switch with IOD enabled, all frames received by a specific ingress port and destined to a certain egress port are always delivered in the same order in which they were received.

Use IOD only if your environment cannot support out-of-order frame delivery.

If you enable the in-order delivery feature, the graceful shutdown feature is not implemented.

About Reordering Network Frames

When you experience a route change in the network, the new selected path may be faster or less congested than the old route.

Figure 29: Route Change Delivery

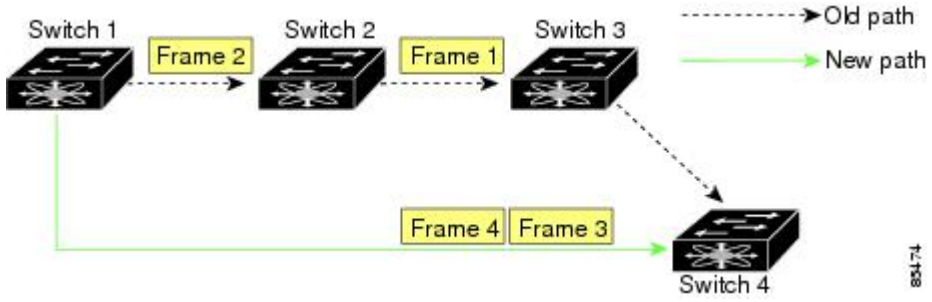

In the figure above, the new path from Switch 1 to Switch 4 is faster. In this scenario, Frame 3 and Frame 4 may be delivered before Frame 1 and Frame 2.

If the in-order guarantee feature is enabled, the frames within the network are delivered as follows:

- Frames in the network are delivered in the order in which they are transmitted.
- Frames that cannot be delivered in order within the network latency drop period are dropped inside the network.

About Reordering SAN Port Channel Frames

When a link change occurs in a SAN port channel, the frames for the same exchange or the same flow can switch from one path to another faster path.

Figure 30: Link Congestion Delivery

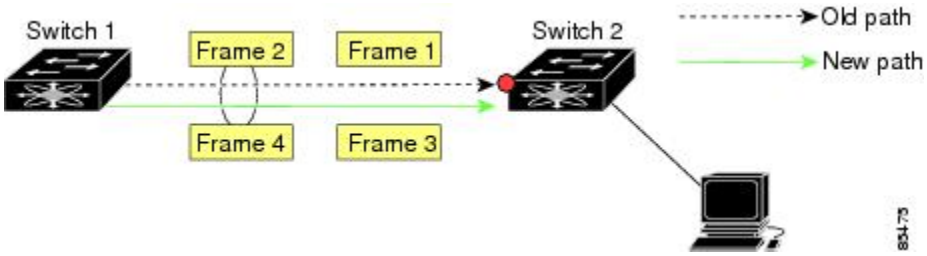

In the figure above, the port of the old path (red dot) is congested. In this scenario, Frame 3 and Frame 4 can be delivered before Frame 1 and Frame 2.

When the in-order delivery feature is enabled and a port channel link change occurs, the frames crossing the SAN port channel are delivered as follows:

- Frames using the old path are delivered before new frames are accepted.
- The new frames are delivered through the new path after the network latency drop period has elapsed and all old frames are flushed.

Frames that cannot be delivered in order through the old path within the network latency drop period are dropped.

Related Topics

Configuring the Drop Latency Time, on page 157

About Enabling In-Order Delivery

You can enable the in-order delivery feature for a specific VSAN or for the entire switch. By default, in-order delivery is disabled on switches.

We recommend that you only enable this feature when devices that cannot handle any out-of-order frames are present in the switch. Load-balancing algorithms in the switch ensure that frames are delivered in order during normal fabric operation. The load-balancing algorithms based on source FC ID, destination FC ID, and exchange ID are enforced in hardware without any performance degradation. However, if the fabric encounters a failure and the in-order delivery feature is enabled, the recovery will be delayed because of an

intentional pausing of fabric forwarding to purge the fabric of resident frames that could potentially be forwarded out-of-order.

Enabling In-Order Delivery Globally

To ensure that the in-order delivery parameters are uniform across all VSANs on the switch, enable in-order delivery globally.

Only enable in-order delivery globally if this is a requirement across your entire fabric. Otherwise, enable IOD only for the VSANs that require this feature.

To enable in-order delivery for the switch, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **in-order-guarantee**
3. switch(config)# **no in-order-guarantee**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **in-order-guarantee**	Enables in-order delivery in the switch.
Step 3	switch(config)# **no in-order-guarantee**	Reverts the switch to the factory defaults and disables the in-order delivery feature.

Enabling In-Order Delivery for a VSAN

When you create a VSAN, that VSAN automatically inherits the global in-order guarantee value. You can override this global value by enabling or disabling in-order guarantee for the new VSAN.

To use the lowest domain switch for the multicast tree computation, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **in-order-guarantee vsan** *vsan-id*
3. switch(config)# **no in-order-guarantee vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **in-order-guarantee vsan** *vsan-id*	Enables in-order delivery in the specified VSAN.
Step 3	switch(config)# **no in-order-guarantee vsan** *vsan-id*	Reverts the switch to the factory defaults and disables the in-order delivery feature in the specified VSAN.

Displaying the In-Order Delivery Status

Use the **show in-order-guarantee** command to display the present configuration status:

```
switch# show in-order-guarantee
global inorder delivery configuration:guaranteed
VSAN specific settings
vsan 1 inorder delivery:guaranteed
vsan 101 inorder delivery:not guaranteed
vsan 1000 inorder delivery:guaranteed
vsan 1001 inorder delivery:guaranteed
vsan 1682 inorder delivery:guaranteed
vsan 2001 inorder delivery:guaranteed
vsan 2009 inorder delivery:guaranteed
vsan 2456 inorder delivery:guaranteed
vsan 3277 inorder delivery:guaranteed
vsan 3451 inorder delivery:guaranteed
vsan 3452 inorder delivery:guaranteed
```

Configuring the Drop Latency Time

You can change the default latency time for a network, a specified VSAN in a network, or for the entire switch.

To configure the network and the switch drop latency time, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcdroplatency network** *value*
3. switch(config)# **fcdroplatency network** *value* **vsan** *vsan-id*
4. switch(config)# **no fcdroplatency network** *value*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fcdroplatency network** *value*	Configures network drop latency time for the network. The valid range is 0 to 60000 msec. The default is 2000 msec.

	Command or Action	Purpose
		Note The network drop latency must be computed as the sum of all switch latencies of the longest path in the network.
Step 3	switch(config)# **fcdroplatency network** *value* **vsan** *vsan-id*	Configures network drop latency time for the specified VSAN.
Step 4	switch(config)# **no fcdroplatency network** *value*	Removes the current fcdroplatecy network configuration and reverts the switch to the factory defaults.

Displaying Latency Information

You can view the configured latency parameters using the **show fcdroplatency** command, The following example shows how to display network latency information:

```
switch# show fcdroplatency
switch latency value:500 milliseconds
global network latency value:2000 milliseconds
VSAN specific network latency settings
vsan 1 network latency:5000 milliseconds
vsan 2 network latency:2000 milliseconds
vsan 103 network latency:2000 milliseconds
vsan 460 network latency:500 milliseconds
```

Flow Statistics Configuration

Flow statistics count the ingress traffic in the aggregated statistics table. You can collect two kinds of statistics:

- Aggregated flow statistics to count the traffic for a VSAN.
- Flow statistics to count the traffic for a source and destination ID pair in a VSAN.

About Flow Statistics

If you enable flow counters, you can enable a maximum of 1000 entries for aggregate flow and flow statistics. Be sure to assign an unused flow index for each new flow. The number space for flow index is shared between the aggregate flow statistics and the flow statistics.

Counting Aggregated Flow Statistics

To count the aggregated flow statistics for a VSAN, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcflow stats aggregated index** *value* **vsan** *vsan-id*
3. switch(config)# **no fcflow stats aggregated index** *value* **vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fcflow stats aggregated index** *value* **vsan** *vsan-id*	Enables the aggregated flow counter.
Step 3	switch(config)# **no fcflow stats aggregated index** *value* **vsan** *vsan-id*	Disables the aggregated flow counter.

Counting Individual Flow Statistics

To count the flow statistics for a source and destination FC ID in a VSAN, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcflow stats index** *value dest-fcid source-fcid netmask* **vsan** *vsan-id*
3. switch(config)# **no fcflow stats aggregated index** *value* **vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fcflow stats index** *value dest-fcid source-fcid netmask* **vsan** *vsan-id*	Enables the flow counter. **Note** The source ID and the destination ID are specified in FC ID hex format (for example, 0x123aff). The mask can be one of 0xff0000 or 0xffffff.
Step 3	switch(config)# **no fcflow stats aggregated index** *value* **vsan** *vsan-id*	Disables the flow counter.

Clearing FIB Statistics

Use the **clear fcflow stats** command to clear the aggregated flow counter. The following example clears the aggregated flow counters:

```
switch# clear fcflow stats aggregated index 1
```

The following example clears the flow counters for source and destination FC IDs:

```
switch# clear fcflow stats index 1
```

Displaying Flow Statistics

Use the **show fcflow stats** commands to view flow statistics. The following example displays the aggregated flow summary:

```
switch# show fcflow stats aggregated
Idx       VSAN     frames
---------- ---------- ----------
         6          1      42871
```

The following example displays the flow statistics:

```
switch# show fcflow stats
```

The following example displays flow index usage:

```
switch# show fcflow stats usage
2 flows configured
Configured flows : 3,7
```

The following example shows how to display global FSPF information for a specific VSAN:

```
switch# show fspf vsan 1
```

The following example shows how to display a summary of the FSPF database for a specified VSAN. If no additional parameters are specified, all LSRs in the database are displayed:

```
switch# show fspf database vsan 1
```

The following example shows how to display FSPF interface information:

```
switch# show fspf vsan 1 interface fc2/1
```

Default FSPF Settings

The following table lists the default settings for FSPF features.

Table 28: Default FSPF Settings

Parameters	Default
FSPF	Enabled on all E ports and TE ports.
SPF computation	Dynamic.
SPF hold time	0.
Backbone region	0.
Acknowledgment interval (RxmtInterval)	5 seconds.
Refresh time (LSRefreshTime)	30 minutes.
Maximum age (MaxAge)	60 minutes.
Hello interval	20 seconds.
Dead interval	80 seconds.
Distribution tree information	Derived from the principal switch (root node).

Parameters	Default
Routing table	FSPF stores up to 16 equal cost paths to a given destination.
Load balancing	Based on destination ID and source ID on different, equal cost paths.
In-order delivery	Disabled.
Drop latency	Disabled.
Static route cost	If the cost (metric) of the route is not specified, the default is 10.
Remote destination switch	If the remote destination switch is not specified, the default is direct.
Multicast routing	Uses the principal switch to compute the multicast tree.

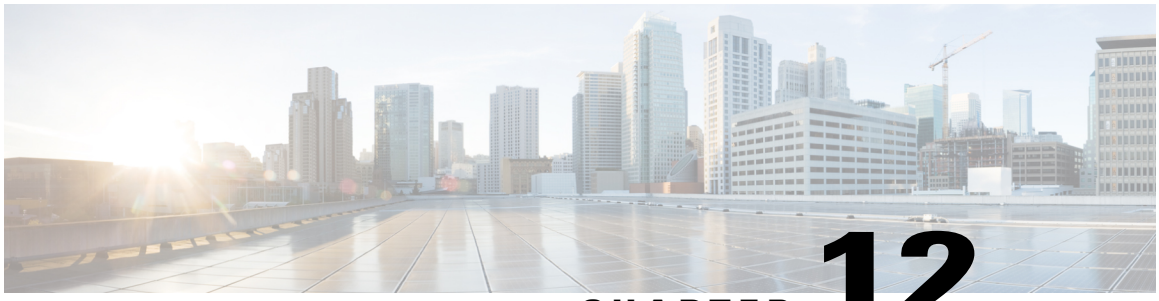

CHAPTER 12

Managing FLOGI, Name Server, FDMI, and RSCN Databases

This chapter contains the following sections:

Managing FLOGI, Name Server, FDMI, and RSCN Databases

Information About 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 following 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.
```

The following 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.
```

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

```
switch# show flogi database vsan 1
```

Name Server Proxy

The name server functionality maintains a database containing the attributes for all hosts and storage devices in each VSAN. Name servers allow a database entry to be modified by a device that originally registered the information.

The proxy feature is useful when you need to modify (update or delete) the contents of a database entry that was previously registered by a different device.

About Registering Name Server Proxies

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

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

Registering Name Server Proxies

To register the name server proxy, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)#**fcns proxy-port** *wwn-id* **vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)#**fcns proxy-port** *wwn-id* **vsan** *vsan-id*	Configures a proxy port for the specified VSAN.

About Rejecting Duplicate pWWNs

You can prevent malicious or accidental log in using another device's pWWN by enabling the reject-duplicate-pwwn option. If you disable this option, these pWWNs are allowed to log in to the fabric and replace the first device in the name server database.

Rejecting Duplicate pWWNs

To reject duplicate pWWNs, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcns reject-duplicate-pwwn vsan** *vsan-id*
3. switch(config)# **no fcns reject-duplicate-pwwn vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fcns reject-duplicate-pwwn vsan** *vsan-id*	Logs out devices when they log into the fabric if the pWWNs already exist.
Step 3	switch(config)# **no fcns reject-duplicate-pwwn vsan** *vsan-id*	Overwrites the first device's entry in the name server database with the new device having the same pWWN (default).

About Name Server Database Entries

The name server stores name entries for all hosts in the FCNS database. The name server permits an Nx port to register attributes during a PLOGI (to the name server) to obtain attributes of other hosts. These attributes are deregistered when the Nx port logs out either explicitly or implicitly.

In a multiswitch fabric configuration, the name server instances running on each switch shares information in a distributed database. One instance of the name server process runs on each switch.

Displaying Name Server Database Entries

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

```
switch# show fcns database
--------------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0x010000    N     50:06:0b:00:00:10:a7:80                   scsi-fcp fc-gs

0x010001    N     10:00:00:05:30:00:24:63 (Cisco)           ipfc

0x010002    N     50:06:04:82:c3:a0:98:52 (Company 1)       scsi-fcp 250

0x010100    N     21:00:00:e0:8b:02:99:36 (Company A)       scsi-fcp

0x020000    N     21:00:00:e0:8b:08:4b:20 (Company A)
0x020100    N     10:00:00:05:30:00:24:23 (Cisco)           ipfc

0x020200    N     21:01:00:e0:8b:22:99:36 (Company A)       scsi-fcp
```

The following example shows how to display the name server database and statistical information for a specified VSAN:

```
switch# show fcns database vsan 1
VSAN 1:
--------------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0x030001    N     10:00:00:05:30:00:25:a3 (Cisco)         ipfc
0x030101    NL    10:00:00:00:77:99:60:2c (Interphase)
0x030200    N     10:00:00:49:c9:28:c7:01
0xec0001    NL    21:00:00:20:37:a6:be:14 (Seagate)       scsi-fcp
Total number of entries = 4
```

The following example shows how to display the name server database details for all VSANs:

```
switch# show fcns database detail
```

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

```
switch# show fcns statistics
```

FDMI

Cisco Nexus 5000 Series switches provide support for the Fabric-Device Management Interface (FDMI) functionality, as described in the FC-GS-4 standard. FDMI enables management of devices such as Fibre Channel host bus adapters (HBAs) through in-band communications. This addition complements the existing Fibre Channel name server and management server functions.

Using the FDMI functionality, the 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

The following example shows how to display all HBA details for a specified VSAN:

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

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

About RSCN Information

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
```

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.

About the multi-pid Option

If the RSCN multi-pid option is enabled, then RSCNs generated to the registered Nx ports may contain more than one affected port IDs. In this case, zoning rules are applied before putting the multiple affected port IDs together in a single RSCN. By enabling this option, you can reduce the number of RSCNs. For example, 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.

Configuring the multi-pid Option

To configure the **multi-pid** option, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **rscn multi-pid vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **rscn multi-pid vsan** *vsan-id*	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.

To suppress the transmission of these SW-RSCNs over an ISL, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **rscn suppress domain-swrscn vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **rscn suppress domain-swrscn vsan** *vsan-id*	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.

The following 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. Upon time-out, all the events are dequeued and coalesced RSCNs are sent to registered users. The default timer values minimize the number of coalesced RSCNs sent to registered users. Some deployments require smaller event timer values to track changes in the fabric.

Note The RSCN timer value must be the same on all switches in the VSAN.

Note Before performing a downgrade, make sure that you revert the RCSN timer value in your network to the default value. Failure to do so will disable the links across your VSANs and other devices.

To configure the RSCN timer, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **rscn distribute**
3. switch(config)# **rscn event-tov** *timeout* **vsan** *vsan-id*
4. switch(config)# **no rscn event-tov** *timeout* **vsan** *vsan-id*
5. switch(config)# **rscn commit vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **rscn distribute**	Enables RSCN timer configuration distribution.
Step 3	switch(config)# **rscn event-tov** *timeout* **vsan** *vsan-id*	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	switch(config)# **no rscn event-tov** *timeout* **vsan** *vsan-id*	Reverts to the default value (2000 milliseconds for Fibre Channel VSANs).

	Command or Action	Purpose
Step 5	switch(config)# **rscn commit vsan** *vsan-id*	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. The following 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. This means different N-ports in a network can receive RSCNs at different times. Cisco Fabric Services (CFS) infrastructure alleviates this situation by automatically distributing the RSCN timer configuration information to all switches in a fabric. This also reduces the number of SW-RSCNs.

RSCN supports two modes, distributed and nondistributed. In distributed mode, RSCN uses CFS to distribute configuration to all switches in the fabric. In nondistributed mode, only the configuration commands on the local switch are affected.

Note All configuration commands are not distributed. Only the **rscn event-tov tov vsan vsan** command is distributed.

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.

For additional information, refer to Using Cisco Fabric Services in the *Cisco Nexus 5000 Series System Management Configuration Guide*.

Enabling RSCN Timer Configuration Distribution

To enable RSCN timer configuration distribution, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **rscn distribute**
3. switch(config)# **no rscn distribute**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **rscn distribute**	Enables RSCN timer distribution.
Step 3	switch(config)# **no rscn distribute**	Disables (default) RSCN timer distribution.

Locking the Fabric

The first action that modifies the database creates the pending database and locks the feature in the VSAN. Once you lock the fabric, the following situations apply:

- No other user can make any configuration changes to this feature.
- A copy of the configuration database becomes the pending database along with the first active change.

Committing the RSCN Timer Configuration Changes

If you commit the changes made to the active database, the configuration is committed to all the switches in the fabric. On a successful commit, the configuration change is applied throughout the fabric and the lock is released.

To commit RSCN timer configuration changes, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **rscn commit vsan** *timeout*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **rscn commit vsan** *timeout*	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.

To discard RSCN timer configuration changes, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **rscn abort vsan** *timeout*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **rscn abort vsan** *timeout*	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 DPVM session, use the **clear rscn session vsan** command in EXEC mode. The following example shows how to clear the RSCN session for VSAN 10:

```
switch# clear rscn session vsan 10
```

Displaying RSCN Configuration Distribution Information

The following 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.

The following 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
rscn event-tov 2000 ms vsan 1
rscn event-tov 2000 ms vsan 2
rscn event-tov 300 ms vsan 10
```

The following example shows how to display the difference between pending and active configurations:

```
switch# show rscn pending-diff vsan 10
- rscn event-tov 2000 ms vsan 10
+ rscn event-tov 300 ms vsan 10
```

Default RSCN Settings

The following table lists the default settings for RSCN.

Table 29: Default RSCN Settings

Parameters	Default
RSCN timer value	2000 milliseconds for Fibre Channel VSANs
RSCN timer configuration distribution	Disabled

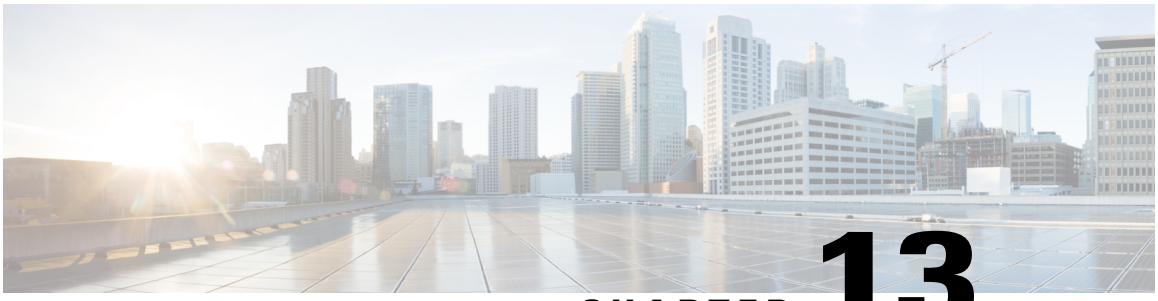

CHAPTER 13

Discovering SCSI Targets

This chapter contains the following sections:

Discovering SCSI Targets

Information About SCSI LUN Discovery

Small Computer System Interface (SCSI) targets include disks, tapes, and other storage devices. These targets do not register logical unit numbers (LUNs) with the name server.

The name server requires LUN information for the following reasons:

- To display LUN storage device information so that a Network Management System (NMS) can access this information.
- To report device capacity, serial number, and device ID information.
- To register the initiator and target features with the name server.

The SCSI LUN discovery feature uses the local domain controller Fibre Channel address. It uses the local domain controller as the source FC ID, and performs SCSI INQUIRY, REPORT LUNS, and READ CAPACITY commands on SCSI devices.

The SCSI LUN discovery feature is initiated on demand, through CLI or SNMP. This information is also synchronized with neighboring switches, if those switches belong to the Cisco Nexus 5000 Series.

About Starting SCSI LUN Discovery

SCSI LUN discovery is done on demand.

Only Nx ports that are present in the name server database and that are registered as FC4 Type = SCSI_FCP are discovered.

Starting SCSI LUN Discovery

To start SCSI LUN discovery, perform this task:

SUMMARY STEPS

1. switch# **discover scsi-target** {**custom-list** | **local** | **remote** | **vsan** *vsan-id* **fcid** *fc-id*} **os** {**aix** | **hpux** | **linux** | **solaris** | **windows**} [**lun** | **target**]

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **discover scsi-target** {**custom-list** | **local** | **remote** | **vsan** *vsan-id* **fcid** *fc-id*} **os** {**aix** | **hpux** | **linux** | **solaris** | **windows**} [**lun** | **target**]	Discovers SCSI targets for the specified operating system (OS).

Examples of Starting SCSI LUN Discovery

The following example discovers local SCSI targets for all operating systems (OSs):

```
switch# discover scsi-target local os all
discovery started
```

The following example discovers remote SCSI targets assigned to the AIX OS:

```
switch# discover scsi-target remote os aix
discovery started
```

The following example discovers SCSI targets for VSAN 1 and FC ID 0x9c03d6:

```
switch# discover scsi-target vsan 1 fcid 0x9c03d6
discover scsi-target vsan 1 fcid 0x9c03d6
VSAN:    1 FCID: 0x9c03d6 PWWN: 00:00:00:00:00:00:00:00
  PRLI RSP: 0x01 SPARM: 0x0012
  SCSI TYPE: 0 NLUNS: 1
  Vendor: Company 4 Model: ST318203FC    Rev: 0004
  Other: 00:00:02:32:8b:00:50:0a
```

The following example discovers SCSI targets from the customized list assigned to the Linux OS:

```
switch# discover scsi-target custom-list os linux
discovery started
```

About Initiating Customized Discovery

Customized discovery consists of a list of VSAN and domain pairs that are selectively configured to initiate a discovery. Use the custom-list option to initiate this discovery. The domain ID is a number from 0 to 255 in decimal or a number from 0x0 to 0xFF in hex.

Initiating Customized Discovery

To initiate a customized discovery, perform this task:

SUMMARY STEPS

1. switch# **discover custom-list add vsan** *vsan-id* **domain** *domain-id*
2. switch# **discover custom-list delete vsan** *vsan-id* **domain** *domain-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **discover custom-list add vsan** *vsan-id* **domain** *domain-id*	Adds the specified entry to the custom list.
Step 2	switch# **discover custom-list delete vsan** *vsan-id* **domain** *domain-id*	Deletes the specified domain ID from the custom list.

Displaying SCSI LUN Information

Use the **show scsi-target**and **show fcns database**commands to display the results of the discovery.

The following example displays the discovered targets:

```
switch# show scsi-target status
discovery completed
```

Note This command takes several minutes to complete, especially if the fabric is large or if several devices are slow to respond.

The following example displays the FCNS database:

```
switch# show fcns database
```

The following example displays the SCSI target disks:

```
switch# show scsi-target disk
```

The following example displays the discovered LUNs on all operating systems:

```
switch# show scsi-target lun os all
```

The following example displays the port WWN that is assigned to each operating system (Windows, AIX, Solaris, Linux, or HPUX):

```
switch# show scsi-target pwwn
```

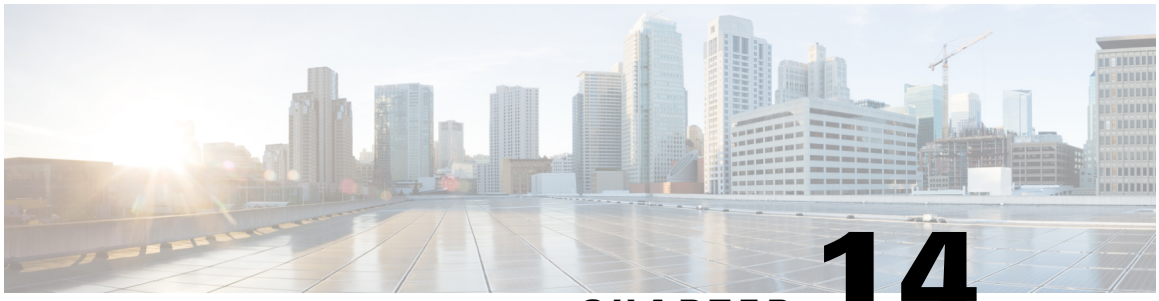

CHAPTER 14

Advanced Fibre Channel Features and Concepts

This chapter contains the following sections:

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. The default is 5,000 milliseconds.
- Error detect TOV (E_D_TOV)—The valid range is from 1,000 to 10,000 milliseconds. The default is 2,000 milliseconds. This value is matched with the other end during port initialization.
- Resource allocation TOV (R_A_TOV)—The valid range is from 5,000 to 10,000 milliseconds. The default is 10,000 milliseconds. This value is matched with the other end during port initialization.

Note The fabric stability TOV (F_S_TOV) constant cannot be configured.

Timer Configuration Across All VSANs

You can modify Fibre Channel protocol related timer values for the switch.

Caution The D_S_TOV, E_D_TOV, and R_A_ TOV values cannot be globally changed unless all VSANs in the switch are suspended.

Note If a VSAN is not specified when you change the timer value, the changed value is applied to all VSANs in the switch.

To configure Fibre Channel timers across all VSANs, perform this task:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **fctimer R_A_TOV** *timeout*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters configuration mode.
Step 2	switch(config)# **fctimer R_A_TOV** *timeout*	Configures the R_A_TOV timeout value for all VSANs. The units 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.

To configure per-VSAN Fibre Channel timers, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config#)# **fctimer D_S_TOV** *timeout* **vsan** *vsan-id*

DETAILED STEPS

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

	Command or Action	Purpose
Step 2	switch(config#)# **fctimer D_S_TOV** *timeout* **vsan** *vsan-id*	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.

The following example configures 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
```

About 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.

For additional information, refer to Using Cisco Fabric Services in the *Cisco Nexus 5000 Series System Management Configuration Guide*.

Enabling or Disabling fctimer Distribution

To enable or disable fctimer fabric distribution, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fctimer distribute**
3. switch(config)# **no fctimer distribute**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fctimer distribute**	Enables fctimer configuration distribution to all switches in the fabric. Acquires a fabric lock and stores all future configuration changes in the pending database.

	Command or Action	Purpose
Step 3	switch(config)# **no fctimer distribute**	Disables (default) fctimer configuration distribution to all switches in the fabric.

Committing fctimer Changes

When you commit the fctimer configuration changes, the effective database is overwritten by the configuration changes in the pending database and all the switches in the fabric receive the same configuration. When you commit the fctimer configuration changes without implementing the session feature, the fctimer configurations are distributed to all the switches in the physical fabric.

To commit the fctimer configuration changes, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fctimer commit**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fctimer commit**	Distributes the fctimer configuration changes to all switches in the fabric and releases the lock. Overwrites the effective database with the changes made to the pending database.

Discarding fctimer Changes

After making the configuration changes, you can choose to discard the changes by discarding the changes instead of committing them. In either case, the lock is released.

To discard the fctimer configuration changes, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fctimer abort**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fctimer abort**	Discards the fctimer configuration changes in the pending database and releases the fabric lock.

Fabric Lock Override

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

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.

For additional information, refer to CFS Merge Support in the *Cisco Nexus 5000 Series System Management Configuration Guide*.

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 Nexus 5000 Series switches support three network address authority (NAA) address formats. (see the following table).

Table 30: 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 WWN Information

Use the **show wwn** commands to display the status of the WWN configuration. The following example displays 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
```

The following example displays 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
```

The following example displays the WWN for a specific switch:

```
switch# show wwn switch
Switch WWN is 20:00:ac:16:5e:52:00:00
```

Link Initialization WWN Usage

Exchange Link Protocol (ELP) and Exchange Fabric Protocol (EFP) use WWNs during link initialization. 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

To allocate secondary MAC addresses, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **wwn secondary-mac** *wwn-id* **range** *value*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **wwn secondary-mac** *wwn-id* **range** *value*	Configures the secondary MAC address. This command cannot be undone.

The following 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 Nexus 5000 Series 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 Nexus 5000 Series 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.

To allocate company IDs, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcid-allocation area company-id** *value*
3. switch(config)# **no fcid-allocation area company-id** *value*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fcid-allocation area company-id** *value*	Adds a new company ID to the default list.
Step 3	switch(config)# **no fcid-allocation area company-id** *value*	Deletes a company ID from the default list.

The following example adds a new company ID to the default list:

```
switch(config)# fcid-allocation area company-id 0x003223
```

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.

The following example displays the list of default and configured company IDs:

```
switch# show fcid-allocation area
FCID area allocation company id info:
00:50:2E <--------------- Default entry
00:50:8B
00:60:B0
00:A0:B8
00:E0:69
00:30:AE + <------------- User-added entry
00:32:23 +
00:E0:8B * <------------- Explicitly deleted entry (from the original default list)
Total company ids: 7
+ - Additional user configured company ids.
* - Explicitly deleted company ids from default list.
```

You can implicitly derive the default entries shipped with a specific release by combining the list of Company IDs displayed without any identification with the list of deleted entries.

You can also view or obtain the company IDs in a specific WWN by 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.

The following example displays the company ID for the specified WWN:

```
switch# show fcid-allocation company-id-from-wwn 20:00:00:05:30:00:21:60
Extracted Company ID: 0x000530
```

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.

Note For more information on configuring interoperability for Cisco Nexus 5000 Series switches, see the *Cisco MDS 9000 Family Switch-to-Switch Interoperability Configuration Guide*

About Interop Mode

Cisco NX-OS software supports the following four interop modes:

- Mode 1— Standards-based interop mode that requires all other vendors in the fabric to be in interop mode.
- Mode 2—Brocade native mode (Core PID 0).
- Mode 3—Brocade native mode (Core PID 1).
- Mode 4—McData native mode.

For information about configuring interop modes 2, 3, and 4, see the *Cisco MDS 9000 Family Switch-to-Switch Interoperability Configuration Guide*, available at the following location: http://www.cisco.com/en/US/docs/storage/san_switches/mds9000/interoperability/guide/intopgd.html

The following table lists the changes in switch operation when you enable interoperability mode. These changes are specific to Cisco Nexus 5000 Series switches while in interop mode.

Table 31: Changes in Switch Operation When Interoperability Is Enabled

Switch Feature	Changes if Interoperability Is Enabled
Domain IDs	Some vendors cannot use the full range of 239 domains within a fabric. Domain IDs are restricted to the range 97 to 127, to accommodate McData's nominal restriction to this same range. 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 on a per port or per switch basis.
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.

Switch Feature	Changes if Interoperability Is Enabled
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 Nexus 5000 Series switches if they are part of the same fabric. You must explicitly save the configuration on each Cisco Nexus 5000 Series switch.
Zone propagation	Some vendors do not pass the full zone configuration to other switches, only the active zone set gets passed. Verify that the active zone set or zone configuration has correctly propagated to the other switches in the fabric.
VSAN	Interop mode only affects the specified VSAN. **Note** Interop modes cannot be enabled on FICON-enabled VSANs.
TE ports and SAN port channel s	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 Nexus 5000 Series 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 1

The interop mode1 in Cisco Nexus 5000 Series switches can be enabled disruptively or nondisruptively.

Note Brocade's **msplmgmtdeactivate** command must explicitly be run prior to connecting from a Brocade switch to either Cisco Nexus 5000 Series switches or to McData switches. This command uses Brocade proprietary frames to exchange platform information, which Cisco Nexus 5000 Series switches or McData switches do not understand. Rejecting these frames causes the common E ports to become isolated.

To configure interop mode 1 in any switch in the Cisco Nexus 5000 Series, perform this task:

SUMMARY STEPS

1. Place the VSAN of the E ports that connect to the OEM switch in interoperability mode.
2. Assign a domain ID in the range of 97 (0x61) through 127 (0x7F).
3. Change the Fibre Channel timers (if they have been changed from the system defaults).
4. When making changes to the domain, you may or may not need to restart the domain manager function for the altered VSAN.

DETAILED STEPS

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 1 interop 1
switch(config-vsan-db)# exit
```

Step 2 Assign a domain ID in the range of 97 (0x61) through 127 (0x7F).

Note This is an limitation imposed by the McData switches.

In Cisco Nexus 5000 Series switches, the default is to request an ID from the principal switch. If the preferred option is used, Cisco Nexus 5000 Series switches request a specific ID, but still join the fabric if the principal switch assigns a different ID. If the static option is used, the Cisco Nexus 5000 Series switches do not join the fabric unless the principal switch agrees and assigns the requested ID.

Note When changing the domain ID, the FC IDs assigned to N ports also change.

Step 3 Change the Fibre Channel timers (if they have been changed from the system defaults).

Note The Cisco Nexus 5000 Series, 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 ?
  <5000-100000> R_A_TOV in milliseconds(5000-100000)
```

Step 4 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 1
```

Verifying Interoperating Status

This section highlights the commands used to verify if the fabric is up and running in interoperability mode.

To verify the resulting status of entering the interoperability command in any switch in the Cisco Nexus 5000 Series, perform this task:

SUMMARY STEPS

1. Verify the software version.
2. Verify if the interface states are as required by your configuration.
3. Verify if you are running the desired configuration.
4. Verify if the interoperability mode is active.
5. Verify the domain ID.
6. Verify the local principal switch status.
7. Verify the next hop and destination for the switch.
8. Verify the name server information.

DETAILED STEPS

Step 1 Verify the software version.

Example:

```
switch# show version
Cisco Nexus Operating System (NX-OS) Software

TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2008, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained herein are owned by
other third parties and are used and distributed under license.

Some parts of this software are covered under the GNU Public

License. A copy of the license is available at
```

```
http://www.gnu.org/licenses/gpl.html.

Software
  BIOS:      version 1.2.0
  loader:    version N/A
  kickstart: version 4.0(1a)N1(1)
  system:    version 4.0(1a)N1(1)
  BIOS compile time:       06/19/08
  kickstart image file is: bootflash:/n5000-uk9-kickstart.4.0.1a.N1.latest.bin
  kickstart compile time:  11/25/2008 6:00:00 [11/25/2008 14:17:12]
  system image file is:    bootflash:/n5000-uk9.4.0.1a.N1.latest.bin
  system compile time:     11/25/2008 6:00:00 [11/25/2008 14:59:49]
Hardware
  cisco Nexus5020 Chassis ("40x10GE/Supervisor")
  Intel(R) Celeron(R) M CPU with 2074308 kB of memory.
  Processor Board ID JAB120900PJ
  Device name: switch
  bootflash: 1003520 kB

Kernel uptime is 0 day(s), 1 hour(s), 29 minute(s), 55 second(s)

Last reset at 510130 usecs after Wed Nov 26 18:12:23 2008
  Reason: Reset Requested by CLI command reload
  System version: 4.0(1a)N1(1)
  Service:

plugin
  Core Plugin, Ethernet Plugin
```

Step 2 Verify if the interface states are as required by your configuration.

Example:

```
switch# show interface brief

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

Interface  Vsan   Admin  Admin   Status          SFP    Oper  Oper   Port

                  Mode   Trunk                          Mode  Speed  Channel

                         Mode                                 (Gbps)

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

fc3/1      1      E      on      trunking         swl    TE      2    --

fc3/2      1      auto   on      sfpAbsent        --     --           --

fc3/3      1      E      on      trunking         swl    TE      2    --

fc3/4      1      auto   on      sfpAbsent        --     --           --

fc3/5      1      auto   auto    notConnected     swl    --           --

fc3/6      1      auto   on      sfpAbsent        --     --           --

fc3/7      1      auto   auto    sfpAbsent        --     --           --

fc3/8      1      auto   auto    sfpAbsent        --     --           --
```

Step 3 Verify if you are running the desired configuration.

Example:

```
switch# show running-config

Building Configuration...
```

```
 interface fc2/1
no shutdown
 interface fc2/2
no shutdown
 interface fc2/3
 interface fc2/4
<snip>
interface mgmt0
ip address 6.1.1.96 255.255.255.0
switchport encap default
no shutdown
vsan database
vsan 1 interop
boot system bootflash:/nx5000-system-23e.bin
boot kickstart bootflash:/nx5000-kickstart-23e.bin
callhome
fcdomain domain 100 preferred vsan 1
ip route 6.1.1.0 255.255.255.0 6.1.1.1
ip routing
line console
  databits 5
  speed 110
logging linecard
ssh key rsa 512 force
ssh server enable
switchname switch
username admin password 5 $1$Li8/fBYX$SNc72.xt4nTXpSnR9OUFB/ role network-admin
```

Step 4 Verify if the interoperability mode is active.

Example:

```
switch# show vsan 1
vsan 1 information
         name:VSAN0001  state:active
         interoperability mode:yes <---------------------- verify mode
         loadbalancing:src-id/dst-id/oxid
         operational state:up
```

Step 5 Verify the domain ID.

Example:

```
switch# show fcdomain vsan 1

The local switch is a Subordinated Switch.

Local switch run time information:

        State: Stable

        Local switch WWN:    20:01:00:05:30:00:51:1f

        Running fabric name: 10:00:00:60:69:22:32:91

        Running priority: 128

        Current domain ID: 0x64(100) <---------------verify domain id

Local switch configuration information:

        State: Enabled

        Auto-reconfiguration: Disabled

        Contiguous-allocation: Disabled

        Configured fabric name: 41:6e:64:69:61:6d:6f:21

        Configured priority: 128

        Configured domain ID: 0x64(100) (preferred)

Principal switch run time information:

        Running priority: 2

Interface               Role          RCF-reject

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

fc2/1               Downstream       Disabled

fc2/2               Downstream       Disabled

fc2/4               Upstream         Disabled

----------------    -------------    ------------
```

Step 6 Verify the local principal switch status.

Example:

```
switch# show fcdomain domain-list vsan 1

Number of domains: 5

Domain ID              WWN

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

 0x61(97)    10:00:00:60:69:50:0c:fe

 0x62(98)    20:01:00:05:30:00:47:9f

 0x63(99)    10:00:00:60:69:c0:0c:1d

0x64(100)    20:01:00:05:30:00:51:1f [Local]
```

```
0x65(101)    10:00:00:60:69:22:32:91 [Principal]
---------    -----------------------
```

Step 7 Verify the next hop and destination for the switch.

Example:

```
switch# show fspf internal route vsan 1

FSPF Unicast Routes
---------------------------
 VSAN Number  Dest Domain   Route Cost    Next hops
---------------------------------------------
           1     0x61(97)          500      fc2/2
           1     0x62(98)         1000      fc2/1
                                          fc2/2
           1     0x63(99)          500      fc2/1
           1    0x65(101)         1000      fc2/4
```

Step 8 Verify the name server information.

Example:

```
switch# show fcns data vsan 1

VSAN 1:
-----------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR) FC4-TYPE:FEATURE
-----------------------------------------------------------------------
0x610400    N     10:00:00:00:c9:24:3d:90 (Emulex)     scsi-fcp
0x6105dc    NL    21:00:00:20:37:28:31:6d (Seagate)    scsi-fcp
0x6105e0    NL    21:00:00:20:37:28:24:7b (Seagate)    scsi-fcp
0x6105e1    NL    21:00:00:20:37:28:22:ea (Seagate)    scsi-fcp
0x6105e2    NL    21:00:00:20:37:28:2e:65 (Seagate)    scsi-fcp
0x6105e4    NL    21:00:00:20:37:28:26:0d (Seagate)    scsi-fcp
0x630400    N     10:00:00:00:c9:24:3f:75 (Emulex)     scsi-fcp
0x630500    N     50:06:01:60:88:02:90:cb              scsi-fcp
0x6514e2    NL    21:00:00:20:37:a7:ca:b7 (Seagate)    scsi-fcp
0x6514e4    NL    21:00:00:20:37:a7:c7:e0 (Seagate)    scsi-fcp
0x6514e8    NL    21:00:00:20:37:a7:c7:df (Seagate)    scsi-fcp
0x651500    N     10:00:00:e0:69:f0:43:9f (JNI)

Total number of entries = 12
```

Note The Cisco switch name server shows both local and remote entries, and does not time out the entries.

Default Settings for Advanced Features

The following table lists the default settings for the features included in this chapter.

Table 32: Default Settings for Advanced Features

Parameters	Default
CIM server	Disabled
CIM server security protocol	HTTP
D_S_TOV	5,000 milliseconds
E_D_TOV	2,000 milliseconds
R_A_TOV	10,000 milliseconds
Timeout period to invoke fctrace	5 seconds
Number of frame sent by the fcping feature	5 frames
Remote capture connection protocol	TCP
Remote capture connection mode	Passive
Local capture frame limits	10 frames
FC ID allocation mode	Auto mode
Loop monitoring	Disabled
Interop mode	Disabled

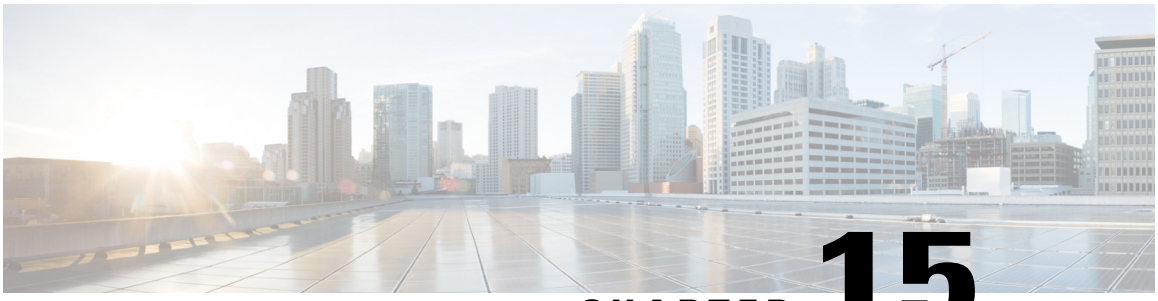

CHAPTER 15

Configuring FC-SP and DHCHAP

This chapter contains the following sections:

Configuring FC-SP and DHCHAP

Fibre Channel Security Protocol (FC-SP) capabilities provide switch-to-switch and host-to-switch authentication to overcome security challenges for enterprise-wide fabrics. Diffie-Hellman Challenge Handshake Authentication Protocol (DHCHAP) is an FC-SP protocol that provides authentication between Cisco Nexus 5000 Series switches and other devices. DHCHAP consists of the CHAP protocol combined with the Diffie-Hellman exchange.

Information About Fabric Authentication

All Cisco Nexus 5000 Series switches enable fabric-wide authentication from one switch to another switch, or from a switch to a host. These switch and host authentications are performed locally or remotely in each fabric. As storage islands are consolidated and migrated to enterprise-wide fabrics, new security challenges arise. The approach of securing storage islands cannot always be guaranteed in enterprise-wide fabrics. For example, in a campus environment with geographically distributed switches, someone could maliciously interconnect incompatible switches or you could accidentally do so, resulting in Inter-Switch Link (ISL) isolation and link disruption.

Cisco Nexus 5000 Series switches support authentication features to address physical security (see the following figure).

Figure 31: Switch and Host Authentication

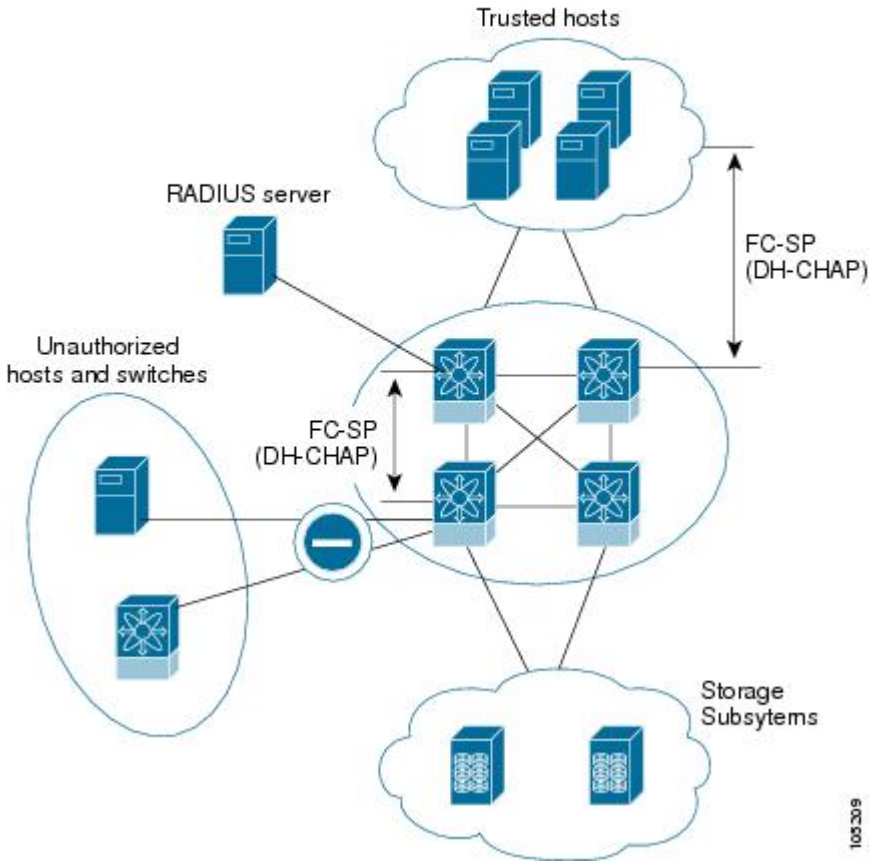

Note Fibre Channel Host Bus Adapters (HBAs) with appropriate firmware and drivers are required for host-switch authentication.

DHCHAP

DHCHAP is an authentication protocol that authenticates the devices connecting to a switch. Fibre Channel authentication allows only trusted devices to be added to a fabric, which prevents unauthorized devices from accessing the switch.

Note The terms FC-SP and DHCHAP are used interchangeably in this chapter.

DHCHAP is a mandatory password-based, key-exchange authentication protocol that supports both switch-to-switch and host-to-switch authentication. DHCHAP negotiates hash algorithms and DH groups before performing authentication. It supports MD5 and SHA-1 algorithm-based authentication.

To configure DHCHAP authentication using the local password database, perform this task:

SUMMARY STEPS

1. Enable DHCHAP.
2. Identify and configure the DHCHAP authentication modes.
3. Configure the hash algorithm and DH group.
4. Configure the DHCHAP password for the local switch and other switches in the fabric.
5. Configure the DHCHAP timeout value for reauthentication.
6. Verify the DHCHAP configuration.

DETAILED STEPS

Step 1 Enable DHCHAP.

Step 2 Identify and configure the DHCHAP authentication modes.

Step 3 Configure the hash algorithm and DH group.

Step 4 Configure the DHCHAP password for the local switch and other switches in the fabric.

Step 5 Configure the DHCHAP timeout value for reauthentication.

Step 6 Verify the DHCHAP configuration.

DHCHAP Compatibility with Fibre Channel Features

This section identifies the impact of configuring the DHCHAP feature along with existing Cisco NX-OS features:

- SAN port channel interfaces—If DHCHAP is enabled for ports belonging to a SAN port channel, DHCHAP authentication is performed at the physical interface level, not at the port channel level.
- Port security or fabric binding—Fabric-binding policies are enforced based on identities authenticated by DHCHAP.
- VSANs—DHCHAP authentication is not done on a per-VSAN basis.

About Enabling DHCHAP

By default, the DHCHAP feature is disabled in all switches.

You must explicitly enable the DHCHAP feature to access the configuration and verification commands for fabric authentication. When you disable this feature, all related configurations are automatically discarded.

Enabling DHCHAP

To enable DHCHAP for a Cisco Nexus 5000 Series switch, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcsp enable**
3. switch(config)# **no fcsp enable**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fcsp enable**	Enables the DHCHAP in this switch.
Step 3	switch(config)# **no fcsp enable**	Disables (default) the DHCHAP in this switch.

About DHCHAP Authentication Modes

The DHCHAP authentication status for each interface depends on the configured DHCHAP port mode.

When the DHCHAP feature is enabled in a switch, each Fibre Channel interface or FCIP interface may be configured to be in one of four DHCHAP port modes:

- On—During switch initialization, if the connecting device supports DHCHAP authentication, the software performs the authentication sequence. If the connecting device does not support DHCHAP authentication, the link is placed in an isolated state.
- Auto-Active—During switch initialization, if the connecting device supports DHCHAP authentication, the software performs the authentication sequence. If the connecting device does not support DHCHAP authentication, the software continues with the rest of the initialization sequence.
- Auto-Passive (default)—The switch does not initiate DHCHAP authentication, but participates in DHCHAP authentication if the connecting device initiates DHCHAP authentication.
- Off—The switch does not support DHCHAP authentication. Authentication messages sent to ports in this mode return error messages to the initiating switch.

Note Whenever DHCHAP port mode is changed to a mode other than the Off mode, reauthentication is performed.

The following table identifies switch-to-switch authentication between two Cisco switches in various modes.

Table 33: DHCHAP Authentication Status Between Two MDS Switches

Switch N DHCHAP Modes	Switch 1 DHCHAP Modes			
	on	**auto-active**	**auto-passive**	**off**
on	FC-SP authentication is performed.	FC-SP authentication is performed.	FC-SP authentication is performed.	Link is brought down.
auto-Active				FC-SP authentication is *not* performed.
auto-Passive			FC-SP authentication is *not* performed.	
off	Link is brought down.	FC-SP authentication is *not* performed.		

Configuring the DHCHAP Mode

To configure the DHCHAP mode for a particular interface, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface fc** *slot*/*port* - *slot*/*port*
3. switch(config-if)# **fcsp on**
4. switch(config-if)# **no fcsp on**
5. switch(config-if)# **fcsp auto-active 0**
6. switch(config-if)# **fcsp auto-active** *timeout-period*
7. switch(config-if)# **fcsp auto-active**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **interface fc** *slot*/*port* - *slot*/*port*	Selects a range of interfaces and enters the interface configuration mode.
Step 3	switch(config-if)# **fcsp on**	Sets the DHCHAP mode for the selected interfaces to be in the on state.
Step 4	switch(config-if)# **no fcsp on**	Reverts to the factory default of auto-passive for these three interfaces.
Step 5	switch(config-if)# **fcsp auto-active 0**	Changes the DHCHAP authentication mode for the selected interfaces to auto-active. Zero (0) indicates that the port does not perform reauthentication.

	Command or Action	Purpose
		Note The reauthorization interval configuration is the same as the default behavior.
Step 6	switch(config-if)# **fcsp auto-active** *timeout-period*	Changes the DHCHAP authentication mode to auto-active for the selected interfaces. The timeout period value (in minutes) sets how often reauthentication occurs after the initial authentication.
Step 7	switch(config-if)# **fcsp auto-active**	Changes the DHCHAP authentication mode to auto-active for the selected interfaces. Reauthentication is disabled (default). **Note** The reauthorization interval configuration is the same as setting it to zero (0).

About the DHCHAP Hash Algorithm

Cisco SAN switches support a default hash algorithm priority list of MD5 followed by SHA-1 for DHCHAP authentication.

If you change the hash algorithm configuration, then change it globally for all switches in the fabric.

Caution RADIUS and TACACS+ protocols always use MD5 for CHAP authentication. Using SHA-1 as the hash algorithm may prevent RADIUS and TACACS+ usage, even if these AAA protocols are enabled for DHCHAP authentication.

Configuring the DHCHAP Hash Algorithm

To configure the hash algorithm, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcsp dhchap hash** [**md5**] [**sha1**]
3. switch(config)# **no fcsp dhchap hash sha1**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fcsp dhchap hash** [**md5**] [**sha1**]	Configures the use of the the MD5 or SHA-1 hash algorithm.

	Command or Action	Purpose
Step 3	switch(config)# **no fcsp dhchap hash sha1**	Reverts to the factory default priority list of the MD5 hash algorithm followed by the SHA-1 hash algorithm.

About the DHCHAP Group Settings

All Cisco Nexus 5000 Series switches support all DHCHAP groups specified in the standard: 0 (null DH group, which does not perform the Diffie-Hellman exchange), 1, 2, 3, or 4.

If you change the DH group configuration, change it globally for all switches in the fabric.

Configuring the DHCHAP Group Settings

To change the DH group settings, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcsp dhchap dhgroup** [**0** | **1** | **2** | **3** | **4**]
3. switch(config)# **no fcsp dhchap dhgroup** [**0** | **1** | **2**| **3** | **4**]

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fcsp dhchap dhgroup** [**0** | **1** | **2** | **3** | **4**]	Prioritizes the use of DH groups in the configured order.
Step 3	switch(config)# **no fcsp dhchap dhgroup** [**0** | **1** | **2**| **3** | **4**]	Reverts to the DHCHAP factory default order of 0, 4, 1, 2, and 3.

About the DHCHAP Password

DHCHAP authentication in each direction requires a shared secret password between the connected devices. To do this, you can use one of three configurations to manage passwords for all switches in the fabric that participate in DHCHAP:

- Configuration 1—Use the same password for all switches in the fabric. This is the simplest configuration. When you add a new switch, you use the same password to authenticate that switch in this fabric. It is also the most vulnerable configuration if someone from the outside maliciously attempts to access any one switch in the fabric.

- Configuration 2—Use a different password for each switch and maintain that password list in each switch in the fabric. When you add a new switch, you create a new password list and update all switches with the new list. Accessing one switch yields the password list for all switches in that fabric.
- Configuration 3—Use different passwords for different switches in the fabric. When you add a new switch, multiple new passwords corresponding to each switch in the fabric must be generated and configured in each switch. Even if one switch is compromised, the password of other switches are still protected. This configuration requires considerable password maintenance by the user.

Note All passwords are restricted to 64 alphanumeric characters and can be changed, but not deleted.

We recommend using RADIUS or TACACS+ for fabrics with more than five switches. If you need to use a local password database, you can continue to do so using Configuration 3 and using the Cisco MDS 9000 Family Fabric Manager to manage the password database.

Configuring DHCHAP Passwords for the Local Switch

To configure the DHCHAP password for the local switch, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcsp dhchap password** [**0** | **7**] *password* [**wwn** *wwn-id*]

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fcsp dhchap password** [**0** | **7**] *password* [**wwn** *wwn-id*]	Configures a clear text password for the local switch.

About Password Configuration for Remote Devices

You can configure passwords in the local authentication database for other devices in a fabric. The other devices are identified by their device name, which is also known as the switch WWN or device WWN. The password is restricted to 64 characters and can be specified in clear text (0) or in encrypted text (7).

Note The switch WWN identifies the physical switch. This WWN is used to authenticate the switch and is different from the VSAN node WWN.

Configuring DHCHAP Passwords for Remote Devices

To locally configure the remote DHCHAP password for another switch in the fabric, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcsp dhchap devicename** *switch-wwn* **password** *password*
3. switch(config)# **no fcsp dhchap devicename** *switch-wwn* **password** *password*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fcsp dhchap devicename** *switch-wwn* **password** *password*	Configures a password for another switch in the fabric that is identified by the switch WWN device name.
Step 3	switch(config)# **no fcsp dhchap devicename** *switch-wwn* **password** *password*	Removes the password entry for this switch from the local authentication database.

About the DHCHAP Timeout Value

During the DHCHAP protocol exchange, if the Cisco Nexus 5000 Series switch does not receive the expected DHCHAP message within a specified time interval, authentication failure is assumed. The time ranges from 20 (no authentication is performed) to 1000 seconds. The default is 30 seconds.

When changing the timeout value, consider the following factors:

- The existing RADIUS and TACACS+ timeout values.
- The same value must also be configured on all switches in the fabric.

Configuring the DHCHAP Timeout Value

To configure the DHCHAP timeout value, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcsp timeout** *timeout*
3. switch(config)# **no fcsp timeout** *timeout*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fcsp timeout** *timeout*	Configures the reauthentication timeout to the specified value. The unit is seconds.
Step 3	switch(config)# **no fcsp timeout** *timeout*	Reverts to the factory default of 30 seconds.

Configuring DHCHAP AAA Authentication

You can configure AAA authentication to use a RADIUS or TACACS+ server group. If AAA authentication is not configured, local authentication is used by default.

Displaying Protocol Security Information

Use the **show fcsp** commands to display configurations for the local database.

The following example shows how to display the DHCHAP configuration for the specified interface:

```
switch# show fcsp interface fc2/4
fc2/4:
        fcsp authentication mode:SEC_MODE_ON
        Status: Successfully authenticated
```

The following example shows how to display DHCHAP statistics for the specified interface:

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

The following example shows how to display the FC-SP WWN of the device connected to the specified interface:

```
switch# show fcsp interface fc2/1 wwn
```

The following example shows how to display the hash algorithm and DHCHAP groups configured in the switch:

```
switch# show fcsp dhchap
```

The following example shows how to display the DHCHAP local password database:

```
switch# show fcsp dhchap database
```

Use the ASCII representation of the device WWN to configure the switch information on RADIUS and TACACS+ servers.

Sample Configuration

This section provides the steps to configure the example illustrated in the following figure.

Figure 32: Sample DHCHAP Authentication

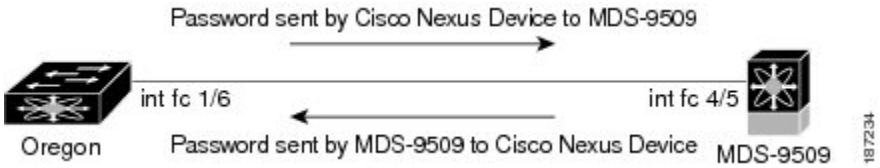

To configure the authentication setup shown in the above figure, perform this task:

SUMMARY STEPS

1. Obtain the device name of the Cisco Nexus 5000 Series switch in the fabric. The Cisco Nexus 5000 Series switch in the fabric is identified by the switch WWN.
2. Explicitly enable DHCHAP in this switch.
3. Configure a clear text password for this switch. This password will be used by the connecting device.
4. Configure a password for another switch in the fabric that is identified by the switch WWN device name.
5. Enable the DHCHAP mode for the required Fibre Channel interface.
6. Verify the protocol security information configured in this switch by displaying the DHCHAP local password database.
7. Display the DHCHAP configuration in the Fibre Channel interface.
8. Repeat these steps on the connecting MDS 9509 switch.

DETAILED STEPS

Step 1 Obtain the device name of the Cisco Nexus 5000 Series switch in the fabric. The Cisco Nexus 5000 Series switch in the fabric is identified by the switch WWN.

Example:
```
switch# show wwn switch
Switch WWN is 20:00:00:05:30:00:54:de
```

Step 2 Explicitly enable DHCHAP in this switch.

Note When you disable DHCHAP, all related configurations are automatically discarded.

Example:
```
switch(config)# fcsp enable
```

Step 3 Configure a clear text password for this switch. This password will be used by the connecting device.

Example:
```
switch(config)# fcsp dhchap password rtp9216
```

Step 4 Configure a password for another switch in the fabric that is identified by the switch WWN device name.

Example:
```
switch(config)# fcsp dhchap devicename 20:00:00:05:30:00:38:5e password rtp9509
```

Step 5 Enable the DHCHAP mode for the required Fibre Channel interface.

Note Whenever DHCHAP port mode is changed to a mode other than the Off mode, reauthentication is performed.

Example:
```
switch(config)# interface fc2/4
switch(config-if)# fcsp on
```

Step 6 Verify the protocol security information configured in this switch by displaying the DHCHAP local password database.

Example:
```
switch# show fcsp dhchap database
DHCHAP Local Password:
        Non-device specific password:*******
Other Devices' Passwords:
        Password for device with WWN:20:00:00:05:30:00:38:5e is *******
```

Step 7 Display the DHCHAP configuration in the Fibre Channel interface.

Example:
```
switch# show fcsp interface fc2/4
fc2/4
        fcsp authentication mode:SEC_MODE_ON
        Status:Successfully authenticated
```

Step 8 Repeat these steps on the connecting MDS 9509 switch.

Example:
```
MDS-9509# show wwn switch
Switch WWN is 20:00:00:05:30:00:38:5e
MDS-9509(config)# fcsp enable
MDS-9509(config)# fcsp dhchap password rtp9509
MDS-9509(config)# fcsp dhchap devicename 20:00:00:05:30:00:54:de password rtp9216
MDS-9509(config)# interface fc 4/5
MDS-9509(config-if)# fcsp on
MDS-9509# show fcsp dhchap database
DHCHAP Local Password:
        Non-device specific password:*******
Other Devices' Passwords:
        Password for device with WWN:20:00:00:05:30:00:54:de is *******
MDS-9509# show fcsp interface fc2/4
Fc2/4
        fcsp authentication mode:SEC_MODE_ON
        Status:Successfully authenticated
```

You have now enabled and configured DHCHAP authentication for the sample setup in shown in the figure above.

Default Fabric Security Settings

The following table lists the default settings for all fabric security features in any switch.

Table 34: Default Fabric Security Settings

Parameters	Default
DHCHAP feature	Disabled
DHCHAP hash algorithm	A priority list of MD5 followed by SHA-1 for DHCHAP authentication
DHCHAP authentication mode	Auto-passive
DHCHAP group default priority exchange order	0, 4, 1, 2, and 3, respectively
DHCHAP timeout value	30 seconds

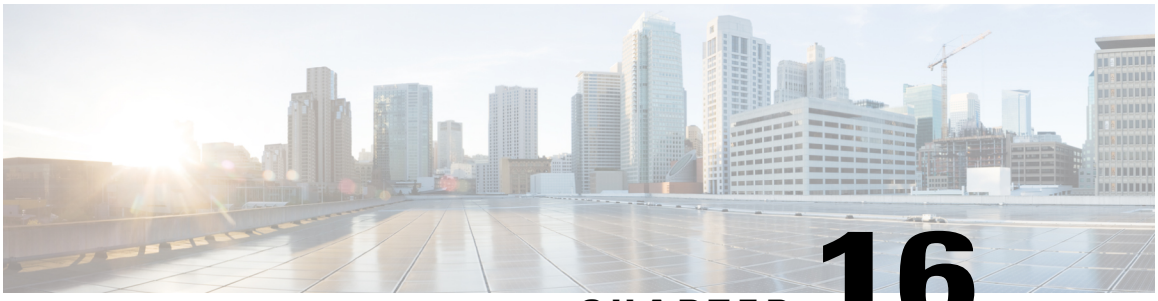

CHAPTER 16

Configuring Port Security

This chapter contains the following sections:

Configuring Port Security

Cisco Nexus 5000 Series switches provide port security features that reject intrusion attempts and report these intrusions to the administrator.

Note Port security is supported on virtual Fibre Channel ports and physical Fibre Channel ports.

Information About Port Security

Typically, any Fibre Channel device in a SAN can attach to any SAN switch port and access SAN services based on zone membership. Port security features prevent unauthorized access to a switch port in the Cisco Nexus 5000 Series switch, using the following methods:

- Login requests from unauthorized Fibre Channel devices (N ports) and switches (xE ports) are rejected.
- All intrusion attempts are reported to the SAN administrator through system messages.
- Configuration distribution uses the CFS infrastructure, and is limited to those switches that are CFS capable. Distribution is disabled by default.
- Configuring the port security policy requires the Storage Protocol Services license.

Port Security Enforcement

To enforce port security, configure the devices and switch port interfaces through which each device or switch is connected, and activate the configuration.

- Use the port world wide name (pWWN) or the node world wide name (nWWN) to specify the N port connection for each device.
- Use the switch world wide name (sWWN) to specify the xE port connection for each switch.

Each N and xE port can be configured to restrict a single port or a range of ports.

Enforcement of port security policies are done on every activation and when the port tries to come up.

The port security feature uses two databases to accept and implement configuration changes.

- Configuration database—All configuration changes are stored in the configuration database.
- Active database—The database currently enforced by the fabric. The port security feature requires all devices connecting to a switch to be part of the port security active database. The software uses this active database to enforce authorization.

About Auto-Learning

You can instruct the switch to automatically learn (auto-learn) the port security configurations over a specified period. This feature allows any Cisco Nexus 5000 Series switch to automatically learn about devices and switches that connect to it. Use this feature when you activate the port security feature for the first time as it saves tedious manual configuration for each port. You must configure auto-learning on a per-VSAN basis. If enabled, devices and switches that are allowed to connect to the switch are automatically learned, even if you have not configured any port access.

When auto-learning is enabled, learning happens only for the devices or interfaces that were not already logged into the switch. Learned entries on a port are cleaned up after you shut down that port if auto-learning is still enabled.

Learning does not override the existing configured port security policies. For example, if an interface is configured to allow a specific pWWN, then auto-learning will not add a new entry to allow any other pWWN on that interface. All other pWWNs will be blocked even in auto-learning mode.

No entries are learned for a port in the shutdown state.

When you activate the port security feature, auto-learning is also automatically enabled.

Note If you enable auto-learning before activating port security, you cannot activate port security until auto-learning is disabled.

Port Security Activation

By default, the port security feature is not activated in Cisco Nexus 5000 Series switches.

When you activate the port security feature, the following operations occur:

- Auto-learning is also automatically enabled, which means:
  - From this point, auto-learning happens only for the devices or interfaces that were not logged into the switch.
  - You cannot activate the database until you disable auto-learning.

- All the devices that are already logged in are learned and are added to the active database.
- All entries in the configured database are copied to the active database.

After the database is activated, subsequent device login is subject to the activated port bound WWN pairs, excluding the auto-learned entries. You must disable auto-learning before the auto-learned entries become activated.

When you activate the port security feature, auto-learning is also automatically enabled. You can choose to activate the port security feature and disable auto-learning.

If a port is shut down because of a denied login attempt, and you subsequently configure the database to allow that login, the port does not come up automatically. You must explicitly enter a **no shutdown** CLI command to bring that port back online.

Configuring Port Security

Configuring Port Security with Auto-Learning and CFS Distribution

To configure port security, using auto-learning and CFS distribution, perform this task:

SUMMARY STEPS

1. Enable port security.
2. Enable CFS distribution.
3. Activate port security on each VSAN.
4. Issue a CFS commit to copy this configuration to all switches in the fabric.
5. Wait until all switches and all hosts are automatically learned.
6. Disable auto-learn on each VSAN.
7. Issue a CFS commit to copy this configuration to all switches in the fabric.
8. Copy the active database to the configure database on each VSAN.
9. Issue a CFS commit to copy this configuration to all switches in the fabric.
10. Copy the running configuration to the startup configuration, using the fabric option.

DETAILED STEPS

Step 1 Enable port security.

Step 2 Enable CFS distribution.

Step 3 Activate port security on each VSAN.
This action turns on auto-learning by default.

Step 4 Issue a CFS commit to copy this configuration to all switches in the fabric.
All switches have port security activated with auto-learning enabled.

Step 5 Wait until all switches and all hosts are automatically learned.

Step 6 Disable auto-learn on each VSAN.

Step 7 Issue a CFS commit to copy this configuration to all switches in the fabric.

The auto-learned entries from every switch are combined into a static active database that is distributed to all switches.

Step 8 Copy the active database to the configure database on each VSAN.

Step 9 Issue a CFS commit to copy this configuration to all switches in the fabric.
This ensures that the configure database is the same on all switches in the fabric.

Step 10 Copy the running configuration to the startup configuration, using the fabric option.

Related Topics

Activating Port Security, on page 218
Committing the Changes, on page 227
Copying the Port Security Database, on page 233
Disabling Auto-Learning, on page 221
Enabling Port Security, on page 217

Configuring Port Security with Auto-Learning without CFS

To configure port security using auto-learning without CFS, perform this task:

SUMMARY STEPS

1. Enable port security.
2. Activate port security on each VSAN, which turns on auto-learning by default.
3. Wait until all switches and all hosts are automatically learned.
4. Disable auto-learn on each VSAN.
5. Copy the active database to the configure database on each VSAN.
6. Copy the running configuration to the startup configuration, which saves the port security configuration database to the startup configuration.
7. Repeat the above steps for all switches in the fabric.

DETAILED STEPS

Step 1 Enable port security.

Step 2 Activate port security on each VSAN, which turns on auto-learning by default.

Step 3 Wait until all switches and all hosts are automatically learned.

Step 4 Disable auto-learn on each VSAN.

Step 5 Copy the active database to the configure database on each VSAN.

Step 6 Copy the running configuration to the startup configuration, which saves the port security configuration database to the startup configuration.

Step 7 Repeat the above steps for all switches in the fabric.

Related Topics

Activating Port Security, on page 218
Copying the Port Security Database, on page 233
Disabling Auto-Learning, on page 221
Enabling Port Security, on page 217

Configuring Port Security with Manual Database Configuration

To configure port security and manually configure the port security database, perform this task:

SUMMARY STEPS

1. Enable port security.
2. Manually configure all port security entries into the configure database on each VSAN.
3. Activate port security on each VSAN. This turns on auto-learning by default.
4. Disable auto-learn on each VSAN.
5. Copy the running configuration to the startup configuration, which saves the port security configuration database to the startup configuration.
6. Repeat the above steps for all switches in the fabric.

DETAILED STEPS

Step 1	Enable port security.
Step 2	Manually configure all port security entries into the configure database on each VSAN.
Step 3	Activate port security on each VSAN. This turns on auto-learning by default.
Step 4	Disable auto-learn on each VSAN.
Step 5	Copy the running configuration to the startup configuration, which saves the port security configuration database to the startup configuration.
Step 6	Repeat the above steps for all switches in the fabric.

Enabling Port Security

By default, the port security feature is disabled in Cisco Nexus 5000 Series switches.

To enable port security, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **port-security enable**
3. switch(config)# **no port-security enable**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **port-security enable**	Enables port security on that switch.
Step 3	switch(config)# **no port-security enable**	Disables (default) port security on that switch.

Port Security Activation

Activating Port Security

To activate port security, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **port-security activate vsan** *vsan-id*
3. switch(config)# **port-security activate vsan** *vsan-id* **no-auto-learn**
4. switch(config)# **no port-security activate vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **port-security activate vsan** *vsan-id*	Activates the port security database for the specified VSAN, and automatically enables auto-learning.
Step 3	switch(config)# **port-security activate vsan** *vsan-id* **no-auto-learn**	Activates the port security database for the specified VSAN, and disables auto-learning.
Step 4	switch(config)# **no port-security activate vsan** *vsan-id*	Deactivates the port security database for the specified VSAN, and automatically disables auto-learning.

Database Activation Rejection

Database activation is rejected in the following cases:

- Missing or conflicting entries exist in the configuration database but not in the active database.

- The auto-learning feature was enabled before the activation. To reactivate a database in this state, disable auto-learning.
- The exact security is not configured for each port channel member.
- The configured database is empty but the active database is not.

If the database activation is rejected due to one or more conflicts listed in the previous section, you may decide to proceed by forcing the port security activation.

Forcing Port Security Activation

If the port security activation request is rejected, you can force the activation.

Note If you force the activation, existing devices are logged out if they violate the active database.

You can view missing or conflicting entries using the **port-security database diff active vsan** command in EXEC mode.

To forcefully activate the port security database, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **port-security activate vsan** *vsan-id* **force**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **port-security activate vsan** *vsan-id* **force**	Forces the port security database to activate for the specified VSAN even if conflicts occur.

Database Reactivation

Tip If auto-learning is enabled, you cannot activate the database without the force option until you disable auto-learning.

To reactivate the port security database, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **no port-security auto-learn vsan** *vsan-id*
3. switch(config)# **exit**
4. switch# **port-security database copy vsan** *vsan-id*
5. switch# **configuration terminal**
6. switch(config)# **port-security activate vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **no port-security auto-learn vsan** *vsan-id*	Disables auto-learning and stops the switch from learning about new devices accessing the switch. Enforces the database contents based on the devices learned up to this point.
Step 3	switch(config)# **exit**	
Step 4	switch# **port-security database copy vsan** *vsan-id*	Copies from the active to the configured database.
Step 5	switch# **configuration terminal**	Re-enters configuration mode.
Step 6	switch(config)# **port-security activate vsan** *vsan-id*	Activates the port security database for the specified VSAN, and automatically enables auto-learning.

Auto-Learning

About Enabling Auto-Learning

The state of the auto-learning configuration depends on the state of the port security feature:

- If the port security feature is not activated, auto-learning is disabled by default.
- If the port security feature is activated, auto-learning is enabled by default (unless you explicitly disabled this option).

Tip If auto-learning is enabled on a VSAN, you can only activate the database for that VSAN by using the force option.

Enabling Auto-Learning

To enable auto-learning, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **port-security auto-learn vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **port-security auto-learn vsan** *vsan-id*	Enables auto-learning so the switch can learn about any device that is allowed to access VSAN 1. These devices are logged in the port security active database.

Disabling Auto-Learning

To disable auto-learning, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **no port-security auto-learn vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **no port-security auto-learn vsan** *vsan-id*	Disables auto-learning and stops the switch from learning about new devices accessing the switch. Enforces the database contents based on the devices learned up to this point.

Auto-Learning Device Authorization

The following table summarizes the authorized connection conditions for device requests.

Table 35: Authorized Auto-Learning Device Requests

Condition	Device (pWWN, nWWN, sWWN)	Requests Connection to	Authorization
1	Configured with one or more switch ports	A configured switch port	Permitted
2		Any other switch port	Denied
3	Not configured	A switch port that is not configured	Permitted if auto-learning enabled
4			Denied ifauto-learning disabled
5	Configured or not configured	A switch port that allows any device	Permitted
6	Configured to log in to any switch port	Any port on the switch	Permitted
7	Not configured	A port configured with some other device	Denied

Authorization Scenario

Assume that the port security feature is activated and the following conditions are specified in the active database:

- A pWWN (P1) is allowed access through interface fc2/1 (F1).
- A pWWN (P2) is allowed access through interface fc2/2 (F1).
- A nWWN (N1) is allowed access through interface fc2/2 (F2).
- Any WWN is allowed access through interface vfc3/1 (F3).
- A nWWN (N3) is allowed access through any interface.
- A pWWN (P3) is allowed access through interface fc2/4 (F4).
- A sWWN (S1) is allowed access through interface fc3/1-3 (F10 to F13).
- A pWWN (P10) is allowed access through interface vfc4/1 (F11).

The following table summarizes the port security authorization results for this active database.

Table 36: Authorization Results for Scenario

Device Connection Request	Authorization	Condition	Reason
P1, N2, F1	Permitted	1	No conflict.

Device Connection Request	Authorization	Condition	Reason
P2, N2, F1	Permitted	1	No conflict.
P3, N2, F1	Denied	2	F1 is bound to P1/P2.
P1, N3, F1	Permitted	6	Wildcard match for N3.
P1, N1, F3	Permitted	5	Wildcard match for F3.
P1, N4, F5	Denied	2	P1 is bound to F1.
P5, N1, F5	Denied	2	N1 is only allowed on F2.
P3, N3, F4	Permitted	1	No conflict.
S1, F10	Permitted	1	No conflict.
S2, F11	Denied	7	P10 is bound to F11.
P4, N4, F5 (auto-learning on)	Permitted	3	No conflict.
P4, N4, F5 (auto-learning off)	Denied	4	No match.
S3, F5 (auto-learning on)	Permitted	3	No conflict.
S3, F5 (auto-learning off)	Denied	4	No match.
P1, N1, F6 (auto-learning on)	Denied	2	P1 is bound to F1.
P5, N5, F1 (auto-learning on)	Denied	7	Only P1 and P2 bound to F1.
S3, F4 (auto-learning on)	Denied	7	P3 paired with F4.
S1, F3 (auto-learning on)	Permitted	5	No conflict.
P5, N3, F3	Permitted	6	Wildcard (*) match for F3 and N3.
P7, N3, F9	Permitted	6	Wildcard (*) match for N3.

Related Topics

Auto-Learning Device Authorization, on page 221

Port Security Manual Configuration

To configure port security on a Cisco Nexus 5000 Series switch, perform this task:

SUMMARY STEPS

1. Identify the WWN of the ports that need to be secured.
2. Secure the fWWN to an authorized nWWN or pWWN.
3. Activate the port security database.
4. Verify your configuration.

DETAILED STEPS

Step 1 Identify the WWN of the ports that need to be secured.

Step 2 Secure the fWWN to an authorized nWWN or pWWN.

Step 3 Activate the port security database.

Step 4 Verify your configuration.

WWN Identification Guidelines

If you decide to manually configure port security, note the following guidelines:

- Identify switch ports by the interface or by the fWWN.
- Identify devices by the pWWN or by the nWWN.
- If an N port is allowed to log in to SAN switch port F, then that N port can only log in through the specified F port.
- If an N port's nWWN is bound to an F port WWN, then all pWWNs in the N port are implicitly paired with the F port.
- TE port checking is done on each VSAN in the allowed VSAN list of the VSAN trunk port.
- All port channel xE ports must be configured with the same set of WWNs in the same SAN port channel.
- E port security is implemented in the port VSAN of the E port. In this case, the sWWN is used to secure authorization checks.
- Once activated, the configuration database can be modified without any effect on the active database.
- By saving the running configuration, you save the configuration database and activated entries in the active database. Learned entries in the active database are not saved.

Adding Authorized Port Pairs

After identifying the WWN pairs that need to be bound, add those pairs to the port security database.

Tip Remote switch binding can be specified at the local switch. To specify the remote interfaces, you can use either the fWWN or sWWN-interface combination.

To add authorized port pairs for port security, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **port-security database vsan** *vsan-id*
3. switch(config)# **no port-security database vsan** *vsan-id*
4. switch(config-port-security)# **swwn** *swwn-id* **interface san-port-channel 5**
5. switch(config-port-security)# **any-wwn interface fc** *slot*/*port* - **fc** *slot*/*port*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **port-security database vsan** *vsan-id*	Enters the port security database mode for the specified VSAN.
Step 3	switch(config)# **no port-security database vsan** *vsan-id*	Deletes the port security configuration database from the specified VSAN.
Step 4	switch(config-port-security)# **swwn** *swwn-id* **interface san-port-channel 5**	Configures the specified sWWN to only log in through SAN port channel 5.
Step 5	switch(config-port-security)# **any-wwn interface fc** *slot*/*port* - **fc** *slot*/*port*	Configures any WWN to log in through the specified interfaces.

This example enters the port security database mode for VSAN 2:

```
switch(config)# port-security database vsan 2
```

This example configures the specified sWWN to only log in through SAN port channel 5:

```
switch(config-port-security)# swwn 20:01:33:11:00:2a:4a:66 interface san-port-channel 5
```

This example configures the specified pWWN to log in through the specified interface in the specified switch:

```
switch(config-port-security)# pwwn 20:11:33:11:00:2a:4a:66 swwn 20:00:00:0c:85:90:3e:80
interface fc 3/2
```

This example configures any WWN to log in through the specified interface in any switch:

```
switch(config-port-security)# any-wwn interface fc 3/2
```

Port Security Configuration Distribution

The port security feature uses the Cisco Fabric Services (CFS) infrastructure to enable efficient database management, provide a single point of configuration for the entire fabric in the VSAN, and enforce the port security policies throughout the fabric.

For additional information, refer to Using Cisco Fabric Services in the *Cisco Nexus 5000 Series System Management Configuration Guide*.

Enabling Port Security Distribution

All the configurations performed in distributed mode are stored in a pending (temporary) database. If you modify the configuration, you need to commit or discard the pending database changes to the configurations. The fabric remains locked during this period. Changes to the pending database are not reflected in the configurations until you commit the changes.

Note Port activation or deactivation and auto-learning enable or disable do not take effect until after a CFS commit if CFS distribution is enabled. Always follow any one of these operations with a CFS commit to ensure proper configuration.

For example, if you activate port security, follow up by disabling auto-learning, and finally commit the changes in the pending database, then the net result of your actions is the same as entering a **port-security activate vsan** *vsan-id* no-auto-learn command.

Tip We recommend that you perform a commit after you activate port security and after you enable auto learning.

To enable the port security distribution, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **port-security distribute**
3. switch(config)# **no port-security distribute**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **port-security distribute**	Enables distribution.
Step 3	switch(config)# **no port-security distribute**	Disables distribution.

Related Topics

Activation and Auto-Learning Configuration Distribution, on page 228

Locking the Fabric

The first action that modifies the existing configuration 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.

Committing the Changes

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

To commit the port security configuration changes for the specified VSAN, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **port-security commit vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **port-security commit vsan** *vsan-id*	Commits the port security changes in the specified VSAN.

Discarding the Changes

If you discard (abort) the changes made to the pending database, the configuration remains unaffected and the lock is released.

To discard the port security configuration changes for the specified VSAN, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **port-security abort vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **port-security abort vsan** *vsan-id*	Discards the port security changes in the specified VSAN and clears the pending configuration database.

Activation and Auto-Learning Configuration Distribution

Activation and auto-learning configurations in distributed mode are remembered as actions to be performed when you commit the changes in the pending database.

Learned entries are temporary and do not have any role in determining if a login is authorized or not. As such, learned entries do not participate in distribution. When you disable learning and commit the changes in the pending database, the learned entries become static entries in the active database and are distributed to all switches in the fabric. After the commit, the active database on all switches are identical and learning can be disabled.

If the pending database contains more than one activation and auto-learning configuration when you commit the changes, the activation and auto-learning changes are consolidated and the resulting operation may change (see the following table.

Table 37: Scenarios for Activation and Auto-learning Configurations in Distributed Mode

Scenario	Actions	Distribution = OFF	Distribution = ON
A and B exist in the configuration database, activation is not done and devices C,D are logged in.	1. You activate the port security database and enable auto-learning.	configuration database = {A,B} active database = {A,B, C[1], D*}	configuration database = {A,B} active database = {null} pending database = {A,B + activation to be enabled}
	2. A new entry E is added to the configuration database.	configuration database = {A,B, E} active database = {A,B, C*, D*}	configuration database = {A,B} active database = {null} pending database = {A,B, E + activation to be enabled}
	3. You issue a commit.	Not applicable	configuration database = {A,B, E} active database = {A,B, E, C*, D*} pending database = empty
A and B exist in the configuration database, activation is not done and devices C,D are logged in.	1. You activate the port security database and enable auto-learning.	configuration database = {A,B} active database = {A,B, C*, D*}	configuration database = {A,B} active database = {null} pending database = {A,B + activation to be enabled}
	2. You disable learning.	configuration database = {A,B} active database = {A,B, C, D}	configuration database = {A,B} active database = {null} pending database = {A,B + activation to be enabled +learning to be disabled}
	3. You issue a commit.	Not applicable	configuration database = {A,B} active database = {A,B} and devices C and D are logged out. This is equal to an activation with auto-learning disabled. pending database = empty

1 The * (asterisk) indicates learned entries.

Port Security Database Merge Guidelines

A database merge refers to a union of the configuration database and static (unlearned) entries in the active database.

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

- Verify that the activation status and the auto-learning status is the same in both fabrics.
- Verify that the combined number of configurations for each VSAN in both databases does not exceed 2000.

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

For additional information, refer to CFS Merge Support in the *Cisco Nexus 5000 Series System Management Configuration Guide*.

Database Interaction

The following table lists the differences and interaction between the active and configuration databases.

Table 38: Active and Configuration Port Security Databases

Active Database	Configuration Database
Read-only.	Read-write.
Saving the configuration only saves the activated entries. Learned entries are not saved.	Saving the configuration saves all the entries in the configuration database.
Once activated, all devices that have already logged into the VSAN are also learned and added to the active database.	Once activated, the configuration database can be modified without any effect on the active database.
You can overwrite the active database with the configured database by activating the port security database. Forcing an activation may violate the entries already configured in the active database.	You can overwrite the configuration database with the active database.

Note You can overwrite the configuration database with the active database using the **port-security database copy vsan** command. The **port-security database diff active vsan** command in EXEC mode lists the differences between the active database and the configuration database.

Database Scenarios

the follwowing figure illustrates various scenarios showing the active database and the configuration database status based on port security configurations.

Figure 33: Port Security Database Scenarios

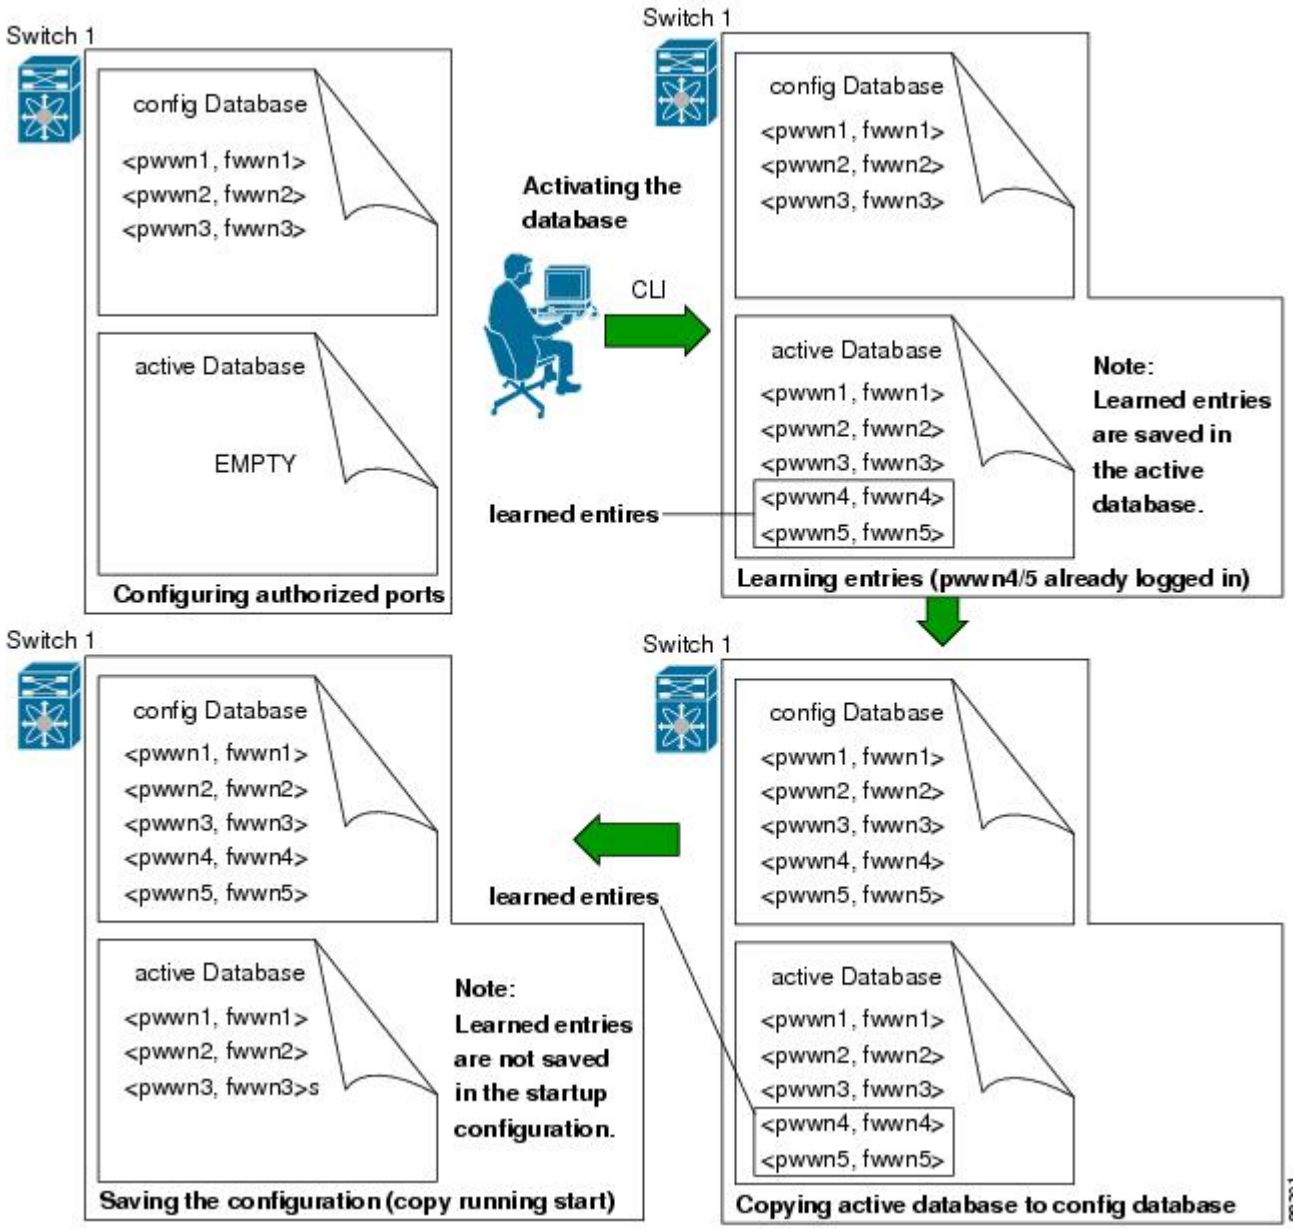

Copying the Port Security Database

Tip We recommend that you copy the active database to the config database after disabling auto-learning. This action will ensure that the configuration database is in synchronization with the active database. If distribution is enabled, this command creates a temporary copy (and consequently a fabric lock) of the configuration database. If you lock the fabric, you need to commit the changes to the configuration databases in all the switches.

Use the **port-security database copy vsan** command to copy from the active to the configured database. If the active database is empty, this command is not accepted.

```
switch# port-security database copy vsan 1
```

Use the **port-security database diff active vsan** command to view the differences between the active database and the configuration database. This command can be used when resolving conflicts.

```
switch# port-security database diff active vsan 1
```

Use the**port-security database diff config vsan** command to obtain information on the differences between the configuration database and the active database.

```
switch# port-security database diff config vsan 1
```

Deleting the Port Security Database

Tip If the distribution is enabled, the deletion creates a copy of the database. An explicit **port-security commit** command is required to actually delete the database.

Use the **no port-security database vsan** command in configuration mode to delete the configured database for a specified VSAN.

```
switch(config)# no port-security database vsan 1
```

Clearing the Port Security Database

Use the **clear port-security statistics vsan** command to clear all existing statistics from the port security database for a specified VSAN.

```
switch# clear port-security statistics vsan 1
```

Use the **clear port-security database auto-learn interface** command to clear any learned entries in the active database for a specified interface within a VSAN.

```
switch# clear port-security database auto-learn interface fc2/1 vsan 1
```

Use the **clear port-security database auto-learn vsan** command to clear any learned entries in the active database for the entire VSAN.

```
switch# clear port-security database auto-learn vsan 1
```

Note The **clear port-security database auto-learn** and **clear port-security statistics** commands are only relevant to the local switch and do not acquire locks. Also, learned entries are only local to the switch and do not participate in distribution.

Use the **port-security clear vsan** command to clear the pending session in the VSAN from any switch in the VSAN.

```
switch# clear port-security session vsan 5
```

Displaying Port Security Configuration

The **show port-security database** commands display the configured port security information. You can optionally specify a fWWN and a VSAN, or an interface and a VSAN in the **show port-security** command to view the output of the activated port security.

Access information for each port can be individually displayed. If you specify the fWWN or interface options, all devices that are paired in the active database (at that point) with the given fWWN or the interface are displayed.

The following example shows how to display the port security configuration database:

```
switch# show port-security database
```

The following example shows how to display the port security configuration database for VSAN 1:

```
switch# show port-security database vsan 1
```

The following example shows how to display the activated database:

```
switch# show port-security database active
```

The following example shows how to display difference between the temporary configuration database and the configuration database:

```
switch# show port-security pending-diff vsan 1
```

The following example shows how to display the configured fWWN port security in VSAN 1:

```
switch# show port-security database fwwn 20:01:00:05:30:00:95:de vsan 1
20:00:00:0c:88:00:4a:e2(swwn)
```

The following example shows how to display the port security statistics:

```
switch# show port-security statistics
```

The following example shows how to verify the status of the active database and the auto-learning configuration:

```
switch# show port-security status
```

Default Port Security Settings

The following table lists the default settings for all port security features in any switch.

Table 39: Default Security Settings

Parameters	Default
Auto-learn	Enabled if port security is enabled.
Port security	Disabled.
Distribution	Disabled. **Note** Enabling distribution enables it on all VSANs in the switch.

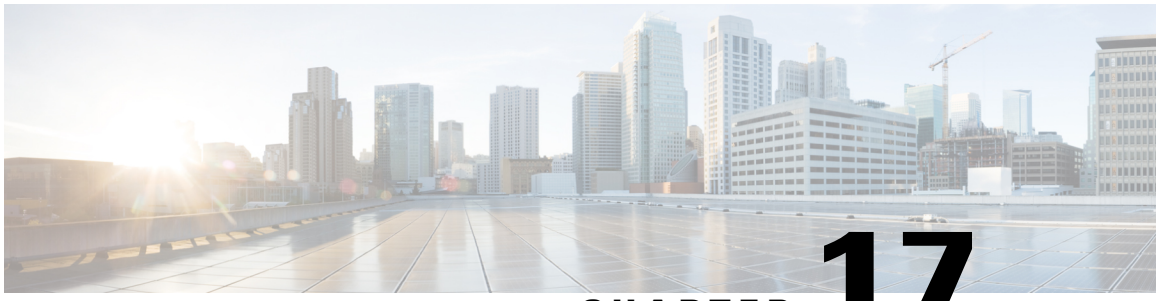

CHAPTER 17

Configuring Fabric Binding

This chapter contains the following sections:

Configuring Fabric Binding

Information About Fabric Binding

The fabric binding feature ensures that ISLs are only enabled between specified switches in the fabric. Fabric binding is configured on a per-VSAN basis.

This feature helps prevent unauthorized switches from joining the fabric or disrupting current fabric operations. It uses the Exchange Fabric Membership Data (EFMD) protocol to ensure that the list of authorized switches is identical in all switches in the fabric.

Licensing Requirements for Fabric Binding

Fabric Binding requires the Storage Protocol Services license.

Port Security Versus Fabric Binding

Port security and fabric binding are two independent features that can be configured to complement each other. The following table compares the two features.

Table 40: Fabric Binding and Port Security Comparison

Fabric Binding	Port Security
Uses a set of sWWNs and a persistent domain ID.	Uses pWWNs/nWWNs or fWWNs/sWWNs.
Binds the fabric at the switch level.	Binds devices at the interface level.

Fabric Binding	Port Security
Authorizes only the configured sWWN stored in the fabric binding database to participate in the fabric.	Allows a preconfigured set of Fibre Channel devices to logically connect to a SAN port. The switch port, identified by a WWN or interface number, connects to a Fibre Channel device (a host or another switch), also identified by a WWN. By binding these two devices, you lock these two ports into a group (or list).
Requires activation on a per VSAN basis.	Requires activation on a per VSAN basis.
Allows specific user-defined switches that are allowed to connect to the fabric, regardless of the physical port to which the peer switch is connected.	Allows specific user-defined physical ports to which another device can connect.
Does not learn about switches that are logging in.	Learns about switches or devices that are logging in if learning mode is enabled.
Cannot be distributed by CFS and must be configured manually on each switch in the fabric.	Can be distributed by CFS.

Port-level checking for xE ports is as follows:

- The switch login uses both port security binding and fabric binding for a given VSAN.
- Binding checks are performed on the port VSAN as follows:
  - E port security binding check on port VSAN
  - TE port security binding check on each allowed VSAN

While port security complements fabric binding, they are independent features and can be enabled or disabled separately.

Fabric Binding Enforcement

To enforce fabric binding, configure the switch world wide name (sWWN) to specify the xE port connection for each switch. Enforcement of fabric binding policies are done on every activation and when the port tries to come up. For a Fibre Channel VSAN, the fabric binding feature requires all sWWNs connected to a switch to be part of the fabric binding active database.

Configuring Fabric Binding

The fabric binding feature ensures ISLs are only enabled between specified switches in the fabric binding configuration. Fabric binding is configured on a per-VSAN basis.

Configuring Fabric Binding

To configure fabric binding in each switch in the fabric, perform this task:

SUMMARY STEPS

1. Enable the fabric configuration feature.
2. Configure a list of sWWNs and their corresponding domain IDs for devices that are allowed to access the fabric.
3. Activate the fabric binding database.
4. Copy the fabric binding active database to the fabric binding configuration database.
5. Save the fabric binding configuration.
6. Verify the fabric binding configuration.

DETAILED STEPS

Step 1 Enable the fabric configuration feature.

Step 2 Configure a list of sWWNs and their corresponding domain IDs for devices that are allowed to access the fabric.

Step 3 Activate the fabric binding database.

Step 4 Copy the fabric binding active database to the fabric binding configuration database.

Step 5 Save the fabric binding configuration.

Step 6 Verify the fabric binding configuration.

Enabling Fabric Binding

The fabric binding feature must be enabled in each switch in the fabric that participates in the fabric binding. By default, this feature is disabled in Cisco Nexus 5000 Series switches. The configuration and verification commands for the fabric binding feature are only available when fabric binding is enabled on a switch. When you disable this configuration, all related configurations are automatically discarded.

To enable fabric binding on any participating switch, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fabric-binding enable**
3. switch(config)# **no fabric-binding enable**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fabric-binding enable**	Enables fabric binding on that switch.
Step 3	switch(config)# **no fabric-binding enable**	Disables (default) fabric binding on that switch.

About Switch WWN Lists

A user-specified fabric binding list contains a list of switch WWNs (sWWNs) within a fabric. If an sWWN attempts to join the fabric, and that sWWN is not on the list or the sWWN is using a domain ID that differs from the one specified in the allowed list, the ISL between the switch and the fabric is automatically isolated in that VSAN and the switch is denied entry into the fabric.

Configuring Switch WWN List

To configure a list of sWWNs and optional domain IDs for a Fibre Channel VSAN, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fabric-binding database vsan** *vsan-id*
3. switch(config)# **no fabric-binding database vsan** *vsan-id*
4. switch(config-fabric-binding)#**swwn** *swwn-id* **domain** *domain-id*
5. switch(config-fabric-binding)#**no swwn** *swwn-id* **domain** *domain-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fabric-binding database vsan** *vsan-id*	Enters the fabric binding submode for the specified VSAN.
Step 3	switch(config)# **no fabric-binding database vsan** *vsan-id*	Deletes the fabric binding database for the specified VSAN.
Step 4	switch(config-fabric-binding)#**swwn** *swwn-id* **domain** *domain-id*	Adds the sWWN of another switch for a specific domain ID to the configured database list.
Step 5	switch(config-fabric-binding)#**no swwn** *swwn-id* **domain** *domain-id*	Deletes the sWWN and domain ID of a switch from the configured database list.

About Fabric Binding Activation and Deactivation

The fabric binding feature maintains a configuration database (config database) and an active database. The config database is a read-write database that collects the configurations you perform. These configurations are only enforced upon activation. This activation overwrites the active database with the contents of the config database. The active database is read-only and is the database that checks each switch that attempts to log in.

By default, the fabric binding feature is not activated. You cannot activate the fabric binding database on the switch if entries existing in the config database conflict with the current state of the fabric. For example, one of the already logged in switches may be denied login by the config database. You can choose to forcefully override these situations.

Note After activation, any already logged in switch that violates the current active database will be logged out, and all switches that were previously denied login because of fabric binding restrictions are reinitialized.

Activating Fabric Binding

To activate the fabric binding feature, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fabric-binding activate vsan** *vsan-id*
3. switch(config)# **no fabric-binding activate vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fabric-binding activate vsan** *vsan-id*	Activates the fabric binding database for the specified VSAN.
Step 3	switch(config)# **no fabric-binding activate vsan** *vsan-id*	Deactivates the fabric binding database for the specified VSAN.

Forcing Fabric Binding Activation

If the database activation is rejected due to one or more conflicts listed in the previous section, you may decide to proceed with the activation by using the force option.

To forcefully activate the fabric binding database, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fabric-binding activate vsan** *vsan-id* **force**
3. switch(config)# **no fabric-binding activate vsan** *vsan-id* **force**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fabric-binding activate vsan** *vsan-id* **force**	Activates the fabric binding database for the specified VSAN forcefully, even if the configuration is not acceptable.
Step 3	switch(config)# **no fabric-binding activate vsan** *vsan-id* **force**	Reverts to the previously configured state or to the factory default (if no state is configured).

Copying Fabric Binding Configurations

When you copy the fabric binding configuration, the config database is saved to the running configuration.

You can use the following commands to copy to the config database:

- Use the **fabric-binding database copy vsan** command to copy from the active database to the config database. If the configured database is empty, this command is not accepted.

  ```
  switch# fabric-binding database copy vsan 1
  ```

- Use the **fabric-binding database diff active vsan** command to view the differences between the active database and the config database. This command can be used when resolving conflicts.

  ```
  switch# fabric-binding database diff active vsan 1
  ```

- Use the **fabric-binding database diff config vsan** command to obtain information on the differences between the config database and the active database.

  ```
  switch# fabric-binding database diff config vsan 1
  ```

- Use the **copy running-config startup-config** command to save the running configuration to the startup configuration so that the fabric binding config database is available after a reboot.

  ```
  switch# copy running-config startup-config
  ```

Clearing the Fabric Binding Statistics

Use the **clear fabric-binding statistics** command to clear all existing statistics from the fabric binding database for a specified VSAN.

```
switch# clear fabric-binding statistics vsan 1
```

Deleting the Fabric Binding Database

Use the **no fabric-binding** command in configuration mode to delete the configured database for a specified VSAN.

```
switch(config)# no fabric-binding database vsan 10
```

Verifying Fabric Binding Information

To display fabric binding information, perform one of the following tasks

SUMMARY STEPS

1. switch# **show fabric-binding database** [**active**]
2. switch# **show fabric-binding database** [**active**] [**vsan** *vsan-id*]
3. switch# **show fabric-binding statistics**
4. switch# **show fabric-binding status**
5. switch# **show fabric-binding violations**
6. switch# **show fabric-binding efmd** [**vsan** *vsan-id*]

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **show fabric-binding database** [**active**]	Displays the configured fabric binding database. Include keyword active to display only the active fabric binding database.
Step 2	switch# **show fabric-binding database** [**active**] [**vsan** *vsan-id*]	Displays the configured fabric binding database for the specified VSAN.
Step 3	switch# **show fabric-binding statistics**	Displays statistics for the fabric binding database.
Step 4	switch# **show fabric-binding status**	Displays fabric binding status for all VSANs.
Step 5	switch# **show fabric-binding violations**	Displays fabric binding violations.
Step 6	switch# **show fabric-binding efmd** [**vsan** *vsan-id*]	Displays the configured fabric binding database for the specified VSAN.

The following example displays the active fabric binding information for VSAN 4:

```
switch# show fabric-binding database active vsan 4
```

The following example displays fabric binding violations:

```
switch# show fabric-binding violations
--------------------------------------------------------------------------------

VSAN Switch WWN [domain]     Last-Time             [Repeat count] Reason
```

```
--------------------------------------------------------------------------------
2    20:00:00:05:30:00:4a:1e [0xeb] Nov 25 05:46:14 2003   [2]    Domain mismatch
3    20:00:00:05:30:00:4a:1e [*] Nov 25 05:44:58 2003      [2]    sWWN not found
4    20:00:00:05:30:00:4a:1e [*] Nov 25 05:46:25 2003      [1]    Database mismatch
```

Note In VSAN 3, the sWWN was not found in the list. In VSAN 2, the sWWN was found in the list, but has a domain ID mismatch.

The following example displays EFMD Statistics for VSAN 4:

```
switch# show fabric-binding efmd statistics vsan 4
```

Default Fabric Binding Settings

The following table lists the default settings for the fabric binding feature.

Table 41: Default Fabric Binding Settings

Parameters	Default
Fabric binding	Disabled

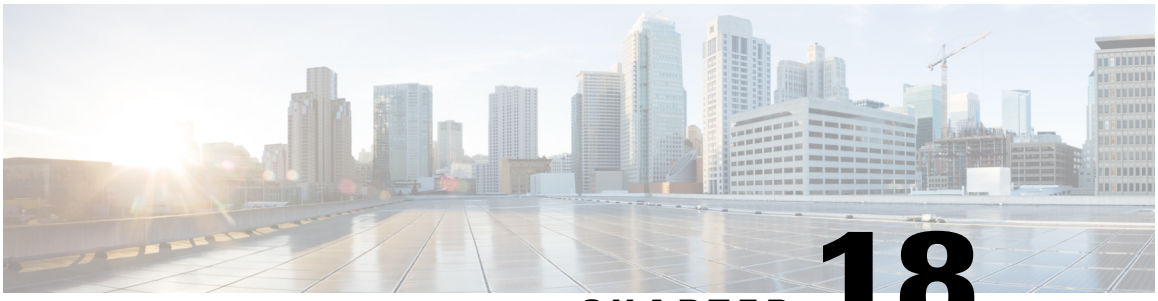

CHAPTER 18

Configuring Fabric Configuration Servers

This chapter contains the following sections:

Configuring Fabric Configuration Servers

Information About FCS

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. The FCS views the entire fabric based on the following objects:

- Interconnect element (IE) object—Each switch in the fabric corresponds to an IE object. One or more IE objects form a fabric.
- Port object—Each physical port in an IE corresponds to a port object. This includes the switch ports (xE and F ports) and their attached N ports.
- Platform object—A set of nodes may be defined as a platform object to make it a single manageable entity. These nodes are end-devices (host systems, storage subsystems) attached to the fabric. Platform objects reside at the edge switches of the fabric.

Each object has its own set of attributes and values. A null value may also be defined for some attributes.

In the Cisco Nexus 5000 Series switch environment, a fabric may consist of multiple VSANs. One instance of the FCS is present per VSAN.

FCS supports the discovery of virtual devices. The **fcs virtual-device-add** command, entered in FCS configuration submode, allows you to discover virtual devices in a particular VSAN or in all VSANs.

If you have attached a management application to a switch, all the frames directed towards the FCS in the switch are part of the port VSAN in the switch port (F port). Your view of the management application is limited only to this VSAN. However, information about other VSANs that this switch is part of can be obtained either through the SNMP or CLI.

In the following figure, Management Application 1 (M1) is connected through an F port with port VSAN ID 1, and Management Application 2 (M2) is connected through an F port with port VSAN ID 2. M1 can query

the FCS information of switches S1 and S3, and M2 can query switches S3 and S4. Switch S2 information is not known to both of them. FCS operations can be done only on those switches that are visible in the VSAN. M2 can send FCS requests only for VSAN 2 even though S3 is also a part of VSAN 1.

Figure 34: FCSs in a VSAN Environment

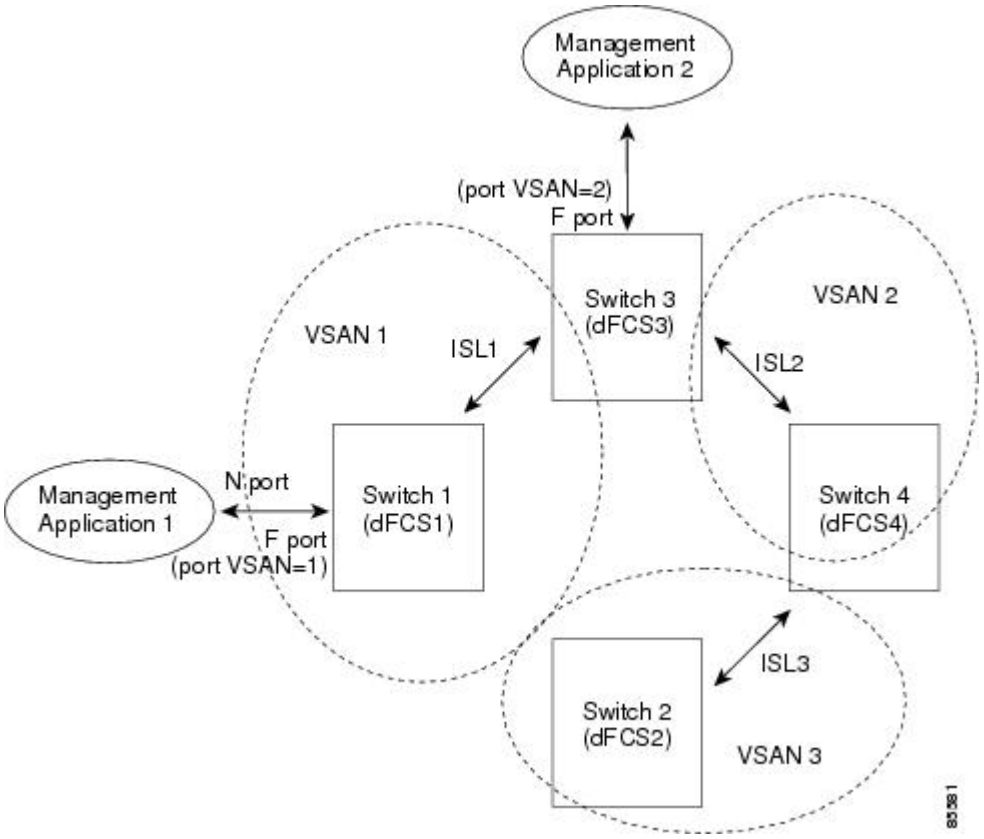

FCS Characteristics

FCSs have the following characteristics:

- Support network management including the following:
  - N port management application can query and obtain information about fabric elements.
  - SNMP manager can use the FCS management information base (MIB) to start discovery and obtain information about the fabric topology.
- Support TE ports in addition to the standard F and E ports.
- Can maintain a group of nodes with a logical name and management address when a platform registers with it. FCSs maintain a backup of all registrations in secondary storage and update it with every change. When a restart or switchover happens, FCSs retrieve the secondary storage information and rebuild its database.
- SNMP manager can query FCSs for all IEs, ports, and platforms in the fabric.

FCS Name Specification

You can specify if the unique name verification is for the entire fabric (globally) or only for locally (default) registered platforms.

Note Set this command globally only if every switch in the fabric belong to the Cisco MDS 9000 Family or Cisco Nexus 5000 Series of switches.

To enable global checking of the platform name, perform this task:

To register platform attributes, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **fcs plat-check-global vsan** *vsan-id*
3. switch(config)# **no fcs plat-check-global vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fcs plat-check-global vsan** *vsan-id*	Enables global checking of the platform name.
Step 3	switch(config)# **no fcs plat-check-global vsan** *vsan-id*	Disables (default) global checking of the platform name.

Displaying FCS Information

You can use the **show fcs** commands to display the status of the WWN configuration.

The following example shows how to display the FCS local database:

```
switch# show fcs database
```

The following example shows how to display a list of all interconnect elements for VSAN 1:

```
switch# show fcs ie vsan 1
```

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

```
switch# show fcs platform name SamplePlatform vsan 1
```

The following example shows how to display port information for a specific pWWN:

```
switch# show fcs port pwwn 20:51:00:05:30:00:16:de vsan 24
```

Default FCS Settings

The following table lists the default FCS settings.

Table 42: Default FCS Settings

Parameters	Default
Global checking of the platform name	Disabled
Platform node type	Unknown

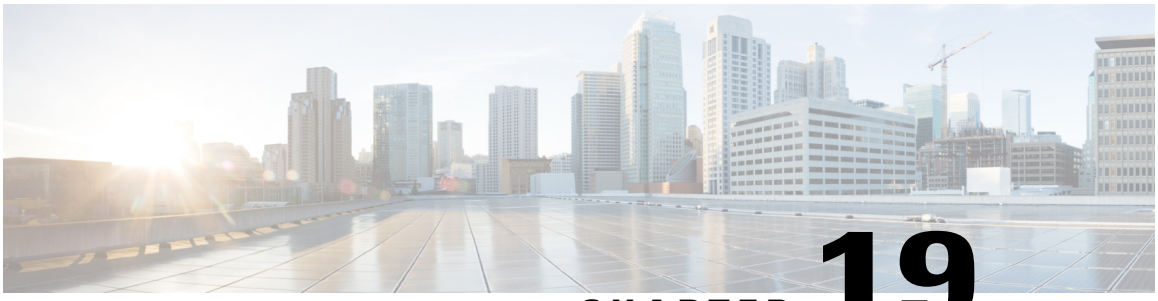

CHAPTER 19

Configuring Port Tracking

This chapter contains the following sections:

Configuring Port Tracking

Cisco Nexus 5000 Series switches offer the port tracking feature on physical Fibre Channel interfaces (but not on virtual Fibre Channel interfaces). This feature uses information about the operational state of the link to initiate a failure in the link that connects the edge device. This process of converting the indirect failure to a direct failure triggers a faster recovery process towards redundant links. When enabled, the port tracking feature brings down the configured links based on the failed link and forces the traffic to be redirected to another redundant link.

Information About Port Tracking

Generally, hosts can instantly recover from a link failure on a link that is immediately (direct link) connected to a switch. However, recovering from an indirect link failure between switches in a WAN or MAN fabric with a keepalive mechanism is dependent on several factors such as the timeout values (TOVs) and on registered state change notification (RSCN) information.

In the following figure, when the direct link 1 to the host fails, recovery can be immediate. However, when the ISL 2 fails between the two switches, recovery depends on TOVs, RSCNs, and other factors.

Figure 35: Traffic Recovery Using Port Tracking

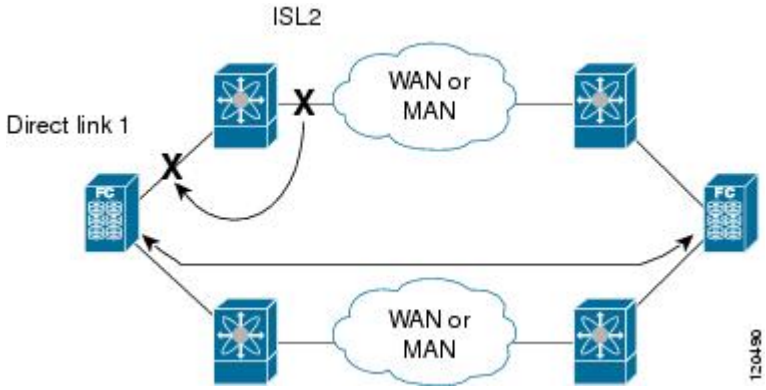

The port tracking feature monitors and detects failures that cause topology changes and brings down the links connecting the attached devices. When you enable this feature and explicitly configure the linked and tracked ports, the switch software monitors the tracked ports and alters the operational state of the linked ports on detecting a link state change.

The following terms are used in this chapter:

- Tracked ports—A port whose operational state is continuously monitored. The operational state of the tracked port is used to alter the operational state of one or more ports. Fibre Channel, VSAN, SAN port channel, or a Gigabit Ethernet port can be tracked. Generally, ports in E and TE port modes can also be F ports.
- Linked ports—A port whose operational state is altered based on the operational state of the tracked ports. Only physical Fibre Channel ports can be linked ports.

Port tracking has the following features:

- The application brings the linked port down when the tracked port goes down. When the tracked port recovers from the failure and comes back up again, the linked port is also brought up automatically (unless otherwise configured).
- You can forcefully continue to keep the linked port down, even though the tracked port comes back up. In this case, you must explicitly bring up the linked port when required.

Configuring Port Tracking

Before configuring port tracking, consider the following guidelines:

- Verify that the tracked ports and the linked ports are on the same Cisco switch.
- Be aware that the linked port is automatically brought down when the tracked port goes down.
- Do not track a linked port back to itself (for example, Port fc2/2 to Port fc2/4 and back to Port fc2/2) to avoid recursive dependency.

Enabling Port Tracking

The port tracking feature is disabled by default in Cisco Nexus 5000 Series switches. When you enable this feature, port tracking is globally enabled for the entire switch.

To configure port tracking, enable the port tracking feature and configure the linked ports for the tracked port.

To enable port tracking, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **port-track enable**
3. switch(config)# **no port-track enable**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **port-track enable**	Enables port tracking.
Step 3	switch(config)# **no port-track enable**	Removes the currently applied port tracking configuration and disables port tracking.

About Configuring Linked Ports

You can link ports using one of two methods:

- Operationally binding the linked ports to the tracked port (default).
- Continuing to keep the linked port down forcefully, even if the tracked port has recovered from the link failure.

Operationally Binding a Tracked Port

When you configure the first tracked port, operational binding is automatically in effect. When you use this method, you have the option to monitor multiple ports or monitor ports in one VSAN.

To operationally bind a tracked port, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface fc** *slot*/*port*
3. switch(config-if)# **port-track interface fc** *slot*/*port* | **san-port-channel** *port*
4. switch(config-if)# **no port-track interface fc** *slot*/*port* | **san-port-channel** *port*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **interface fc** *slot*/*port*	Enters the interface configuration mode for the linked port. You can now configure the tracked ports.
Step 3	switch(config-if)# **port-track interface fc** *slot*/*port* | **san-port-channel** *port*	Specifies the tracked port. When the tracked port goes down, the linked port is also brought down.
Step 4	switch(config-if)# **no port-track interface fc** *slot*/*port* | **san-port-channel** *port*	Removes the port tracking configuration that is currently applied to the interface.

About Tracking Multiple Ports

You can control the operational state of the linked port based on the operational states of multiple tracked ports. When more than one tracked port is associated with a linked port, the operational state of the linked port will be set to down only if all the associated tracked ports are down. Even if one tracked port is up, the linked port will stay up.

In the following figure, only if both ISLs 2 and 3 fail, will the direct link 1 be brought down. Direct link 1 will not be brought down if either 2 or 3 are still functioning as desired.

Figure 36: Traffic Recovery Using Port Tracking

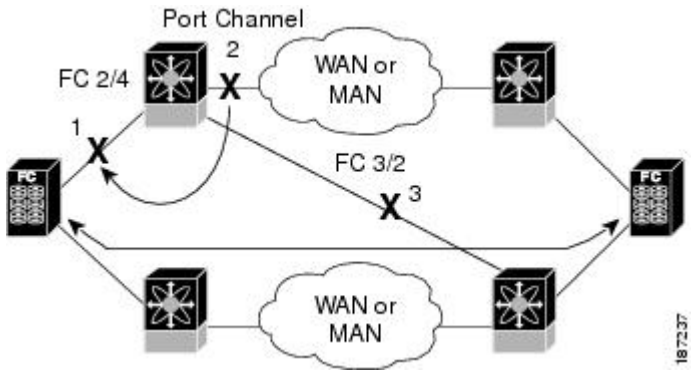

Tracking Multiple Ports

To track multiple ports, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface fc** *slot*/*port*
3. switch(config-if)# **port-track interface interface fc** *slot*/*port* | **san-port-channel** *port*

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 the interface configuration mode. You can now configure tracked ports.
Step 3	switch(config-if)# **port-track interface interface fc** *slot*/*port*	**san-port-channel** *port*	Tracks the linked port with the specified interface. When the tracked port goes down, the linked port is also brought down.

About Monitoring Ports in a VSAN

You can optionally configure one VSAN from the set of all operational VSANs on the tracked port with the linked port by specifying the required VSAN. This level of flexibility provides higher granularity in tracked ports. In some cases, when a tracked port is a TE port, the set of operational VSANs on the port can change dynamically without bringing down the operational state of the port. In such cases, the port VSAN of the linked port can be monitored on the set of operational VSANs on the tracked port.

If you configure this feature, the linked port is up only when the VSAN is up on the tracked port.

The specified VSAN does not have to be the same as the port VSAN of the linked port.

Monitoring Ports in a VSAN

To monitor a tracked port in a specific VSAN, perform this task :

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface fc** *slot*/*port*
3. switch(config-if)# **port-track interface san-port-channel 1 vsan 2**
4. switch(config-if)# **no port-track interface san-port-channel 1 vsan 2**

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 the interface configuration mode. You can now configure tracked ports.
Step 3	switch(config-if)# **port-track interface san-port-channel 1 vsan 2**	Enables tracking of the SAN port channel in VSAN 2.
Step 4	switch(config-if)# **no port-track interface san-port-channel 1 vsan 2**	Removes the VSAN association for the linked port. The SAN port channel link remains in effect.

About Forceful Shutdown

If a tracked port flaps frequently, then tracking ports using the operational binding feature may cause frequent topology change. In this case, you may choose to keep the port in the down state until you are able to resolve the reason for these frequent flaps. Keeping the flapping port in the down state forces the traffic to flow through the redundant path until the primary tracked port problems are resolved. When the problems are resolved and the tracked port is back up, you can explicitly enable the interface.

If you configure this feature, the linked port continues to remain in the shutdown state even after the tracked port comes back up. You must explicitly remove the forced shut state (by administratively bringing up this interface) of the linked port once the tracked port is up and stable.

Forcefully Shutting Down a Tracked Port

To forcefully shut down a tracked port, perform this task:

SUMMARY STEPS

1. switch# **configuration terminal**
2. switch(config)# **interface fc** *slot*/*port*
3. switch(config-if)# **port-track force-shut**
4. switch(config-if)# **no port-track force-shut**

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 the interface configuration mode. You can now configure tracked ports.

	Command or Action	Purpose
Step 3	switch(config-if)# **port-track force-shut**	Forcefully shuts down the tracked port.
Step 4	switch(config-if)# **no port-track force-shut**	Removes the port shutdown configuration for the tracked port.

Displaying Port Tracking Information

The **show** commands display the current port tracking settings for the switch.

The following example shows how to display tracked port configuration for a specific interface:

```
switch# show interface fc2/1
fc2/1 is down (Administratively down)
    Hardware is Fibre Channel, FCOT is short wave laser w/o OFC (SN)
    Port WWN is 20:01:00:05:30:00:0d:de
    Admin port mode is FX
    Port vsan is 1
    Receive data field Size is 2112
    Beacon is turned off
    Port tracked with interface fc2/2 (down)
    Port tracked with interface san-port-channel 1 vsan 2 (down)
5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
...
```

The following example shows how to display tracked port configuration for a SAN port channel:

```
switch# show interface san-port-channel 1
port-channel 1 is down (No operational members)
    Hardware is Fibre Channel
    Port WWN is 24:01:00:05:30:00:0d:de
    Admin port mode is auto, trunk mode is on
    Port vsan is 2
    Linked to 1 port(s)
      Port linked to interface fc2/1
...
```

The following example shows how to display the port track mode:

```
switch# show interface fc 2/4
fc2/4 is up
    Hardware is Fibre Channel, FCOT is short wave laser
...
    Transmit B2B Credit is 64
    Receive B2B Credit is 16
    Receive data field Size is 2112
    Beacon is turned off
    Port track mode is force_shut <-- this port remains shut even if the tracked port is
back up
```

Default Port Tracking Settings

The following table lists the default settings for port tracking parameters.

Table 43: Default Port Tracking Parameters

Parameters	Default
Port tracking	Disabled
Operational binding	Enabled along with port tracking

INDEX

E

F

H

I

L

M

N

O

P

R

S

T

U

V

W

Z