Cisco Nexus 6000 Series NX-OS SAN Switching Configuration Guide, Release 6.x

First Published: 2013-01-30

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-27932-01

CONTENTS

Preface

The preface contains the following sections:

Audience

This publication is for network administrators who configure and maintain Cisco Nexus devices.

Document Conventions

Note As part of our constant endeavor to remodel our documents to meet our customers' requirements, we have modified the manner in which we document configuration tasks. As a result of this, you may find a deviation in the style used to describe these tasks, with the newly included sections of the document following the new format.

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).

Convention	Description
[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.

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 Cisco Nexus 6000 Series NX-OS Software

The entire Cisco NX-OS 6000 Series documentation set is available at the following URL:

http://www.cisco.com/en/US/products/ps12806/tsd_products_support_series_home.html

Release Notes

The release notes are available at the following URL:

http://www.cisco.com/c/en/us/support/switches/nexus-6000-series-switches/products-release-notes-list.html

Configuration Guides

These guides are available at the following URL:

http://www.cisco.com/c/en/us/support/switches/nexus-6000-series-switches/products-installation-and-configuration-guides-list.html

The documents in this category include:

- *Cisco Nexus 6000 Series NX-OS Adapter-FEX Configuration Guide*
- *Cisco Nexus 6000 Series NX-OS FabricPath Configuration Guide*
- *Cisco Nexus 6000 Series NX-OS FCoE Configuration Guide*
- *Cisco Nexus 6000 Series NX-OS Fundamentals Configuration Guide*
- *Cisco Nexus 6000 Series NX-OS Interfaces Configuration Guide*
- *Cisco Nexus 6000 Series NX-OS Layer 2 Switching Configuration Guide*
- *Cisco Nexus 6000 Series NX-OS Multicast Routing Configuration Guide*
- *Cisco Nexus 6000 Series NX-OS Quality of Service Configuration Guide*
- *Cisco Nexus 6000 Series NX-OS SAN Switching Configuration Guide*
- *Cisco Nexus 6000 Series NX-OS Security Configuration Guide*
- *Cisco Nexus 6000 Series NX-OS System Management Configuration Guide*
- *Cisco Nexus 6000 Series NX-OS Unicast Routing Configuration Guide*

Installation and Upgrade Guides

These guides are available at the following URL:

http://www.cisco.com/c/en/us/support/switches/nexus-6000-series-switches/products-installation-guides-list.html

The document in this category include:

- *Cisco Nexus 6000 Series NX-OS Software Upgrade and Downgrade Guides*

Licensing Guide

The *License and Copyright Information for Cisco NX-OS Software* is available at http://www.cisco.com/en/US/docs/switches/datacenter/sw/4_0/nx-os/license_agreement/nx-ossw_lisns.html.

Command References

These guides are available at the following URL:

http://www.cisco.com/c/en/us/support/switches/nexus-6000-series-switches/products-command-reference-list.html

The documents in this category include:

- *Cisco Nexus 6000 Series NX-OS Fabric Extender Command Reference*
- *Cisco Nexus 6000 Series NX-OS FabricPath Command Reference*
- *Cisco Nexus 6000 Series NX-OS Fundamentals Command Reference*
- *Cisco Nexus 6000 Series NX-OS Interfaces Command Reference*
- *Cisco Nexus 6000 Series NX-OS Layer 2 Interfaces Command Reference*
- *Cisco Nexus 6000 Series NX-OS Multicast Routing Command Reference*
- *Cisco Nexus 6000 Series NX-OS Quality of Service Command Reference*
- *Cisco Nexus 6000 Series NX-OS Security Command Reference*
- *Cisco Nexus 6000 Series NX-OS System Management Command Reference*
- *Cisco Nexus 6000 Series NX-OS TrustSec Command Reference*
- *Cisco Nexus 6000 Series NX-OS Unicast Routing Command Reference*
- *Cisco Nexus 6000 Series NX-OS Virtual Port Channel Command Reference*

Technical References

The *Cisco Nexus 6000 Series NX-OS MIB Reference* is available at http://www.cisco.com/en/US/docs/switches/datacenter/nexus6000/sw/mib/reference/NX6000_MIBRef.html.

Error and System Messages

The *Cisco Nexus 6000 Series NX-OS System Message Guide* is available at http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus6000/sw/system_messages/reference/sl_nxos_book.html.

Troubleshooting Guide

The *Cisco Nexus 6000 Series NX-OS Troubleshooting Guide* is available at http://www.cisco.com/c/en/us/support/switches/nexus-6000-series-switches/tsd-products-support-troubleshoot-and-alerts.html.

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to: .

We appreciate your feedback.

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, using the Cisco Bug Search Tool (BST), submitting a service request, and gathering additional information, see What's New in Cisco Product Documentation.

To receive new and revised Cisco technical content directly to your desktop, you can subscribe to the What's New in Cisco Product Documentation RSS feed. RSS feeds are a free service.

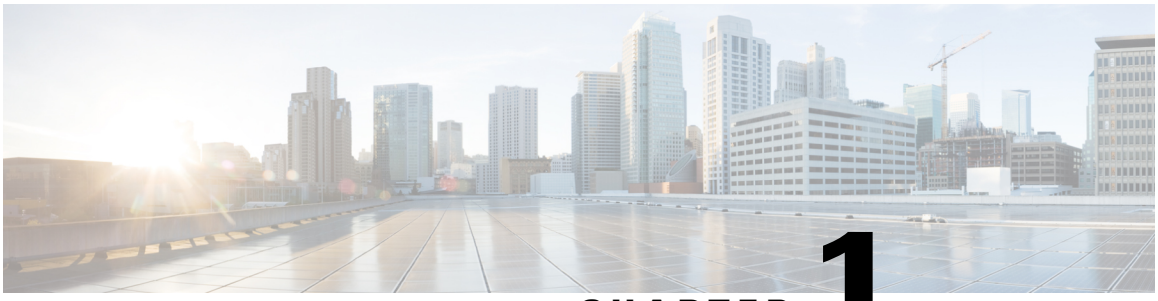

CHAPTER 1

Overview

This chapter contains the following sections:

SAN Switching Overview

This chapter provides an overview of SAN switching for Cisco NX-OS devices. This chapter includes the following sections:

Domain Parameters

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

N Port 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.

VSAN Trunking

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

SAN Port Channels

PortChannels aggregate multiple physical ISLs into one logical link with higher bandwidth and port resiliency for Fibre Channel 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 SANs (VSANs) partition a single physical SAN into multiple VSANs. VSANs allow the Cisco NX-OS software to logically divide a large physical fabric into separate, isolated environments to improve Fibre Channel SAN scalability, availability, manageability, and network security.

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

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

VSANs are supported across Fibre Channel over IP (FCIP) links between SANs, which extends VSANs to include devices at a remote location. The Cisco SAN 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. The Cisco NX-OS software supports the following types of zoning:

- N port zoning-Defines zone members based on the end-device (host and storage) port.
  - WWN
  - Fibre Channel identifier (FC-ID)
- Fx port zoning-Defines zone members based on the switch port.
  - WWN
  - WWN plus the interface index, or domain ID plus the 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, logical unit number (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 useful for sharing volumes across servers for backup, data warehousing, and so on.
- 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 the 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

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

Fibre Channel Routing

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

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

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 device.

Advanced Fibre Channel Features

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

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

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

FC-SP and DHCHAP

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

With FC-SP, switches, storage devices, and hosts can prove their identity through a reliable and manageable authentication mechanism. With FC-SP, Fibre Channel traffic can be secured per frame 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

Fabric binding ensures Inter-Switch Links (ISLs) are enabled only between specified switches in the fabric binding configuration, which prevents 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. Multiple VSANs constitute a fabric, where one instance of the FCS is present per VSAN.

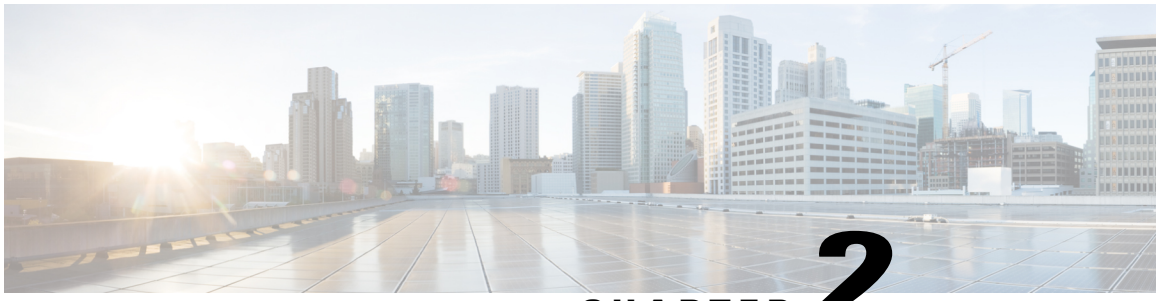

CHAPTER 2

Configuring Fibre Channel Domain Parameters

This chapter describes how to configure Fibre Channel domain parameters.

This chapter includes the following sections:

Information About Domain Parameters

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

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

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

Fibre Channel Domains

The fcdomain has four phases:

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

The following figure shows an example fcdomain configuration.

Figure 1: Sample fcdomain Configuration

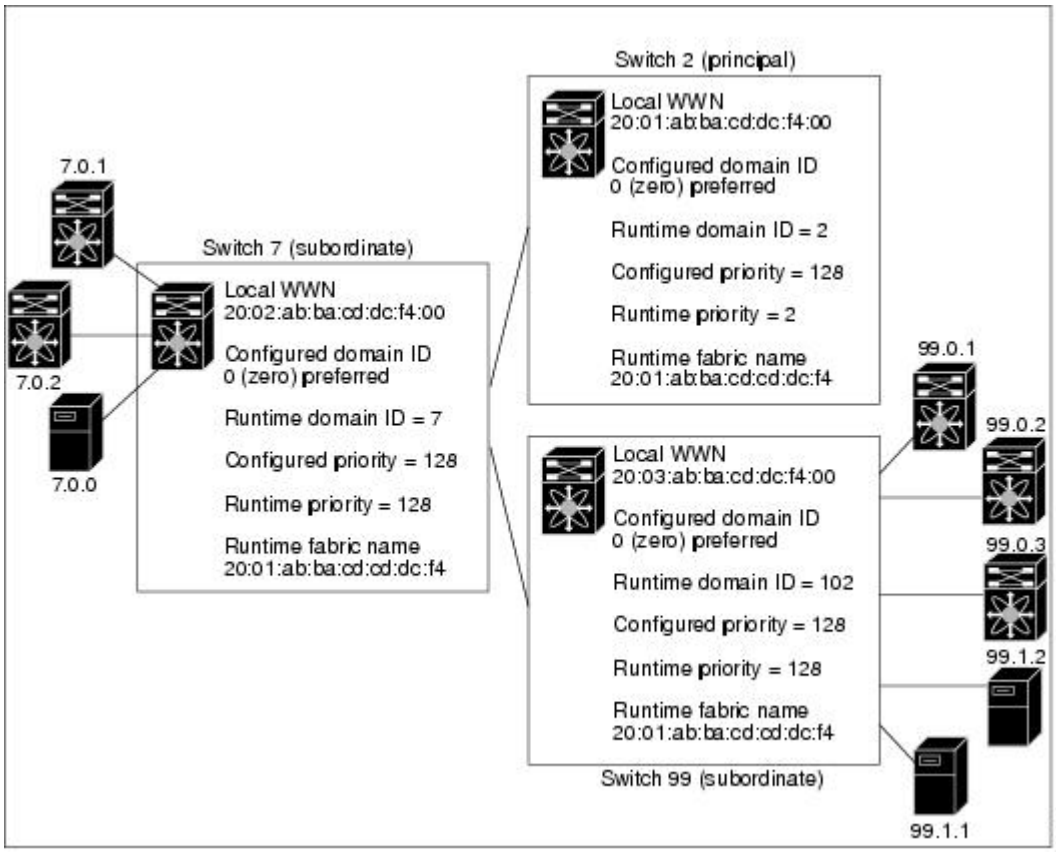

Domain Restarts

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

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

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

You can restart the fabric disruptively or nondisruptively.

Procedure

	Command or Action	Purpose

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

Step 2	**fcdomain restart vsan** *vsan-id* **Example:** `switch (config)# fcdomain restart vsan 100`	Forces the VSAN to reconfigure without traffic disruption. The VSAN ID ranges from 1 to 4093.
Step 3	switch(config)# **fcdomain restart disruptive vsan** *vsan-id* **Example:** `switch (config)# fcdomain restart disruptive vsan 101`	Forces the VSAN to reconfigure with data traffic disruption.

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

You can enable the domain manager fast restart feature.

Procedure

	Command or Action	Purpose

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

Step 2	**fcdomain optimize fast-restart vsan** *vsan-id* **Example:** `switch(config)# fcdomain optimize fast-restart vsan 1`	Enables domain manager fast restart in the specified VSAN. The VSAN ID range is from 1 to 4093.
Step 3	**no fcdomain optimize fast-restart vsan** *vsan-id* **Example:** `switch(config)# no fcdomain optimize fast-restart vsan 1`	Disables (default) domain manager fast restart in the specified VSAN. The VSAN ID range is from 1 to 4093.

Switch Priority

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

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

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

Configuring Switch Priority

You can configure the priority for the principal switch.

Procedure

	Command or Action	Purpose

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

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

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:

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global 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

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

Procedure

	Command or Action	Purpose

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

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

Incoming RCFs

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

The rcf-reject option takes effect immediately.

No fcdomain restart is required.

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

Rejecting Incoming RCFs

You can reject incoming RCF request frames.

Procedure

	Command or Action	Purpose

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

	Command or Action	Purpose
Step 2	switch(config)# **interface vfc** *vfc-id*	Configures the specified interface.
Step 3	**fcdomain rcf-reject vsan** *vsan-id* **Example:** `switch(config-if)# fcdomain rcf-reject vsan 10`	Enables the RCF filter on the specified interface in the specified VSAN. The VSAN ID ranges from 1 to 4093.
Step 4	**no fcdomain rcf-reject vsan** *vsan-id* **Example:** `switch(config-if)# no fcdomain rcf-reject vsan 10`	Disables (default) the RCF filter on the specified interface in the specified VSAN. The VSAN ID ranges from 1 to 4093.

Autoreconfiguring Merged Fabrics

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

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

The autoreconfigure option takes immediate effect at runtime. You do not need to restart the fcdomain. If a domain is currently isolated due to domain overlap, and you later enable the autoreconfigure option on both switches, the fabric continues to be isolated. If you enabled the autoreconfigure option on both switches before connecting the fabric, a disruptive reconfiguration (RCF) will occur. A disruptive reconfiguration can affect data traffic. You can nondisruptively reconfigure the fcdomain by changing the configured domains on the overlapping links and eliminating the domain overlap.

Enabling Autoreconfiguration

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

Procedure

	Command or Action	Purpose

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

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

Domain IDs

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

Domain IDs - Guidelines

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

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

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

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

- The local switch sends a configured domain ID request to the principal switch.

- The principal switch assigns the requested domain ID if available. Otherwise, it assigns another available domain ID.

Figure 2: Configuration Process Using the Preferred Option

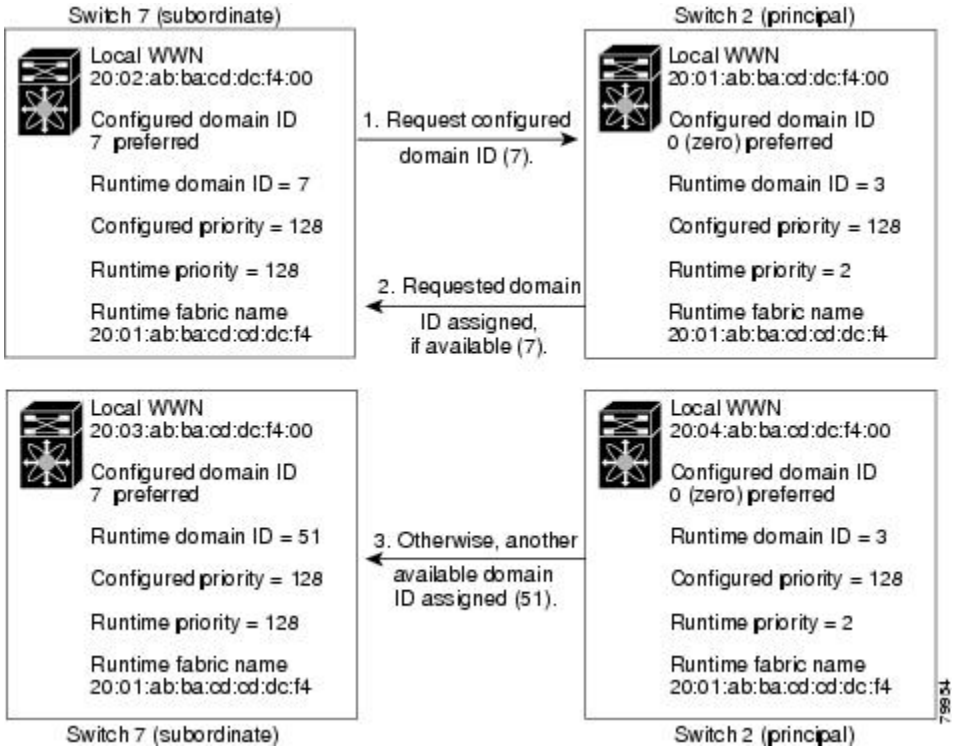

The operation of a subordinate switch changes based on three factors:

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

In specific situations, the changes are as follows:

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

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

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

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

Related Topics

Allowed Domain ID Lists, on page 15

Configuring Static or Preferred Domain IDs

You can specify a static or preferred domain ID.

Procedure

	Command or Action	Purpose

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

Step 2	**fcdomain domain** *domain-id* **static vsan** *vsan-id* **Example:** `switch(config)# fcdomain domain 1 static vsan 3`	Configures the switch in the specified VSAN to accept only a specific value and moves the local interfaces in the specified VSAN to an isolated state if the requested domain ID is not granted. The domain ID range is 1 to 239. The VSAN ID range is 1 to 4093.
Step 3	**no fcdomain domain** *domain-id* **static vsan** *vsan-id* **Example:** `switch(config)# no fcdomain domain 1 static vsan 3`	Resets the configured domain ID to factory defaults in the specified VSAN. The configured domain ID becomes 0 preferred.
Step 4	**fcdomain domain** *domain-id* **preferred vsan** *vsan-id* **Example:** `switch(config)# fcdomain domain 1 preferred vsan 5`	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 ID range is 1 to 239. The VSAN ID range is 1 to 4093.

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

Allowed Domain ID Lists

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

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

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

Configuring Allowed Domain ID Lists

You can configure the allowed domain ID list.

Procedure

	Command or Action	Purpose

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

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

CFS Distribution of Allowed Domain ID Lists

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

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

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

For additional information, refer to Using Cisco Fabric Services in the System Management Configuration Guide for your device.

Enabling Distribution

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

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

Procedure

	Command or Action	Purpose

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

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

Locking the Fabric

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

- No other user can make any configuration changes to this feature.

- A pending configuration is created by copying the active configuration. Subsequent modifications are made to the pending configuration and remain there until you commit the changes to the active configuration (and other switches in the fabric) or discard them.

Committing Changes

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

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

Procedure

	Command or Action	Purpose

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

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

Discarding Changes

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

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

Procedure

	Command or Action	Purpose

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

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

Clearing a Fabric Lock

If you have performed a domain configuration task and have not released the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your pending changes are discarded and the fabric lock is released.

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

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

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

Displaying CFS Distribution Status

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

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

Displaying Pending Changes

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

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

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

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

Displaying Session Status

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

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

Contiguous Domain ID Assignments

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

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

Enabling Contiguous Domain ID Assignments

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

Procedure

	Command or Action	Purpose

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

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

FC IDs

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

- An N port logs into a SAN switch. The WWN of the requesting N port and the assigned FC ID are retained and stored in a volatile cache. The contents of this volatile cache are not saved across reboots.

- The switch is designed to preserve the binding FC ID to the WWN on a best-effort basis. For example, if one N port disconnects from the switch and its FC ID is requested by another device, this request is granted and the WWN with the initial FC ID association is released.
- The volatile cache stores up to 4000 entries of WWN to FC ID binding. If this cache is full, a new (more recent) entry overwrites the oldest entry in the cache. In this case, the corresponding WWN to FC ID association for the oldest entry is lost.
- N ports receive the same FC IDs if disconnected and reconnected to any port within the same switch (as long as it belongs to the same VSAN).

Persistent FC IDs

When persistent FC IDs are enabled, the following occurs:

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

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

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

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

Enabling the Persistent FC ID Feature

You can enable the persistent FC ID feature.

Procedure

	Command or Action	Purpose

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

Step 2	**fcdomain fcid persistent vsan** *vsan-id* **Example:** `switch(config)# fcdomain fcid persistent vsan 78`	Activates (default) persistency of FC IDs in the specified VSAN.

	Command or Action	Purpose
Step 3	**no fcdomain fcid persistent vsan** *vsan-id* **Example:** `switch(config)# no fcdomain fcid persistent vsan 33`	Disables the FC ID persistency feature in the specified VSAN.

Persistent FC ID Configuration Guidelines

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

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

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

Configuring Persistent FC IDs

You can configure persistent FC IDs.

Procedure

	Command or Action	Purpose

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

Step 2	**fcdomain fcid database** **Example:** `switch(config)# fcdomain fcid database`	Enters FC ID database configuration submode.
Step 3	**vsan** *vsan-id* **wwn 33:e8:00:05:30:00:16:df fcid** *fcid* **Example:** `switch(config-fcid-db)# vsan 26 wwn 33:e8:00:05:30:00:16:df fcid 4`	Configures a device WWN (33:e8:00:05:30:00:16:df) with the FC ID 0x070128 in the specified VSAN.

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

Unique Area FC IDs for HBAs

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

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

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

Configuring Unique Area FC IDs for an HBA

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

The following task uses an example configuration with a switch domain of 111(6f hex). The server connects to the switch over FCoE. The HBA port connects to interface vfc20 and the storage port connects to interface fc2/3 on the same switch.

Procedure

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

```
switch# show flogi database
```

```
--------------------------------------------------------------------
INTERFACE VSAN   FCID           PORT NAME                NODE NAME
--------------------------------------------------------------------
 vfc20     3    0x6f7703  50:05:08:b2:00:71:c8:c2   50:05:08:b2:00:71:c8:c0
 vfc23     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 SAN switch.

```
switch# configure terminal
switch(config)# interface vfc 20

switch(config-if)# shutdown

switch(config-if)# end
```

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

```
switch# show fcdomain vsan 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 SAN switch.

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

Step 6 Enable the HBA interface in the SAN switch.

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

switch(config-if)# end
```

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

```
switch# show flogi database

--------------------------------------------------------------------
INTERFACE VSAN   FCID       PORT NAME                NODE NAME
--------------------------------------------------------------------
 vfc20     3    0x6fee00    50:05:08:b2:00:71:c8:c2   50:05:08:b2:00:71:c8:c0
 vfc23     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.

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 1: Purged FC IDs

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

Purging Persistent FC IDs

You can purge persistent FC IDs.

Procedure

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

Verifying the fcdomain Configuration

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

This example shows how to display information about fcdomain configurations:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Default Settings for Fibre Channel Domains

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

Table 2: 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

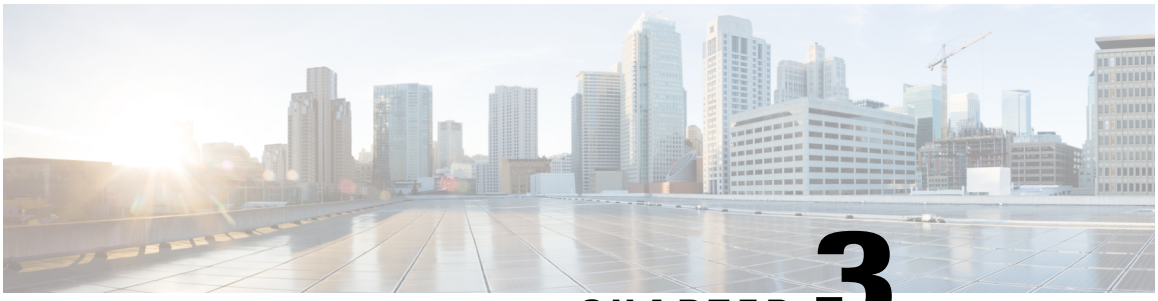

CHAPTER 3

Configuring N Port Virtualization

This chapter contains the following sections:

Configuring N Port Virtualization

Information About NPV

NPV Overview

By default, Cisco Nexus devices 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 following figure shows an interface-level view of an NPV configuration.

Figure 3: NPV Interface Configuration

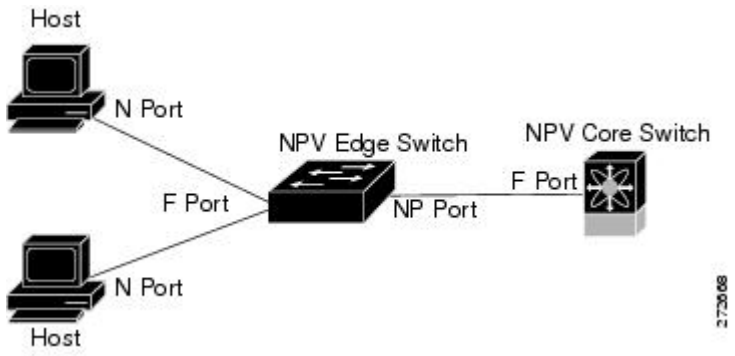

NPV Mode

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.

Note As the NPIV box has multiple FLOGIs from the NPV box, the **disable-feature** command is rejected.

In Cisco Nexus devices, server interfaces can be virtual Fibre Channel interfaces.

Related Topics

Configuring N Port Virtualization, on page 27

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 forwarded as-is to the core switch.

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

In Cisco Nexus devices, NP uplink interfaces are virtual Fibre Channel interfaces.

Related Topics

Fabric Login, on page 119

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 3: 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.

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.

Related Topics

Information About Zones, on page 79

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.

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.
- Servers can be connected to the switch when in NPV mode.
- When initiators and targets are assigned to the same border port (NP or NP-PO), then Cisco Nexus 5000 Series switches in NPIV mode do not support hairpinning.
- Fibre Channel switching is not performed in the edge switch; all traffic is switched in the core switch.
- NPV supports NPIV-capable servers. This capability is called nested NPIV.
- Connecting two Cisco NPV switches together is not supported.
- Only VF and VNP port types are supported in NPV mode.
- For an NPV switch which is configured for trunking on any interface, or for a regular switch where the f port-channel-trunk command is issued to enable the Trunking F Port Channels feature, follow these configuration guidelines for reserved VSANs and isolated VSAN:

- If the trunk mode is enabled for any of the interfaces, or if the NP port channel is up, the reserved VSANs range from 3840 to 4078, which are not available for user configuration.
- The Exchange Virtual Fabric Protocol (EVFP) isolated VSAN is 4079, and it is not available for user configuration.

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:

Procedure

	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** When the switch is reloaded in the NPV mode, only the following configurations are saved: • **switchname** • **management ip** configuration and **vrf** • **boot** variable • **username** / **password** details • **ntp** configuration • **callhome** configuration • **snmp-server** details • **feature fcoe**
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:

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface vfc** *vfc-id*	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:

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface vfc** *vfc-id*	Selects an interface that will be connected to the core NPV switch.
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:

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **npv traffic-map server-interface vfc** *vfc-id* **external-interface vfc** *vfc-id*	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 vfc** *vfc-id* **external-interface vfc** *vfc-id*	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:

Procedure

	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:

Procedure

	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 device:

```
switch# show npv flogi-table
--------------------------------------------------------------------------------
SERVER                                                                          EXTERNAL
INTERFACE VSAN FCID             PORT NAME               NODE NAME               INTERFACE
--------------------------------------------------------------------------------
vfc31     1    0xee0008 10:00:00:00:c9:60:e4:9a 20:00:00:00:c9:60:e4:9a  vfc21
vfc31     1    0xee0009 20:00:00:00:0a:00:00:01 20:00:00:00:c9:60:e4:9a  vfc22
vfc31     1    0xee000a 20:00:00:00:0a:00:00:02 20:00:00:00:c9:60:e4:9a  vfc23
vfc31     1    0xee000b 33:33:33:33:33:33:33:33 20:00:00:00:c9:60:e4:9a  vfc24

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: vfc21, VSAN: 1, FCID: 0x1c0000, State: Up
  Interface: vfc22, VSAN: 1, FCID: 0x040000, State: Up
  Interface: vfc23, VSAN: 1, FCID: 0x260000, State: Up
  Interface: vfc24, VSAN: 1, FCID: 0x1a0000, State: Up

  Number of External Interfaces: 4

Server Interfaces:
==================
  Interface: vfc31, 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)
----------------------------------------
vfc13           vfc110,vfc111
vfc15           vfc11,vfc12
----------------------------------------
```

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: vfc21, VSAN: 2, FCID: 0x1c0000, State: Up
...
```

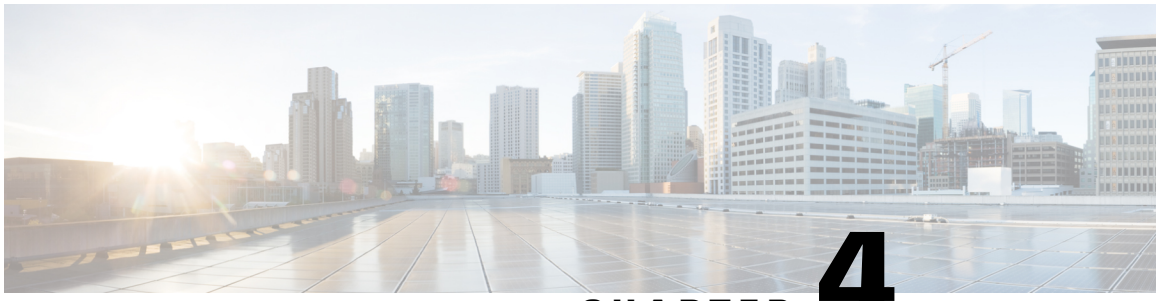

CHAPTER 4

Configuring FCoE NPV

This chapter contains the following sections:

Information About FCoE NPV

FCoE NPV is supported on the Cisco Nexus devices. The FCoE NPV feature is an enhanced form of FIP snooping that provides a secure method to connect FCoE-capable hosts to an FCoE-capable FCoE forwarder (FCF) switch. The FCoE NPV feature provides the following benefits:

- FCoE NPV does not have the management and troubleshooting issues that are inherent to managing hosts remotely at the FCF.

- FCoE NPV implements FIP snooping as an extension to the NPV function while retaining the traffic-engineering, vsan-management, administration and trouble-shooting aspects of NPV.
- FCoE NPV and NPV together allow communication through FC and FCoE ports at the same time. This provides a smooth transition when moving from FC to FCoE topologies.

You can enable FCoE NPV by choosing one of the following methods:

- **Enable FCoE and then enable NPV—**This method requires that you enable FCoE first using the **feature fcoe** command and then you enable NPV by using the **feature npv** command. When FCoE is enabled, the default mode of operation is FC switching and when you enable NPV, the mode changes to NPV mode. Switching to NPV mode automatically performs a write erase and reloads the system. After the reload, the system comes up in NPV mode. To exit NPV mode and return to FC switching mode, enter the **no feature npv** command. Exiting NPV mode also triggers a write erase and a switch reload. This method requires the Storage Protocols Services Package (FC_FEATURES_PKG) license

  .
- **Enable FCoE NPV—**When you enable FCoE NPV using the **feature fcoe-npv** command, the mode changes to NPV. When you use this method, a write erase and reload does not occur. This method requires a separate license package (N6K-FNPV-SSK9). This license is also included in the Storage Protocol Services License.

Method	License	Write Erase	Reload
Enable FCoE and then Enable NPV	Storage Protocols Services Package (FC_FEATURES_PKG)	Yes	Yes
Enable FCoE NPV	(N6K-FNPV-SSK9)	No	No

Interoperability with FCoE-Capable Switches

The Cisco Nexus device interoperates with the following FCoE-capable switches:

- Cisco MDS 9000 Series Multilayer switches enabled to perform FCF functions (EthNPV and VE)
- Cisco Nexus 7000 Series switches enabled to perform FCF functions (EthNPV and VE)
- Cisco Nexus 4000 Series switches enabled for FIP Snooping

For detailed information about switch interoperability, see the Cisco Data Center Interoperability Support Matrix.

Licensing

The following table shows the licensing requirements for FCoE NPV:

Product	License Requirement
NX-OS	FCoE NPV requires a separate license (FCOE_NPV_PKG). The FCoE NPV license is also included in the Storage Protocol Services License. FCoE and NPV require the Storage Protocols Services Package (FC_FEATURES_PKG). For detailed information about features that require licensing and Cisco NX-OS license installation, see the Cisco NX-OS Licensing Guide. For information about troubleshooting licensing issues, see the Troubleshooting Guide for your device.

FCoE NPV Model

The following figure shows the FCoE NPV bridge connecting hosts and FCFs. From a control plane perspective, FCoE NPV performs proxy functions towards the FCF and the hosts in order to load balance logins from the

hosts evenly across the available FCF uplink ports. An FCoE NPV bridge is VSAN-aware and capable of assigning VSANs to the hosts.

Figure 4: FCoE NPV Model

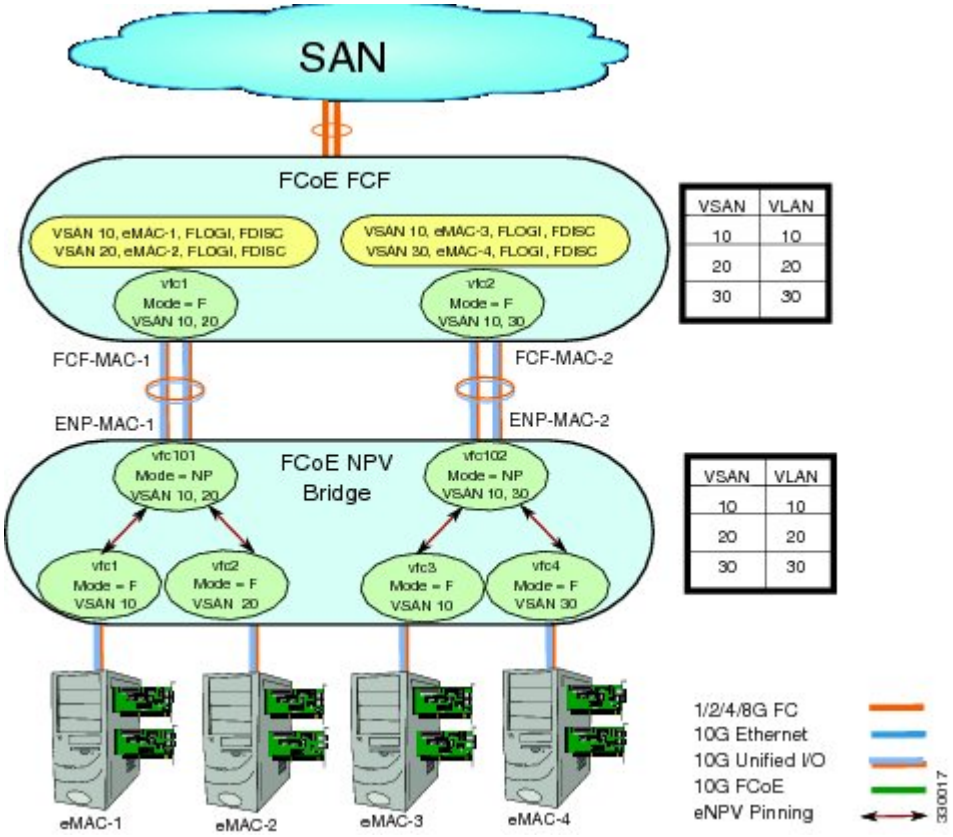

Mapping Requirements

VSANs and VLAN-VSAN Mapping

VSANs from the hosts must be created and for each VSAN, a dedicated VLAN must also be created and mapped. The mapped VLAN is used to carry FIP and FCoE traffic for the corresponding VSAN. The VLAN-VSAN mapping must be configured consistently in the entire fabric. The Cisco Nexus device supports 32 VSANs.

FC Mapping

The FC-MAP value associated with a SAN fabric must be configured on the FCoE NPV bridge which helps the FCoE NPV bridge isolate misconnections to FCFs in other fabrics.

Port Requirements

VF Ports

For each host directly connected over Ethernet interfaces on the FCoE NPV bridge, a virtual Fibre Channel (vFC) interface must be created and bound to the Ethernet interface. By default, the vFC interface is configured in the F mode (VF port).

The VF port must be configured with the following parameters:

- A VF port must be bound to a VLAN trunk Ethernet interface or a port-channel interface. The FCoE VLAN must not be configured as the native VLAN on the Ethernet interface.
- A port VSAN must be configured for the VF port.
- The administrative state must be up.

VNP Ports

Connectivity from an FCoE NPV bridge to the FCF is only supported over point-to-point links. These links can be individual Ethernet interfaces or members of an Ethernet port channel interface. For each FCF connected Ethernet interfaces, a vFC interface must be created and bound to the Ethernet interface. These vFC interfaces must be configured as VNP ports. On the VNP port, an FCoE NPV bridge emulates an FCoE-capable host with multiple enodes, each with a unique enode MAC address. A VNP port interface binding to MAC address is not supported. By default, the VNP port is enabled in trunk mode. Multiple VSANs can be configured on the VNP port. The FCoE VLANs that correspond to the VNP port VSANs must be configured on the bound Ethernet interface.

Note The spanning-tree protocol (STP) is automatically disabled in the FCoE VLAN on the interfaces that the VNP port are bound to.

NPV Features

The following NPV features apply for the FCoE NPV feature:

- Automatic Traffic Mapping
- Static Traffic Mapping
- Disruptive Load Balancing
- FCoE Forwarding in the FCoE NPV Bridge
- FCoE frames received over VNP ports are forwarded only if the L2_DA matches one of the FCoE MAC addresses assigned to hosts on the VF ports otherwise they're discarded.

vPC Topologies

When VNP ports are configured vPC topologies between an FCoE NPV bridge and an FCF, the following limitations apply:

- vPC spanning multiple FCFs in the same SAN fabric is not supported.
- For LAN traffic, dedicated links must be used for FCoE VLANs between the FCoE NPV bridge and the FCF connected over a vPC.
- FCoE VLANs must not be configured on the inter-switch vPC interfaces.
- VF port binding to a vPC member port is not supported for an inter-switch vPC.

Figure 5: VNP Ports in an Inter-Switch vPC Topology

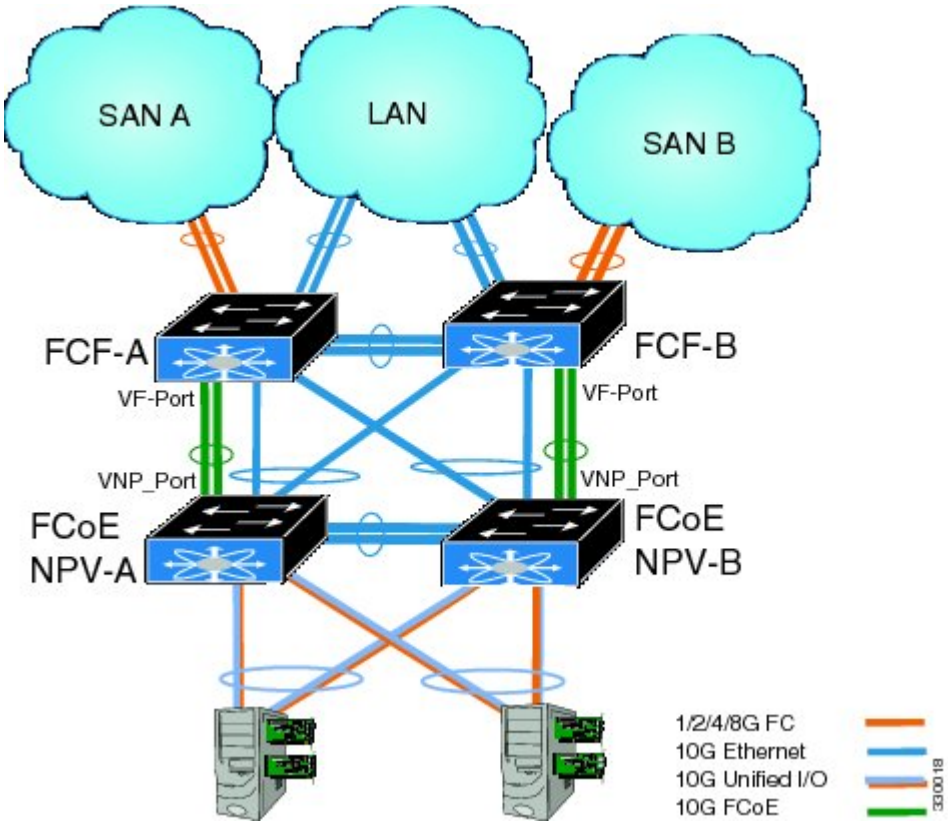

Supported and Unsupported Topologies

FCoE NPV supports the following topologies:

Figure 6: Cisco Nexus Device As An FCoE NPV Device Connected to a Cisco Nexus Device Over A Non- vPC Port Channel

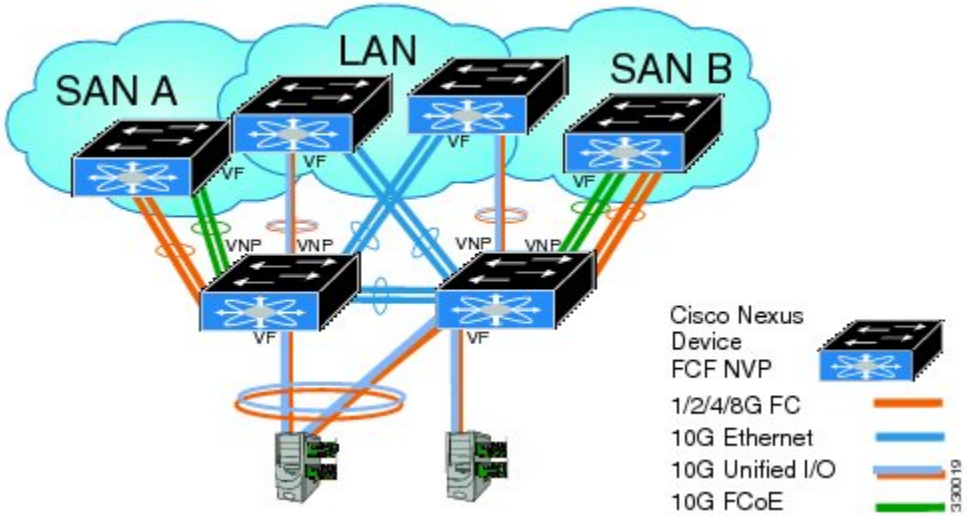

Figure 7: Cisco Nexus Device As An FCoE NPV Device Connected Over a vPC To Another Cisco Nexus Device

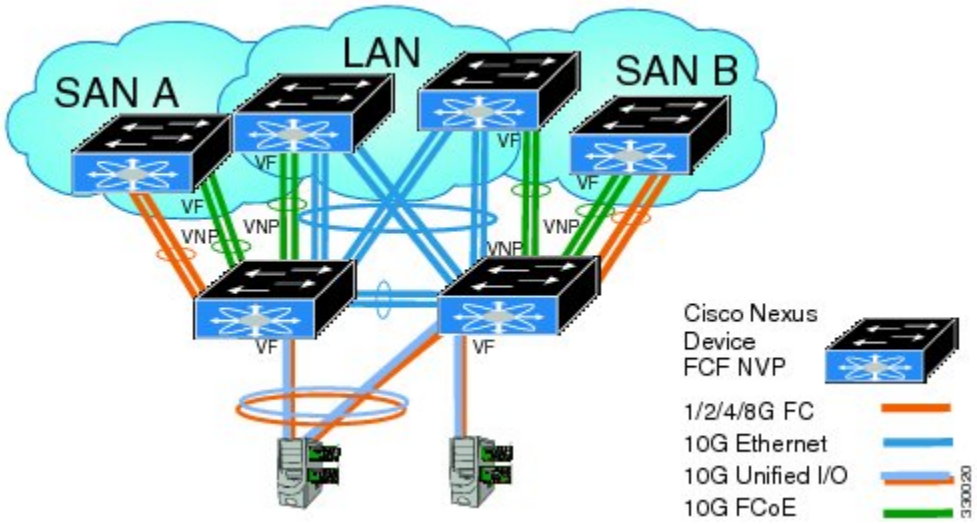

Figure 8: Cisco Nexus Device With A 10GB Fabric Extender As An FCoE NPV Device Connected to a Cisco Nexus Device Over A Non- vPC Port Channel

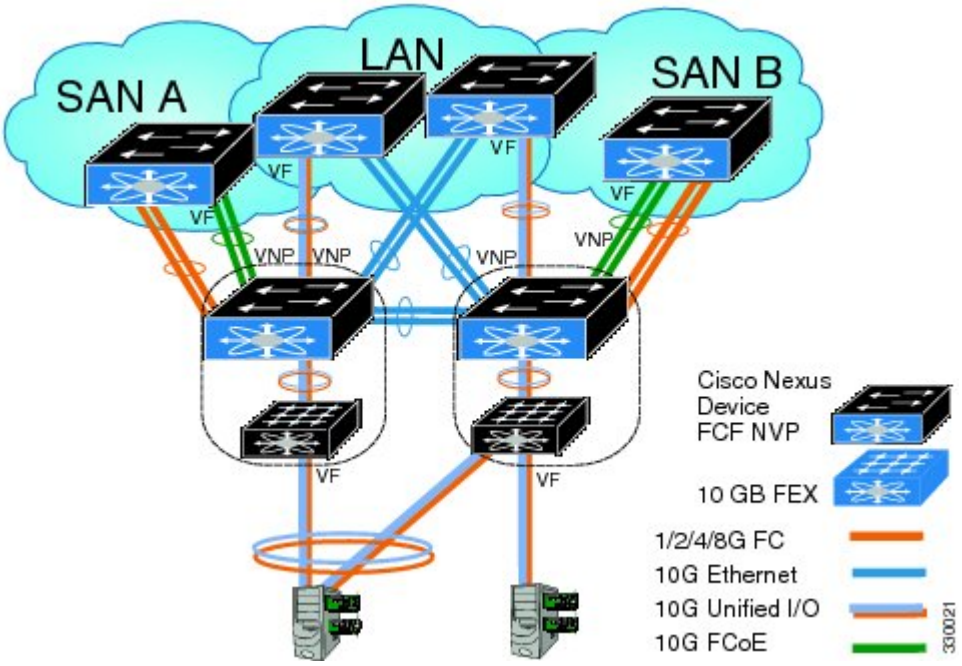

Figure 9: Cisco Nexus Device With A 10GB Fabric Extender as an FCoE NPV Device Connected Over a vPC to Another Cisco Nexus Device

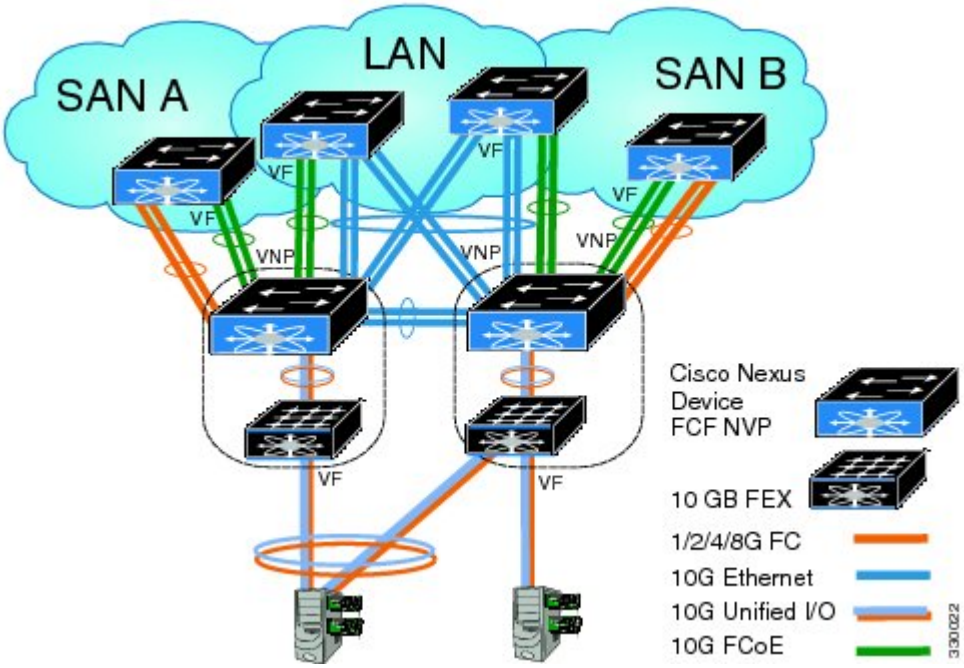

Figure 10: Cisco Nexus Device As An FCoE NPV Bridge Connecting to a FIP Snooping Bridge

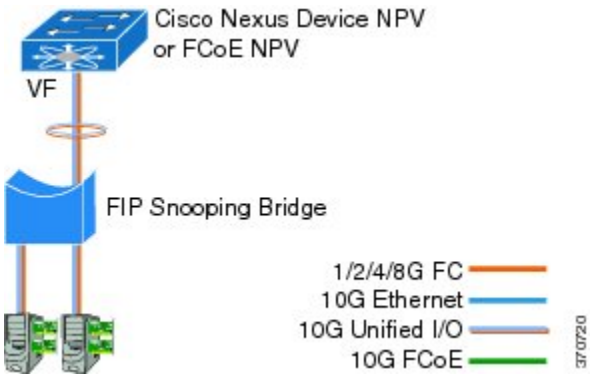
Cisco Nexus Device NPV
or FCoE NPV
VF
FIP Snooping Bridge
1/2/4/8G FC
10G Ethernet
10G Unified I/O
10G FCoE
370720

Unsupported Topologies

FCoE NPV does not support the following topologies:

Figure 11: 10GB Fabric Extender Connecting To The Same FCoE NPV Bridge Over Multiple VF Ports

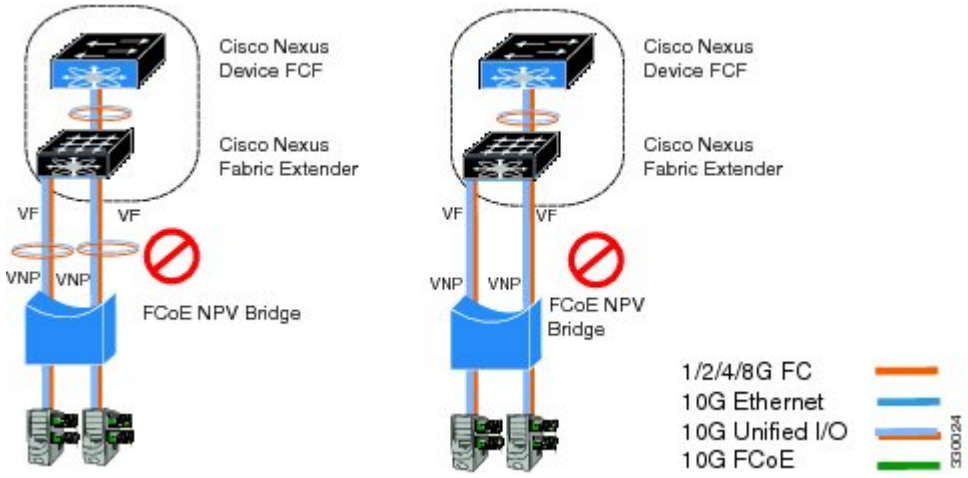

Figure 12: Cisco Nexus Device As An FCoE NPV Bridge Connecting To A FIP Snooping Bridge Or Another FCoE NPV Bridge

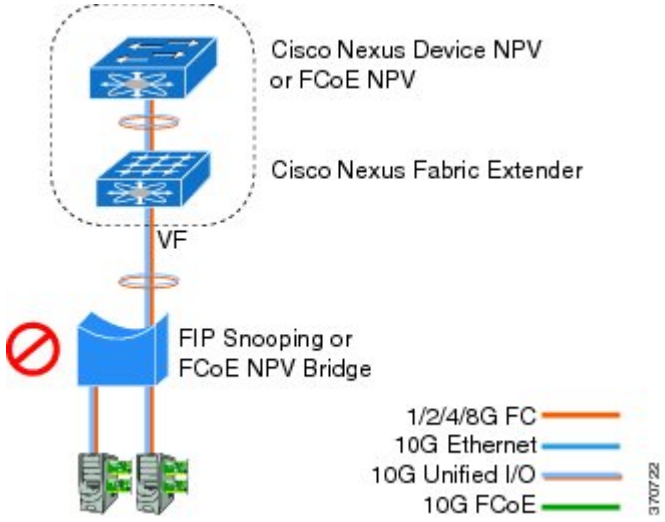

Figure 13: VF Port Trunk To Hosts In FCoE NPV Mode

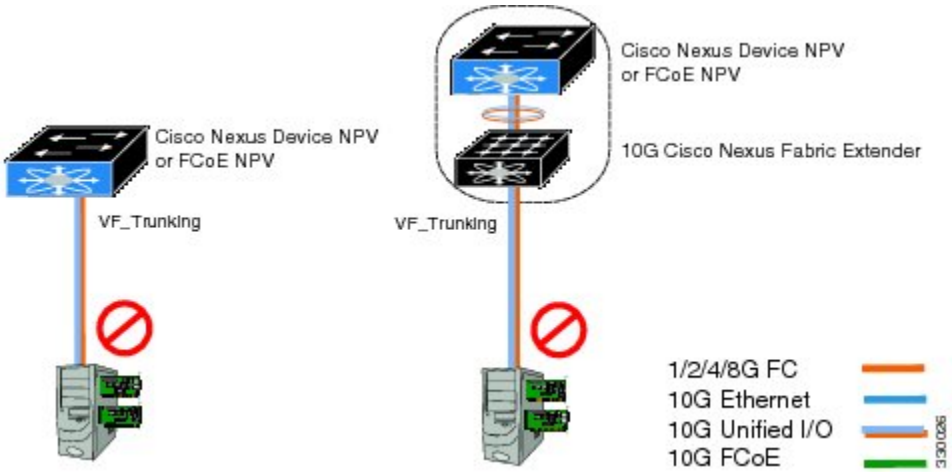

Figure 14: Cisco Nexus Device As An FCoE NPV Bridge Connecting to an FCoE NPV Bridge

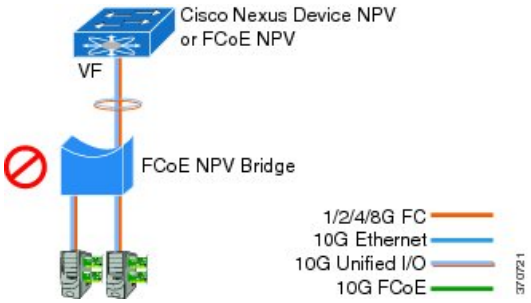

Guidelines and Limitations

The FCoE NPV feature has the following guidelines and limitations:

- When FCoE NPV mode is configured on a switch, the FCoE feature cannot be enabled. A warning is displayed to reload the system first in order to enable FCoE.
- You can not perform an in-service software downgrade (ISSD) to Cisco NX-OS Release 5.0(3)N1(1) or an earlier release if FCoE NPV is enabled and if VNP ports are configured.
- A warning is displayed if an ISSD is performed to Cisco NX-OS Release 5.0(3)N1(1) or an earlier release when FCoE NPV is enabled but VNP ports are not configured.
- Before performing an ISSU on an FCoE NPV bridge, use the **disable-fka** command to disable the timeout value check (FKA check) on the core switch.

FCoE NPV Configuration Limits

The following table lists the FCoE configuration limits over Ethernet, Ethernet port channel, and virtual Ethernet interfaces.

Table 4: VNP Port Configuration Limits

Interface Type	Cisco Nexus 6000 Series	Cisco Nexus 2000 Series (10G interfaces)
VNP port bound to Ethernet interface	4 VNP ports	Not Supported
VNP port bound to Ethernet port channel interface	2 VNP ports	Not Supported
VNP port bound to virtual Ethernet (vEth) interface	Not Supported	Not Supported

The configuration limits guidelines are as follows:

- The number of VF port and VN port interfaces that can be supported between a given FCF and an FCoE NPV bridge also depends on the FCF to MAC advertising capability of the FCF:
  - If an FCF advertises the same FCF-MAC address over all of its interfaces, then the FCoE NPV bridge can connect to it over one VNP Port. In this scenario, we recommend that one port channel interface be used for redundancy.
  - If an FCF advertises multiple FCF-MAC addresses, then the limits in the previous table apply. For additional information, see the best practices recommendations for the FCF switch.
- The total number of supported VSANs is 31 (excluding the EVFP VSAN).
- The total number of supported FCIDs is 2048.

Default Settings

The following table lists the default settings for FCoE NPV parameters.

Table 5: Default FCoE NPV Parameters

Parameters	Default
FCoE NPV	Disabled
FCoE	Disabled
NPV	Disabled
VNP port	Disabled
FIP Keep Alive (FKA)	Disabled

Enabling FCoE and Enabling NPV

You can enable FCoE first and then enable NPV. This method requires the full Storage Services License. A write erase reload occurs when this method is used. This method allows both FCoE and FC upstream and host NPV connections. You must also configure class-fcoe in all QoS policy types.

1 Enable FCoE.

```
switch# configure terminal
switch(config)# feature fcoe
FC license checked out successfully
fc_plugin extracted successfully
FC plugin loaded successfully
FCoE manager enabled successfully
FC enabled on all modules successfully
Warning: Ensure class-fcoe is included in qos policy-maps of all types
```

2 Enable NPV.

```
switch# configure terminal
switch(config)# feature npv
```

Enabling FCoE NPV

You can enable FCoE NPV using the **feature fcoe-npv** command. We recommend this method in topologies that include all FCoE connections. A write erase reload does not occur when you use this method and a storage service license is not required. Enabling FCoE NPV using the **feature fcoe-npv** command requires an installed FCOE_NPV_PKG license.

Before You Begin

FCoE NPV has the following prerequisites:

- Ensure that the correct licenses are installed.
- Configure the VNP ports.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	**feature fcoe-npv**	Enables FCoE NPV.
Step 3	**exit**	Exits configuration mode.
Step 4	switch(config)# **copy running-config startup-config**	(Optional) Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

This example shows how to enable FCoE NPV using the **feature fcoe-npv** command.

```
switch# configure terminal
switch(config)# feature fcoe-npv
FCoE NPV license checked out successfully
fc_plugin extracted successfully
FC plugin loaded successfully
FCoE manager enabled successfully
FCoE NPV enabled on all modules successfully
```

This example shows how to enable FCoE NPV using the **feature fcoe** and **feature npv** commands.

```
switch# configure terminal
switch(config)# feature fcoe
switch(config)# feature npv
```

Configuring NPV Ports for FCoE NPV

You can configure NVP port for FCoE NPV.

1. Create a vFC port.

   ```
   switch# config t
   switch(config)# interface vfc 20
   switch(config-if)#
   ```

2. Bind the vFC to an Ethernet port.

   ```
   switch(config-if)# bind interface ethernet 1/20
   switch(config-if)#
   ```

3. Set the port mode to NP.

   ```
   switch(config-if)# switchport mode NP
   switch(config-if)#
   ```

4. Bring up the port:

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

Verifying FCoE NPV Configuration

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

Command	Purpose
show fcoe database	Displays information about the FCoE database.
show interface Ethernet x/y fcoe	Displays FCoE information for a specified Ethernet interface including the following: • FCF or associated enode MAC address • Status • Associated VFC information
show interface vfc x	Displays information about the specified vFC interface including attributes and status.

Command	Purpose
show npv status	Displays the status of the NPV configuration including information about VNP ports.
show fcoe-npv issu-impact	Displays the impact of FCoE NPV on an ISSU.
show running-config fcoe_mgr	Displays the running configuration information about FCoE.
show startup-config fcoe_mgr	Displays the startup configuration information about FCoE.
show tech-support fcoe	Displays troubleshooting information about FCoE.
show npv flogi-table	Displays information about N port virtualization (NPV) fabric login (FLOGI) session
show fcoe	Displays the status of Fibre Channel over Ethernet (FCoE) configurations.

For detailed information about the fields in the output from these commands, refer to the command reference for your device.

Configuration Examples for FCoE NPV

This example shows how to enable FCoE NPV, LACP, QoS for no drop queuing, and VLAN/VSAN mapping:

```
switch# config t
switch(config)# feature fcoe-npv
FCoE NPV license checked out successfully
fc_plugin extracted successfully
FC plugin loaded successfully
FCoE manager enabled successfully
FCoE NPV enabled on all modules successfully

switch(config)# feature lacp

switch# config t
switch(config)# system qos
switch(config-sys-qos)# service-policy type qos input fcoe-default-in-policy
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 network-qos fcoe-default-nq-policy

switch(config)# vsan database
switch(config-vsan-db)# vsan 50-51
switch(config-vsan-db)# vlan 50
switch(config-vlan)# fcoe vsan 50
switch(config-vlan)# vlan 51
switch(config-vlan)# fcoe vsan 51

This example shows a summary of the interface configuration information for trunked NP
ports:
switch# show interface brief | grep TNP

vfc25      400    NP     on      trunking         swl    TNP     2    --
```

```
vfc26      400    NP     on      trunking         swl    TNP     2    --
vfc130     1      NP     on      trunking         --     TNP     auto --
switch#
```

This example shows the running configuration information about FCoE:

```
switch# show running-config fcoe_mgr

!Command: show running-config fcoe_mgr
!Time: Wed Jan  20 21:59:39 2013

version 6.0(2)N1(1)

interface vfc1
  bind interface Ethernet1/19

interface vfc2
  bind interface Ethernet1/2

interface vfc90
  bind interface Ethernet1/9

interface vfc100
  bind interface Ethernet1/10

interface vfc110
  bind interface port-channel110

interface vfc111
  bind interface Ethernet1/11

interface vfc120
  bind interface port-channel120

interface vfc130
  bind interface port-channel130

interface vfc177
  bind interface Ethernet1/7
fcoe fka-adv-period 16
```

This example shows the FCoE VLAN to VSAN mappings:

```
switch# show vlan fcoe

Original VLAN ID        Translated VSAN ID      Association State
----------------        ------------------      -----------------

      400                     400              Operational
      20                      20               Operational
      100                     100              Operational
      500                     500              Operational
      200                     200              Operational
      300                     300              Operational
```

This example shows the information about the vFC 130 interface including attributes and status:

```
switch# show interface vfc 130
vfc130 is trunking (Not all VSANs UP on the trunk)
    Bound interface is port-channel130
    Hardware is Virtual Fibre Channel
    Port WWN is 20:81:00:05:9b:74:bd:bf
    Admin port mode is NP, trunk mode is on
    snmp link state traps are enabled
    Port mode is TNP
    Port vsan is 1
    Trunk vsans (admin allowed and active) (1,20,100,200,300,400,500)
    Trunk vsans (up)                       (500)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             (1,20,100,200,300,400)
    1 minute input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    1 minute output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      15 frames input, 2276 bytes
```

```
      0 discards, 0 errors
    7 frames output, 1004 bytes
      0 discards, 0 errors
  last clearing of "show interface" counters Tue May 31 20:56:41 2011

  Interface last changed at Wed Jun  1 21:53:08 2011
```

This example shows the information about the vFC 1 interface including attributes and status:

```
switch# show interface vfc 1
vfc1 is trunking (Not all VSANs UP on the trunk)
    Bound interface is Ethernet1/19
    Hardware is Virtual Fibre Channel
    Port WWN is 20:00:00:05:9b:74:bd:bf
    Admin port mode is F, trunk mode is on
    snmp link state traps are enabled
    Port mode is TF
    Port vsan is 20
    Trunk vsans (admin allowed and active) (1,20,100,200,300,400,500)
    Trunk vsans (up)                       (20)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             (1,100,200,300,400,500)
    1 minute input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    1 minute output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      355278397 frames input, 573433988904 bytes
        0 discards, 0 errors
      391579316 frames output, 572319570200 bytes
        0 discards, 0 errors
    last clearing of "show interface" counters Tue May 31 20:56:41 2011

    Interface last changed at Wed Jun  1 20:25:36 2011
```

This example shows the information about the NPV FLOGI session:

```
switch# show npv flogi-table
--------------------------------------------------------------------------------
SERVER                                                                  EXTERNAL
INTERFACE VSAN FCID            PORT NAME               NODE NAME       INTERFACE
--------------------------------------------------------------------------------
vfc1      20   0x670000 21:01:00:1b:32:2a:e5:b8 20:01:00:1b:32:2a:e5:b8 vfc26

Total number of flogi = 1.
```

This example shows the status of the NPV configuration including information about VNP ports:

```
switch# show npv status

npiv is enabled

disruptive load balancing is disabled

External Interfaces:
====================
  Interface:  vfc25, State: Trunking
        VSAN:    1, State: Up
        VSAN:  200, State: Up
        VSAN:  400, State: Up
        VSAN:   20, State: Up
        VSAN:  100, State: Up
        VSAN:  300, State: Up
        VSAN:  500, State: Up, FCID: 0xa10000
  Interface:  vfc26, State: Trunking
        VSAN:    1, State: Up
        VSAN:  200, State: Up
        VSAN:  400, State: Up
        VSAN:   20, State: Up
        VSAN:  100, State: Up
        VSAN:  300, State: Up
        VSAN:  500, State: Up, FCID: 0xa10001
  Interface: vfc90,  State: Down
  Interface: vfc100, State: Down
  Interface: vfc110, State: Down
  Interface: vfc111, State: Down
```

```
  Interface: vfc120, State: Down
  Interface: vfc130, State: Trunking
        VSAN:    1, State: Waiting For VSAN Up
        VSAN:  200, State: Up
        VSAN:  400, State: Up
        VSAN:  100, State: Up
        VSAN:  300, State: Up
        VSAN:  500, State: Up, FCID: 0xa10002

  Number of External Interfaces: 8

Server Interfaces:
==================
  Interface:   vfc1, VSAN:   20, State: Up
  Interface:   vfc2, VSAN: 4094, State: Down
  Interface:   vfc3, VSAN: 4094, State: Down
  Interface: vfc5000, VSAN: 4094, State: Down
  Interface: vfc6000, VSAN: 4094, State: Down
  Interface: vfc7000, VSAN: 4094, State: Down
  Interface: vfc8090, VSAN: 4094, State: Down
  Interface: vfc8191, VSAN: 4094, State: Down

  Number of Server Interfaces: 8
```

This example shows the running configuration of port channel 130:

```
switch# show running-config interface port-channel 130

!Command: show running-config interface port-channel130
!Time: Wed Jan  30 22:01:05 2013

version 6.0(2)N1(1)

interface port-channel130
  switchport mode trunk
  switchport trunk native vlan 2
  no negotiate auto
```

This example shows the impact of FCoE NPV on an ISSU:

```
switch# show fcoe-npv issu-impact
show fcoe-npv issu-impact
-------------------------

Please make sure to enable "disable-fka" on all logged in VFCs
Please increase the FKA duration to 60 seconds on FCF


Active VNP ports with no disable-fka set
----------------------------------------

vfc90
vfc100
vfc110
vfc111
vfc120
vfc130

ISSU downgrade not supported as feature fcoe-npv is enabled
switch#
```

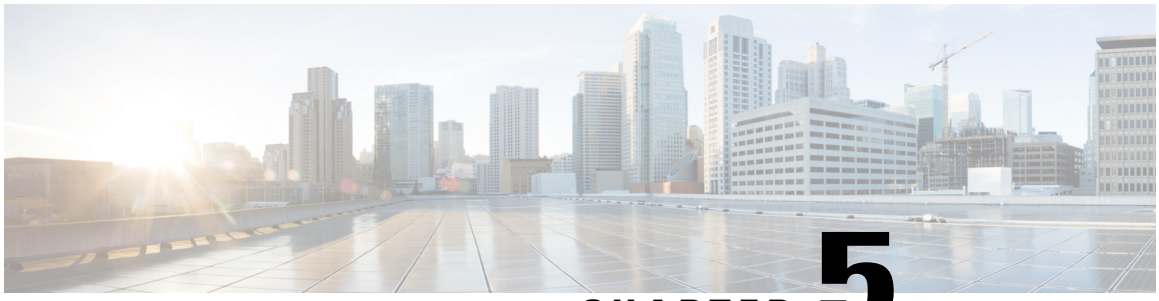

CHAPTER 5

Configuring VSAN Trunking

This chapter describes how to configure VSAN trunking.

This chapter includes the following sections:

Configuring VSAN Trunking

Information About VSAN Trunking

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

VSAN trunking is supported on virtual Fibre Channel interfaces.

The VSAN trunking feature includes the following restrictions:

- Trunking configurations are applicable only 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.

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 15: VSAN Mismatch

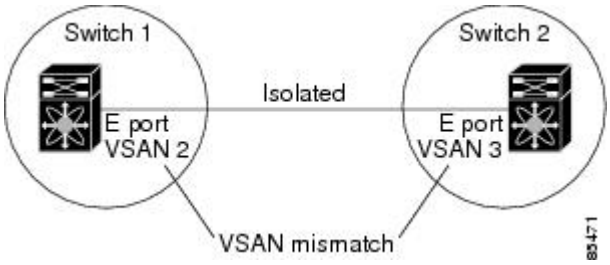

In this example, the trunking protocol detects potential VSAN merging and isolates the ports involved.

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

Figure 16: Third-Party Switch VSAN Mismatch

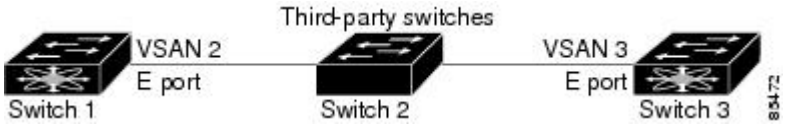

VSAN 2 and VSAN 3 are effectively merged with overlapping entries in the name server and the zone applications. Cisco MDS 9000 Fabric Manager helps detect such topologies.

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 Limitations

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.

Enabling or Disabling the VSAN Trunking Protocol

You can enable or disable the VSAN trunking protocol.

Procedure

	Command or Action	Purpose

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

Step 2	**no trunk protocol enable** **Example:** `switch(config)# no trunk protocol enable`	Disables the trunking protocol.
Step 3	**trunk protocol enable** **Example:** `switch(config)# trunk protocol enable`	Enables trunking protocol (default).

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 6: 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 SAN 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 Inter-Switch Link (ISL) is always in a trunking disabled state.

Configuring Trunk Mode

You can configure trunk mode.

Procedure

	Command or Action	Purpose

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

Step 2	switch(config)# **interface vfc** *vfc-id*	Selects an interface that will be connected to the core NPV switch.
Step 3	**interface vfc** *vfc-id* **Example:** `switch(config)# interface vfc 15`	Configures the specified Fibre Channel or virtual Fibre Channel interface.
Step 4	**switchport trunk mode on** **Example:** `switch(config-if)# switchport trunk mode on`	Enables (default) the trunk mode for the specified interface.

	Command or Action	Purpose
Step 5	**switchport trunk mode off** **Example:** `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 6	**switchport trunk mode auto** **Example:** `switch(config-if)# switchport trunk mode auto`	Configures the trunk mode to **auto** mode, which provides automatic sensing for the interface.

EXAMPLES

This example shows how to configure a vFC interface in trunk mode:

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

This 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
```

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 17: Default Allowed-Active VSAN Configuration

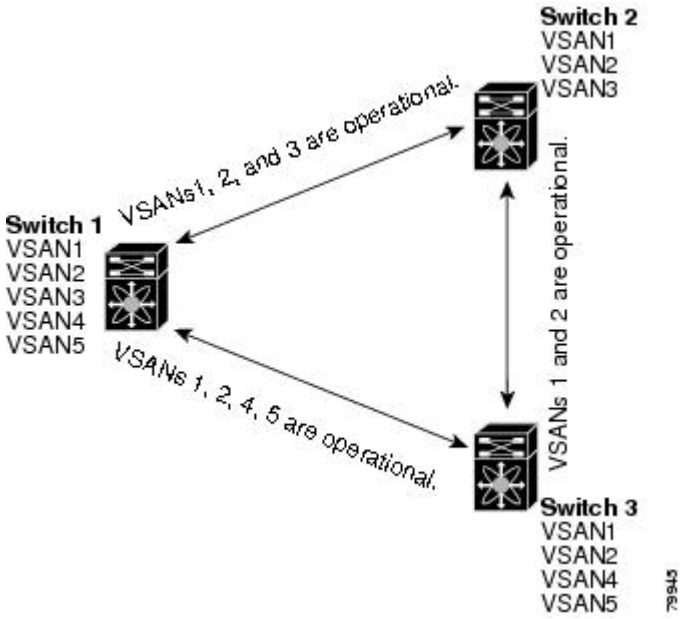

You can configure a selected set of VSANs (from the allowed-active list) to control access to the VSANs specified in a trunking ISL.

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.

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

Figure 18: Operational and Allowed VSAN Configuration

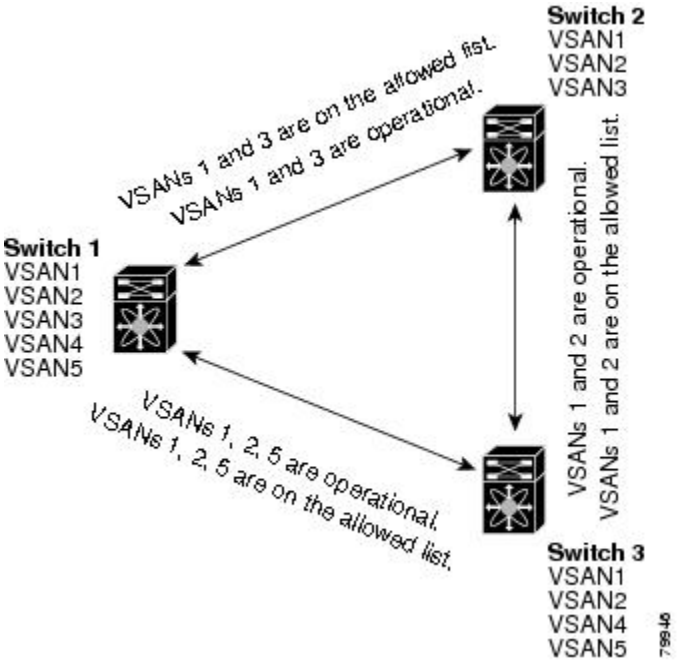

Configuring an Allowed-Active List of VSANs

You can configure an allowed-active list of VSANs for an interface.

Procedure

	Command or Action	Purpose

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

Step 2	**interface vfc** *vfc-id* **Example:** `switch(config)# interface vfc 4`	Configures the specified interface.
Step 3	**switchport trunk allowed vsan** *vsan-id* **-** *vsan-id* **Example:** `switch(config-if)# switchport trunk allowed vsan 35-55`	Changes the allowed list for the specified VSAN range.

	Command or Action	Purpose
Step 4	**switchport trunk allowed vsan add** *vsan-id* **Example:** `switch(config-if)# switchport trunk allowed vsan add 40`	Expands the specified VSAN to the new allowed list.
Step 5	**no switchport trunk allowed vsan** *vsan-id* **-** *vsan-id* **Example:** `switch(config-if)# no switchport trunk allowed vsan 61-65`	Deletes the specified VSAN range.
Step 6	**no switchport trunk allowed vsan add** *vsan-id* **Example:** `switch(config-if)# no switchport trunk allowed vsan add 40`	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 vfc33
vfc33 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
vfc31 is not trunking
...
vfc311 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 Settings for VSAN Trunks

The following table lists the default settings for VSAN trunking parameters.

Table 7: Default VSAN Trunk Configuration Parameters

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

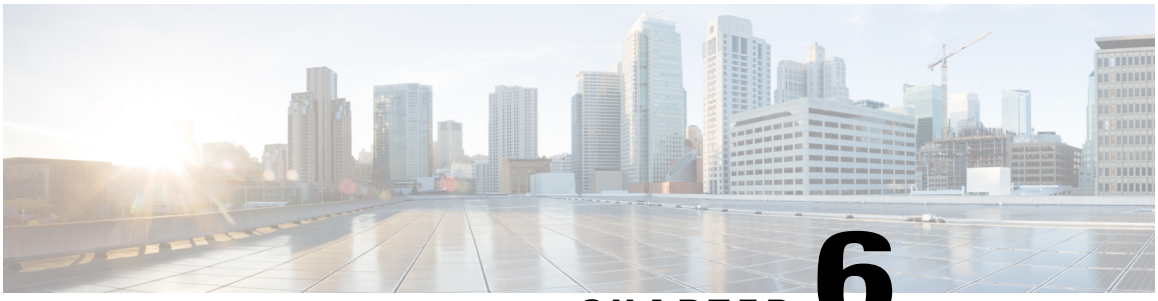

CHAPTER 6

Configuring and Managing VSANs

This chapter describes how to configure and manage VSANs.

This chapter includes 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.

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.

VSAN Topologies

A VSAN has the following additional features:

- Multiple VSANs can share the same physical topology.
- The same Fibre Channel IDs (FC IDs) can be assigned to a host in another VSAN, which increases VSAN scalability.

- Every instance of a VSAN runs all required protocols such as FSPF, domain manager, and zoning.
- Fabric-related configurations in one VSAN do not affect the associated traffic in another VSAN.
- Events causing traffic disruptions in one VSAN are contained within that VSAN and are not propagated to other VSANs.

The following figure shows a fabric with three switches, one on each floor. The geographic location of the switches and the attached devices is independent of their segmentation into logical VSANs. No communication between VSANs is possible. Within each VSAN, all members can talk to one another.

Figure 19: Logical VSAN Segmentation

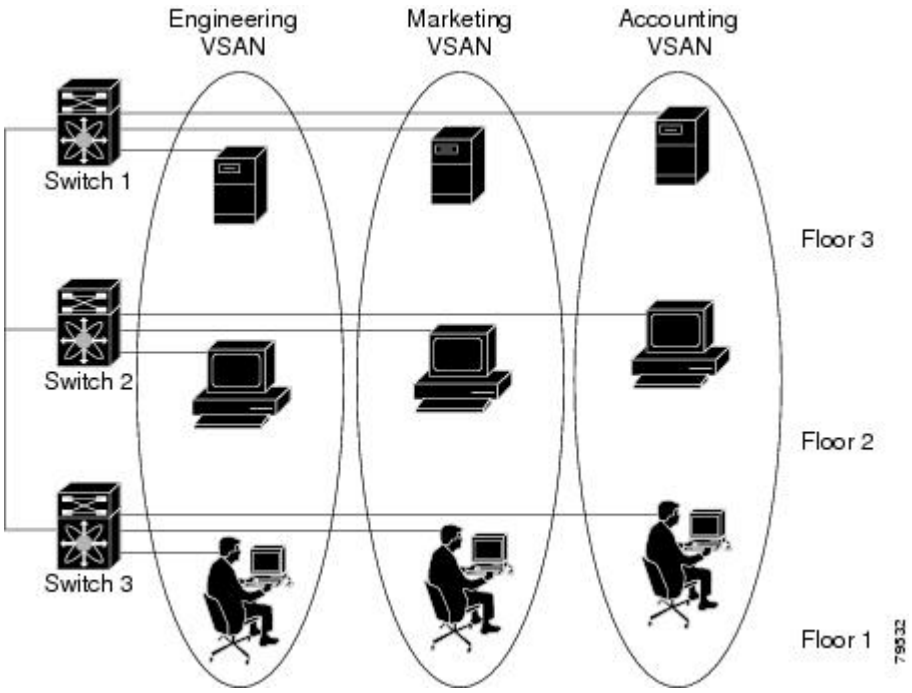

The application servers or storage arrays can be connected to the switch using Fibre Channel or virtual Fibre Channel interfaces. A VSAN can include a mixture of Fibre Channel and virtual Fibre Channel interfaces.

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 20: Example of Two VSANs

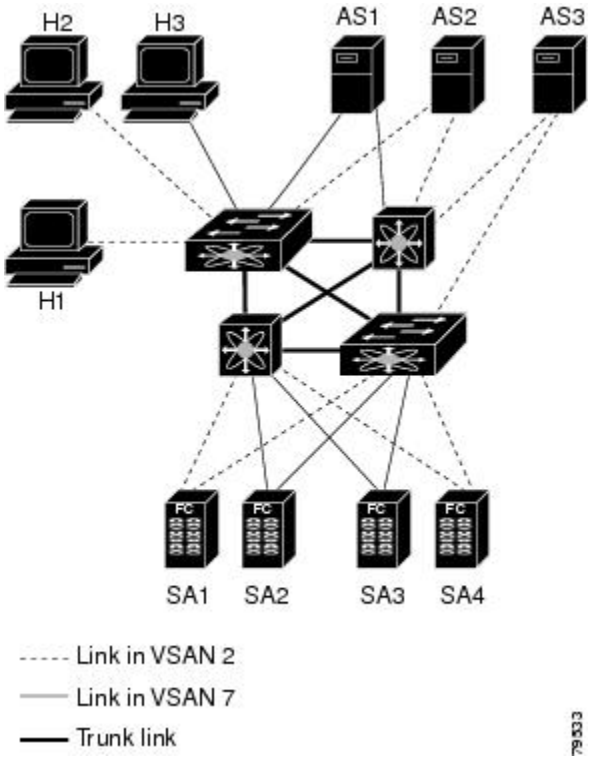

The four switches in this network are interconnected by VSAN trunk links that carry both VSAN 2 and VSAN 7 traffic. You can configure a different inter-switch topology for each VSAN. In the preceding figure, the inter-switch topology is identical for VSAN 2 and VSAN 7.

Without VSANs, a network administrator would need separate switches and links for separate SANs. By enabling VSANs, the same switches and links might be shared by multiple VSANs. VSANs allow SANs to be built on port granularity instead of switch granularity. The preceding figure illustrates that a VSAN is a group of hosts or storage devices that communicate with each other using a virtual topology defined on the physical SAN.

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

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

- VSANs can meet the needs of a particular department or application.

VSAN Advantages

VSANs offer the following advantages:

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

Up to 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 8: 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 21: VSANS with Zoning

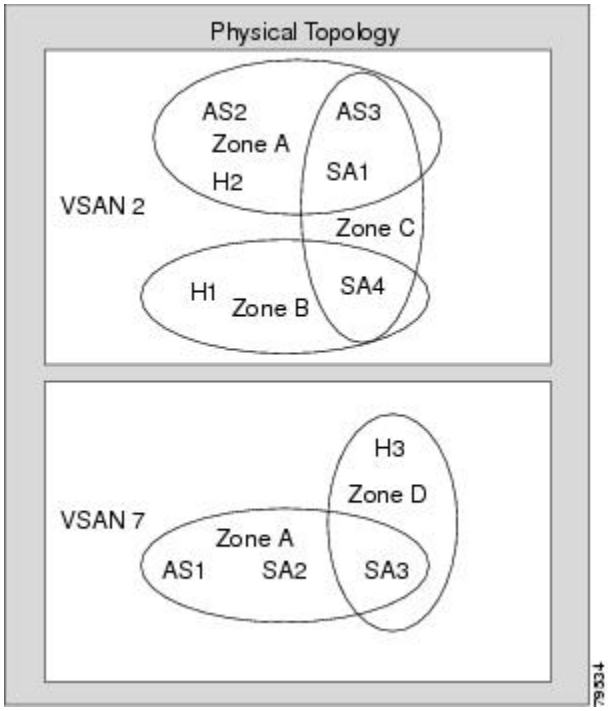

Guidelines and Limitations for VSANs

VSANs have the following configuration guidelines and limitations:

- VSAN ID—The VSAN ID identifies the VSAN as the default VSAN (VSAN 1), user-defined VSANs (VSAN 2 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.
- A VSAN is in the operational state if the VSAN is active and at least one port is up. This state indicates that traffic can pass through this VSAN. This state cannot be configured.
- You can create only 14 VSANs in N5672UP-16G, including the default VSAN 1.
- For an NPV switch which is configured for trunking on any interface, or for a regular switch where the f port-channel-trunk command is issued to enable the Trunking F Port Channels feature, follow these configuration guidelines for reserved VSANs and isolated VSAN:
  - If the trunk mode is enabled for any of the interfaces, or if the NP port channel is up, the reserved VSANs range from 3840 to 4078, which are not available for user configuration.
  - The Exchange Virtual Fabric Protocol (EVFP) isolated VSAN is 4079, and it is not available for user configuration.

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.

Procedure

	Command or Action	Purpose

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

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

Port VSAN Membership

Port VSAN membership on the switch is assigned on a port-by-port basis. By default each port belongs to the default VSAN. You can 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 devices 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 72
Configuring VSAN Trunking, on page 55

Assigning Static Port VSAN Membership

You can statically assign VSAN membership for an interface port.

Procedure

	Command or Action	Purpose

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

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

Step 3	**vsan** *vsan-id* **Example:** `switch(config-vsan-db)# vsan 50`	Creates a VSAN with the specified ID if that VSAN does not exist already.
Step 4	switch(config-vsan-db)# **vsan** *vsan-id* **interface vfc** *vfc-id*	Assigns the membership of the specified interface to the VSAN.
Step 5	switch(config-vsan-db)# **vsan** *vsan-id* **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:
        vfc21   vfc22   vfc23   vfc24

        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:
        vfc21   vfc22   vfc23   vfc24

        san-port-channel 3  vfc31
vsan 2 interfaces:
        vfc23 vfc41
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 vfc21
vfc21
        vsan:1
        allowed list:1-4093
```

Default VSANs

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

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

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

Isolated VSANs

VSAN 4094 is an isolated VSAN. When a VSAN is deleted, all nontrunking ports are transferred to the isolated VSAN to avoid an implicit transfer of ports to the default VSAN or to another configured VSAN. This action ensures that all ports in the deleted VSAN become isolated (disabled).

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

Caution Do not use an isolated VSAN to configure ports.

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

Displaying Isolated VSAN Membership

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

Operational State of a VSAN

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

Static VSAN Deletion

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

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

Figure 22: VSAN Port Membership Details

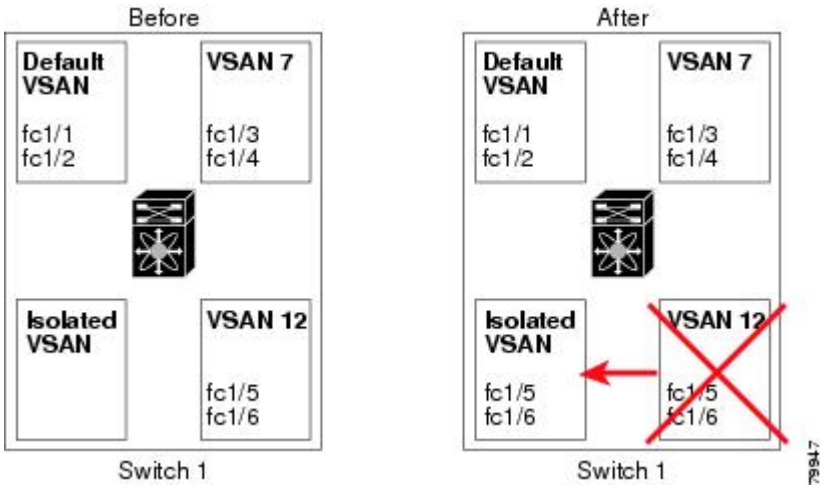

- VSAN-based runtime (name server), zoning, and configuration (static routes) information is removed when the VSAN is deleted.
- Configured VSAN interface information is removed when the VSAN is deleted.

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

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

Related Topics

Configuring VSAN Trunking, on page 55

Deleting Static VSANs

You can delete a VSAN and its various attributes.

Procedure

	Command or Action	Purpose

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

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

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

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

You can configure load balancing on an existing VSAN.

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.

Procedure

	Command or Action	Purpose

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

Step 2	**vsan database**
	
	Example:
	`switch(config)# vsan database`
	`switch(config-vsan-db)#`
		Enters VSAN database configuration submode

Step 3	**vsan** *vsan-id* **Example:** `switch(config-vsan-db)# vsan 15`	Specifies an existing VSAN.
Step 4	**vsan** *vsan-id* **loadbalancing src-dst-id** **Example:** `switch(config-vsan-db)# vsan 15 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	**no vsan** *vsan-id* **loadbalancing src-dst-id** **Example:** `switch(config-vsan-db)# no vsan 15 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	**vsan** *vsan-id* **loadbalancing src-dst-ox-id** **Example:** `switch(config-vsan-db)# vsan 15 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	**vsan** *vsan-id* **suspend** **Example:** `switch(config-vsan-db)# vsan 23 suspend`	Suspends the selected VSAN.
Step 8	**no vsan** *vsan-id* **suspend** **Example:** `switch(config-vsan-db)# no vsan 23 suspend`	Negates the **suspend** command entered in the previous step.
Step 9	**end** **Example:** `switch(config-vsan-db)# end`	Returns you to EXEC mode.

Interop Mode

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

Related Topics

Switch Interoperability

Displaying the Static VSAN Configuration

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

```
switch# show vsan 100
```

The following example shows how to display VSAN usage:

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

The following example shows how to display all VSANs:

```
switch# show vsan
```

Default Settings for VSANs

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

Table 9: Default VSAN Parameters

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

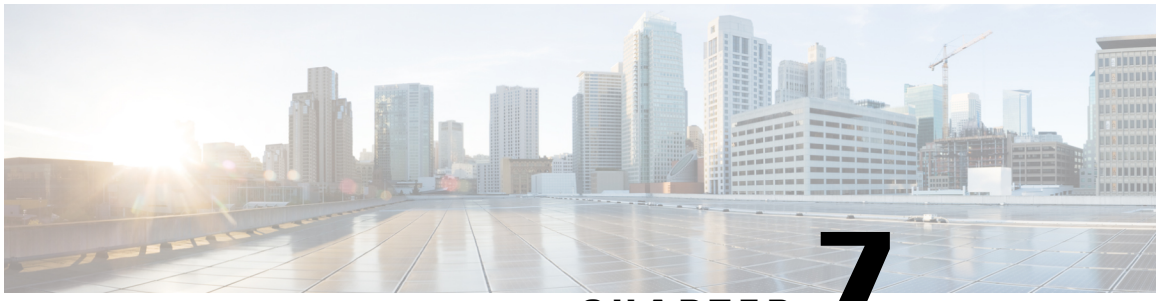

CHAPTER 7

Configuring and Managing Zones

This chapter describes how to configure and manage zones.

This chapter contains the following sections:

Information About Zones

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

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 23: Fabric with Two Zones

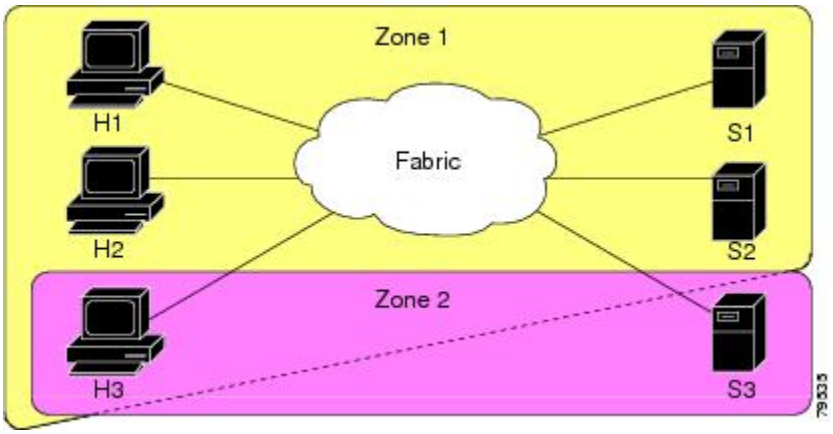

You can use other ways to partition this fabric into zones. The following figure shows another possibility. Assume that there is a need to isolate storage system S2 for the purpose of testing new software. To achieve this, zone 3 is configured, which contains only host H2 and storage S2. You can restrict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

Figure 24: Fabric with Three Zones

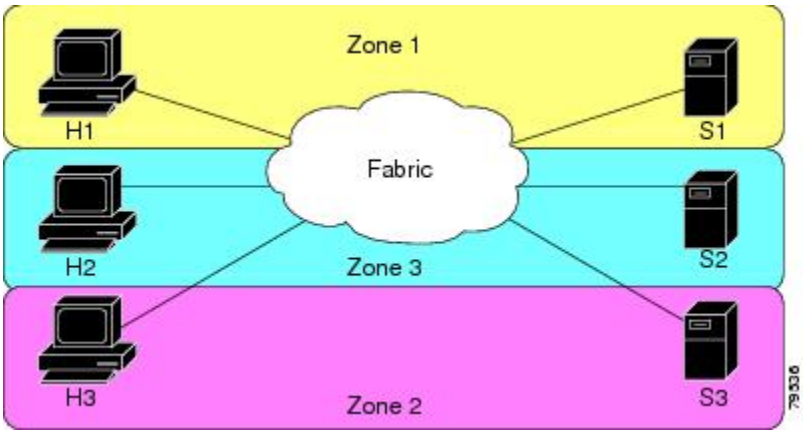

Zone Implementation

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

- Zones are contained in a VSAN.

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

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

- Propagate full zone sets to all switches per VSAN.
- Change the default policy for unzoned members.
- 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 Sets

Before configuring a zone set, consider the following guidelines:

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

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

The following figure shows a zone being added to an activated zone set.

Figure 25: Active and Full Zone Sets

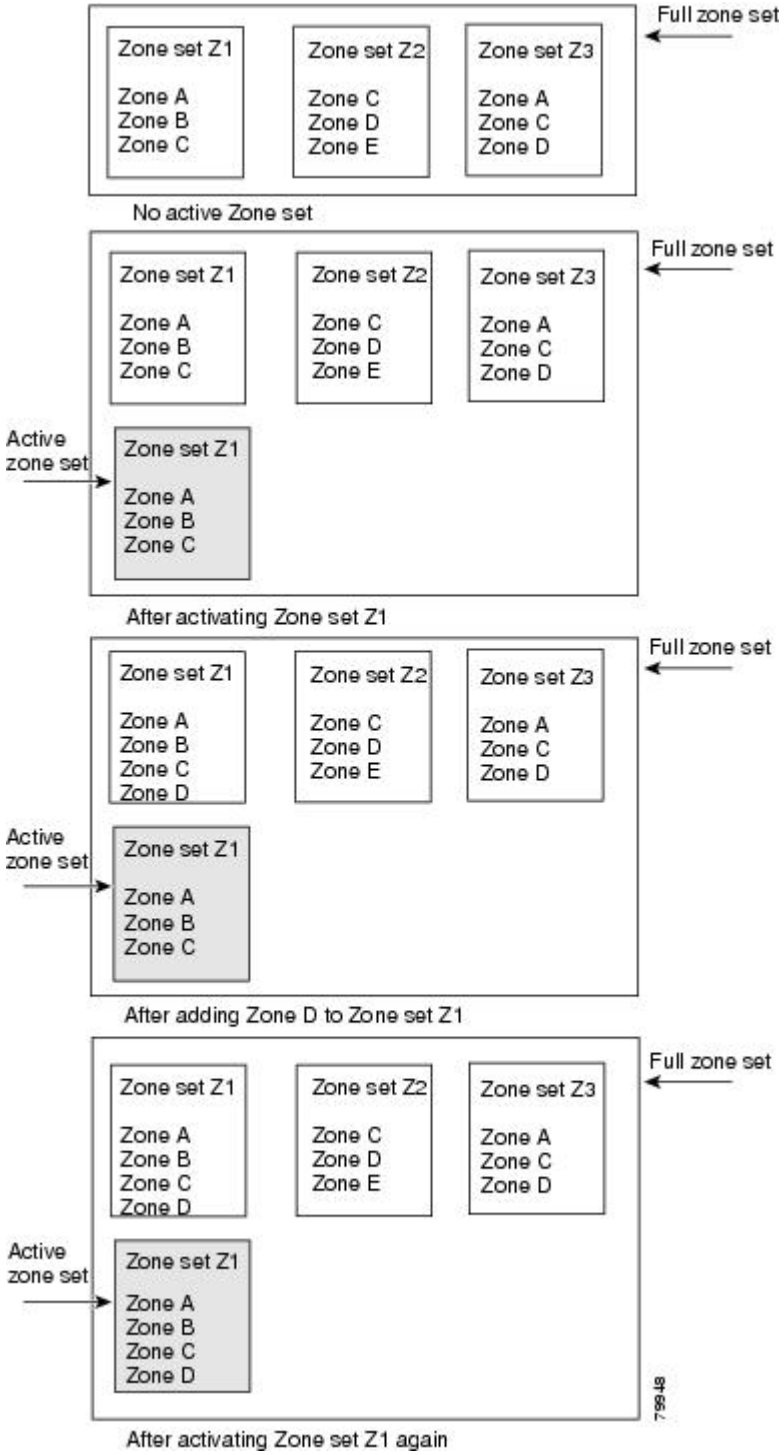

Configuring a Zone

You can configure a zone and assign a zone name.

Procedure

	Command or Action	Purpose

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

Step 2	**zone name** *zone-name* **vsan** *vsan-id* **Example:** `switch(config)# zone name test vsan 5`	Configures a zone in the specified VSAN. **Note** All alphanumeric characters or one of the following symbols ($, -, ^, _) are supported.
Step 3	**member** *type value* **Example:** `switch(config-zone)# member interface 4`	Configures a member for the specified zone based on the type (pWWN, fabric pWWN, FC ID, fcalias, 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, the **show interface** or **show flogi database** commands) to obtain the required value in hex format.

Configuration Examples

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

The following examples show how to configure zone members:

```
switch(config)# zone name MyZone vsan 2
```

pWWN example:

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

Fabric pWWN example:

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

FC ID example:

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

FC alias example:

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

Domain ID example:

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

Show WWN example:

```
switch# show wwn switch
```

Local sWWN interface example:

```
switch(config-zone)# member interface vfc 21
```

Remote sWWN interface example:

```
switch(config-zone)# member interface vfc 21 swwn 20:00:00:05:30:00:4a:de
```

Domain ID interface example:

```
switch(config-zone)# member interface vfc 21 domain-id 25
```

The following example shows how to configure different types of member alias:

```
switch(config)# fcalias name AliasSample vsan 3
```

pWWN example:

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

fWWN example:

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

FC ID example:

```
switch(config-fcalias)# member fcid 0x222222
```

Domain ID example:

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

Device alias example:

```
switch(config-fcalias)# member device-alias devName
```

Zone Sets

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

Figure 26: Hierarchy of Zone Sets, Zones, and Zone Members

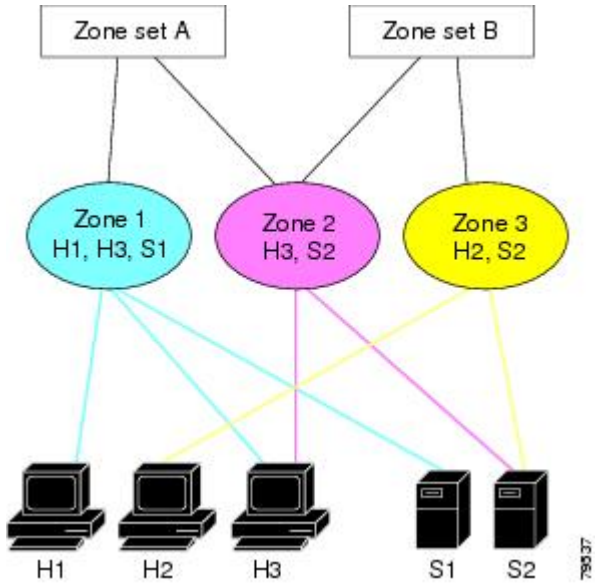

Zones provide a method for specifying access control, while zone sets are a grouping of zones to enforce access control in the fabric. Either zone set A or zone set B can be activated (but not together).

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

Activating a Zone Set

You can activate or deactivate an existing zone set.

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

Procedure

	Command or Action	Purpose

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

Step 2	**zoneset activate name** *zoneset-name* **vsan** *vsan-id* **Example:** `switch(config)# zoneset activate name test vsan 34`	Activates the specified zone set.
Step 3	**no zoneset activate name** *zoneset-name* **vsan** *vsan-id* **Example:** `switch(config)# no zoneset activate name test vsan 30`	Deactivates the specified zone set.

Default Zone

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

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

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

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

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

Note The default settings for default zone configurations can be changed.

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

Configuring the Default Zone Access Permission

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

Procedure

	Command or Action	Purpose

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

Step 2	**zone default-zone permit vsan** *vsan-id* **Example:** `switch(config)# zone default-zone permit vsan 13`	Permits traffic flow to default zone members.
Step 3	**no zone default-zone permit vsan** *vsan-id* **Example:** `switch(config)# no zone default-zone permit vsan 40`	Denies (default) traffic flow to default zone members.

FC Alias Creation

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

- pWWN—The WWN of the N port is in hex format (for example, 10:00:00:23:45:67:89:ab).
- fWWN—The WWN of the fabric port name is in hex format (for example, 10:00:00:23:45:67:89:ab).
- FC ID—The N port ID is in 0xhhhhhh format (for example, 0xce00d1).

- Domain ID—The domain ID is an integer from 1 to 239. A mandatory port number of a non-Cisco switch is required to complete this membership configuration.
- Interface—Interface-based zoning is similar to port-based zoning because the switch interface is used to configure the zone. You can specify a switch interface as a zone member for both local and remote switches. To specify a remote switch, enter the remote switch WWN (sWWN) or the domain ID in the particular VSAN.

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

Creating FC Aliases

You create an alias.

Procedure

	Command or Action	Purpose

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

Step 2	**fcalias name** *alias-name***vsan** *vsan-id* **Example:** `switch(config)# fcalias name testname vsan 50`	Configures an alias name. The alias name can be any case-sensitive, alphanumeric string up to 64 characters.
Step 3	**member** *type value* **Example:** `switch(config-fcalias)# member pwwn 4`	Configures a member for the specified fcalias based on the type (pWWN, fabric pWWN, FC ID, domain ID, or interface) and value specified. **Note** Multiple members can be specified on multiple lines.

Creating FC Aliases Example

Table 10: 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* **Note** If this is a 10G breakout port, the *slot*/*port* syntax is *QSFP-module*/*port*.
Domain ID interface	**member interface** *type slot*/*port* **domain-id** domain-id **Note** If this is a 10G breakout port, the *slot*/*port* syntax is *QSFP-module*/*port*.
Remote sWWN interface	**member interface** *type slot*/*port* **swwn** *swwn-id* **Note** If this is a 10G breakout port, the *slot*/*port* syntax is *QSFP-module*/*port*.
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 vfc 21
```
Remote sWWN interface example:
```
switch(config-fcalias)# member interface vfc 21 swwn 20:00:00:05:30:00:4a:de
```
Domain ID interface example:
```
switch(config-fcalias)# member interface vfc21 domain-id 25
```
Device alias example:
```
switch(config-fcalias)# member device-alias devName
```

Creating Zone Sets and Adding Member Zones

You can create a zone set to include several zones.

Procedure

	Command or Action	Purpose

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

	Command or Action	Purpose
Step 2	**zone set name** *zoneset-name* **vsan** *vsan-id* **Example:** `switch(config)# zone set name new vsan 23`	Configures a zone set with the configured zoneset-name. **Tip** To activate a zone set, you must first create the zone and a zone set.
Step 3	**member** *name* **Example:** `switch(config-zoneset)# member new`	Adds a zone as a member of the previously specified zone set. **Tip** If the specified zone name was not previously configured, this command will return a "zone not present" error message:
Step 4	**zone name** *zone-name* **Example:** `switch(config-zoneset)# zone name trial`	Adds a zone to the specified zone set. **Tip** Execute this step only if you need to create a zone from a zone set prompt.
Step 5	**member fcid** *fcid* **Example:** `switch(config-zoneset-zone)# member fcid 0x222222`	Adds a new member to the new zone. **Tip** Execute this step only if you need to add a member to a zone from a zone set prompt.

Tip You do not have to copy the running configuration to the startup configuration to store the active zone set. However, you need to copy the running configuration to the startup configuration to explicitly store full zone sets.

Zone Enforcement

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

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

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

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

Cisco SAN switches support both hard and soft zoning.

Zone Set Distribution

You can distribute full zone sets using one of two methods: one-time distribution using the **zoneset distribute vsan** command at the EXEC mode level or full zone set distribution using the **zoneset distribute full vsan** command at the configuration mode level. The following table lists the differences between the methods.

Table 11: Zone Set Distribution Differences

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

Enabling Full Zone Set Distribution

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

You can enable full zone set and active zone set distribution to all switches on a per VSAN basis.

Procedure

	Command or Action	Purpose

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

Step 2	**zoneset distribute full vsan** *vsan-id* **Example:** `switch(config)# zoneset distribute full vsan 12`	Enables sending a full zone set along with an active zone set.

Enabling a One-Time Distribution

You can perform a one-time distribution of inactive, unmodified zone sets throughout the fabric.

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

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

This command only distributes the full zone set information, as it does not save the information to the startup configuration. You must explicitly enter the **copy running-config startup-config** command to save the full zone set information to the startup configuration.

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

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

```
switch# show zone status vsan 3
VSAN: 3 default-zone: permit distribute: active only Interop: 100
    mode:basic merge-control:allow

    session:none
    hard-zoning:enabled
Default zone:
    qos:none broadcast:disabled ronly:disabled
Full Zoning Database :
    Zonesets:0  Zones:0 Aliases: 0
Active Zoning Database :
    Name: nozoneset  Zonesets:1  Zones:2
Status: Zoneset distribution completed at 04:01:06 Aug 28 2010
```

Recovering from Link Isolation

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

- Import the neighboring switch's active zone set database and replace the current active zone set (see the figure below).
- Export the current database to the neighboring switch.

- Manually resolve the conflict by editing the full zone set, activating the corrected zone set, and then bringing up the link.

Figure 27: Importing and Exporting the Database

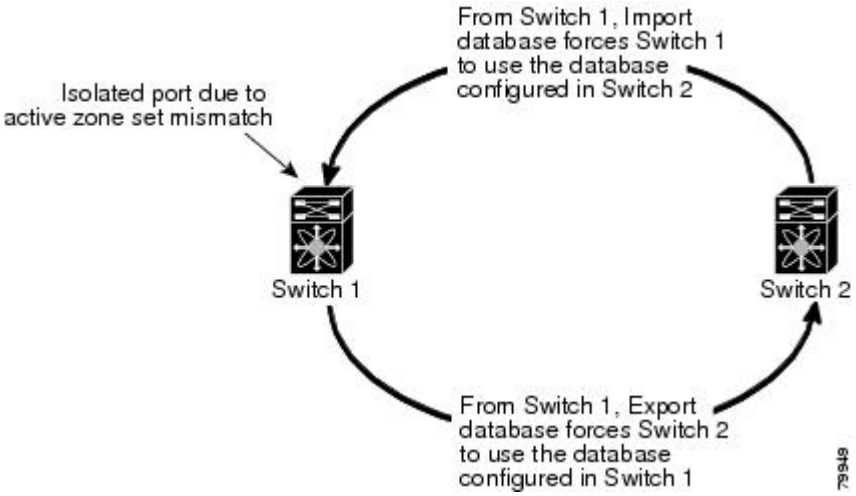

Importing and Exporting Zone Sets

You can import or export the zone set information from or to an adjacent switch.

Procedure

	Command or Action	Purpose
Step 1	switch# **zoneset import interface vfc** *vfc-id* **vsan** *vsan-id*	Imports the zone set from the adjacent switch connected through the specified interface for the VSAN or range of VSANs .
Step 2	**zoneset export vsan** *vsan-id* **Example:** `switch# zoneset export vsan 5`	Exports the zone set to the adjacent switch connected through the specified VSAN or range of VSANs.

Zone Set Duplication

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

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

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

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

Copying Zone Sets

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

Procedure

	Command or Action	Purpose
Step 1	**zone copy active-zoneset full-zoneset vsan** *vsan-id* **Example:** `switch# zone copy active-zoneset full-zoneset vsan 301`	Makes a copy of the active zone set in the specified VSAN to the full zone set.
Step 2	**zone copy vsan** *vsan-id* **active-zoneset scp://guest@myserver/tmp/active_zoneset.txt** **Example:** `switch# zone copy vsan 55 active-zoneset scp://guest@myserver/tmp/active_zoneset.txt`	Copies the active zone in the specified VSAN to a remote location using SCP.

Renaming Zones, Zone Sets, and Aliases

You can rename a zone, zone set, fcalias, or zone-attribute-group.

Procedure

	Command or Action	Purpose

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

Step 2	**zoneset rename** *oldname newname* **vsan** *vsan-id* **Example:** `switch(config)# zoneset rename test myzoneset vsan 60`	Renames a zone set in the specified VSAN.

	Command or Action	Purpose
Step 3	**zone rename** *oldname newname* **vsan** *vsan-id* **Example:** `switch(config)# zone rename test myzone vsan 50`	Renames a zone in the specified VSAN.
Step 4	**fcalias rename** *oldname newname* **vsan** *vsan-id* **Example:** `switch(config)# fcalias rename test myfc vsan 200`	Renames a fcalias in the specified VSAN.
Step 5	**zone-attribute-group rename** *oldname newname* **vsan** *vsan-id* **Example:** `switch(config)#  zone-attribute-group rename test mygroup vsan 12`	Renames a zone attribute group in the specified VSAN.
Step 6	**zoneset activate name** *newname* **vsan** *vsan-id* **Example:** `switch(config)# zoneset activate name myzone vsan 50`	Activates the zone set and updates the new zone name in the active zone set.

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

You can clone a zone, zone set, fcalias, or zone-attribute-group.

Procedure

	Command or Action	Purpose

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

Step 2	**zoneset clone** *oldname newname* **vsan** *vsan-id* **Example:** `switch(config)# zoneset clone test myzoneset2 vsan 2`	Clones a zone set in the specified VSAN.
Step 3	**zone clone** *oldname newname* **vsan** *number* **Example:** `switch(config)# zone clone test myzone3 vsan 3`	Clones a zone in the specified VSAN.

	Command or Action	Purpose
Step 4	**fcalias clone** *oldname newname* **vsan** *vsan-id* **Example:** `switch(config)# fcalias clone test myfcalias vsan 30`	Clones a fcalias in the specified VSAN.
Step 5	**zone-attribute-group clone** *oldname newname* **vsan** *vsan-id* **Example:** `switch(config)# zone-attribute-group clone test mygroup2 vsan 10`	Clones a zone attribute group in the specified VSAN.
Step 6	**zoneset activate name** *newname* **vsan** *vsan-id* **Example:** `switch(config)# zoneset activate name myzonetest1 vsan 3`	Activates the zone set and updates the new zone name in the active zone set.

Clearing the Zone Server Database

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

To clear the zone server database, use the following command:

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

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

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

Verifying the Zone Configuration

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

Command	Purpose
show zone	Displays zone information for all VSANs.
show zone vsan *vsan-id*	Displays zone information for a specific VSAN.
show zoneset vsan *vsan-id* - *vsan-id*	Displays the configured zone sets for a range of VSANs.

Command	Purpose
show zone nam*zone-name*	Displays the members of a specific zone.
show fcalias vsan *vsan-id*	Displays the fcalias configuration.
show zone member pwwn *pwwn-id*	Displays all zones to which a member belongs.
show zone statistics	Displays the number of control frames exchanged with other switches.
show zoneset active	Displays the active zone set.
show zone active	Displays the active zones.
show zone status	Displays the zone status.

Enhanced Zoning

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.

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.

The following table lists the advantages of the enhanced zoning feature in all switches in the Cisco SAN switches.

Table 12: Advantages of Enhanced Zoning

Basic Zoning	Enhanced Zoning	Enhanced Zoning Advantages
Administrators can make simultaneous configuration changes. Upon activation, one administrator can overwrite another administrator's changes.	Performs all configurations within a single configuration session. When you begin a session, the switch locks the entire fabric to implement the change.	One configuration session for the entire fabric to ensure consistency within the fabric.
If a zone is part of multiple zone sets, you create an instance of this zone in each zone set.	References to the zone are used by the zone sets as required once you define the zone.	Reduced payload size as the zone is referenced. The size is more significant with bigger databases.
The default zone policy is defined per switch. To ensure smooth fabric operation, all switches in the fabric must have the same default zone setting.	Enforces and exchanges the default zone setting throughout the fabric.	Fabric-wide policy enforcement reduces troubleshooting time.

Basic Zoning	Enhanced Zoning	Enhanced Zoning Advantages
To retrieve the results of the activation per switch basis, the managing switch provides a combined status about the activation. It does not identify the failure switch.	Retrieves the activation results and the nature of the problem from each remote switch.	Enhanced error reporting eases the troubleshooting process
To distribute the zoning database, you must reactivate the same zone set. The reactivation may affect hardware changes for hard zoning on the local switch and on remote switches.	Implements changes to the zoning database and distributes it without reactivation.	Distribution of zone sets without activation avoids hardware changes for hard zoning in the switches.
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

You can change to the enhanced zoning mode from the basic mode.

Procedure

Step 1 Verify that all switches in the fabric can operate in the enhanced mode.

Step 2 If one or more switches cannot operate in the enhanced mode, then your request to move to enhanced mode is rejected.

Step 3 Set the operation mode to enhanced zoning mode.

Changing from Enhanced Zoning to Basic Zoning

Cisco 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.

Procedure

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

Step 2 If such configurations exist, delete them before proceeding with this procedure. If you do not delete the existing configuration, the switch software automatically removes them.

Step 3 Set the operation mode to basic zoning mode.

Enabling Enhanced Zoning

You can enable enhanced zoning in a VSAN.

By default, the enhanced zoning feature is disabled in all Cisco SAN switches.

Procedure

	Command or Action	Purpose

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

Step 2	**zone mode enhanced vsan** *vsan-id* **Example:** `switch(config)# zone mode enhanced vsan 22`	Enables enhanced zoning in the specified VSAN.
Step 3	**no zone mode enhanced vsan** *vsan-id* **Example:** `switch(config)# no zone mode enhanced vsan 30`	Disables enhanced zoning in the specified VSAN.

Modifying the Zone Database

Ypou can commit or discard changes to the zoning database in a VSAN.

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

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

Procedure

	Command or Action	Purpose

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

Step 2	**zone commit vsan** *vsan-id* **Example:** `switch(config)# zone commit vsan 679`	Applies the changes to the enhanced zone database and closes the session.
Step 3	switch(config)# **zone commit vsan** *vsan-id* **force** **Example:** `switch(config)# zone commit vsan 34 force`	Forcefully applies the changes to the enhanced zone database and closes the session created by another user.
Step 4	switch(config)# **no zone commit vsan** *vsan-id* **Example:** `switch(config)# no zone commit vsan 22`	Discards the changes to the enhanced zone database and closes the session.
Step 5	**no zone commit vsan** *vsan-id* **force** **Example:** `switch(config)# no zone commit vsan 34 force`	Forcefully discards the changes to the enhanced zone database and closes the session created by another user.

Releasing Zone Database Locks

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

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

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

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

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

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 13: Database Zone Merge Status

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

The merge process operates as follows:

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

Configuring Zone Merge Control Policies

You can configure merge control policies.

Procedure

	Command or Action	Purpose

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

Step 2	**zone merge-control restrict vsan** *vsan-id* **Example:** `switch(config)# zone merge-control restrict vsan 24`	Configures a restricted merge control setting for this VSAN.
Step 3	**no zone merge-control restrict vsan** *vsan-id* **Example:** `switch(config)# no zone merge-control restrict vsan 33`	Defaults to using the allow merge control setting for this VSAN.
Step 4	**zone commit vsan** *vsan-id* **Example:** `switch(config)# zone commit vsan 20`	Commits the changes made to the specified VSAN.

Default Zone Policies

You can permit or deny traffic in the default zone.

Procedure

	Command or Action	Purpose

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

Step 2	**zone default-zone permit vsan** *vsan-id* **Example:** `switch(config)# zone default-zone permit vsan 12`	Permits traffic flow to default zone members.
Step 3	**no zone default-zone permit vsan** *vsan-id* **Example:** `switch(config)# no zone default-zone permit vsan 12`	Denies traffic flow to default zone members and reverts to factory default.

	Command or Action	Purpose
Step 4	**zone commit vsan** *vsan-id* **Example:** `switch(config)# zone commit vsan 340`	Commits the changes made to the specified VSAN.

Configuring System Default Zoning Settings

You can configure default settings for default zone policies and full zone distribution for new VSANs on the switch.

Procedure

	Command or Action	Purpose

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

Step 2	**system default zone default-zone permit** **Example:** `switch(config)# system default zone default-zone permit`	Configures permit as the default zoning policy for new VSANs on the switch.
Step 3	**no system default zone default-zone permit** **Example:** `switch(config)# no system default zone default-zone permit`	Configures deny (default) as the default zoning policy for new VSANs on the switch.
Step 4	**system default zone distribute full** **Example:** `switch(config)# system default zone distribute full`	Enables full zone database distribution as the default for new VSANs on the switch.
Step 5	**no system default zone distribute full** **Example:** `switch(config)# no system default zone distribute full`	Disables (default) full zone database distribution as the default for new VSANs on the switch. Only the active zone database is distributed.

Verifying Enhanced Zone Information

This 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.

Procedure

	Command or Action	Purpose

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

Step 2	**no zone name** *zone-name* **vsan** *vsan-id* **Example:** `switch(config)# no zone name myzone vsan 35`	Deletes a zone to reduce the number of zones to 2000 or fewer.
Step 3	**zone compact vsan** *vsan-id* **Example:** `switch(config)# zone compact vsan 42`	Compacts the zone database for the specified VSAN to recover the zone ID released when a zone was deleted.

Analyzing the Zone and Zone Set

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 command reference for your device for the description of the information displayed in the command output.

Default Settings for Zones

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

Table 14: 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.

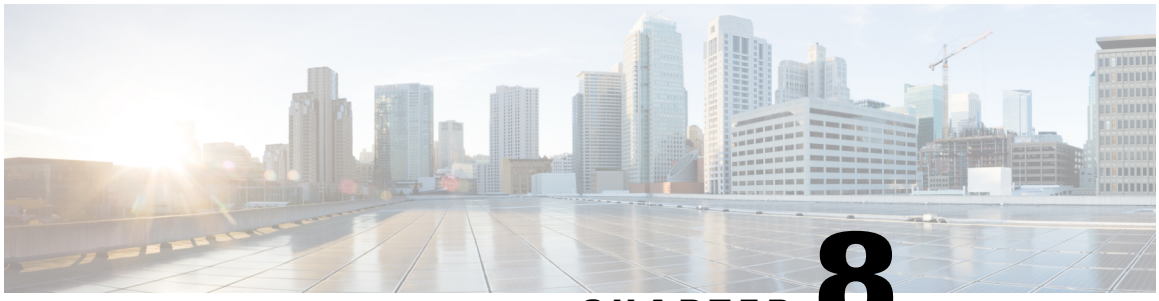

CHAPTER 8

Distributing Device Alias Services

This chapter describes how to distribute device alias services.

This chapter contains the following sections:

Distributing Device Alias Services

Cisco SAN switches support Distributed Device Alias Services (device aliases) on a fabric-wide basis.

Information About Device Aliases

Cisco SAN switches support Distributed Device Alias Services (device aliases) on a fabric-wide basis.

When the port WWN (pWWN) of a device must be specified to configure features (for example, zoning, DPVM, or port security) in a Cisco SAN switch, you must assign the correct device name each time you configure these features. An inaccurate device name may cause unexpected results. You can circumvent this problem if you define a user-friendly name for a pWWN and use this name in all the configuration commands as required. These user-friendly names are referred to as *device aliases*.

Device Alias Features

Device aliases have the following features:

- The device alias information is independent of the VSAN configuration.
- The device alias configuration and distribution is independent of the zone server and the zone server database.
- 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 System Management Configuration Guide for your device.

Related Topics

Device Alias Modes, on page 110

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 15: 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.

Zone-Based Aliases	Device Aliases
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

You can create a device alias in the pending database.

Procedure

	Command or Action	Purpose

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

Step 2	**device-alias database**
	
	Example:
	`switch(config)# device-alias database`
	`switch(config-device-alias-db)#`
		Enters the pending database configuration submode.

Step 3	**device-alias name** *device-name* **pwwn** *pwwn-id*
	
	Example:
	`switch(config-device-alias-db)#`
	`device-alias name mydevice pwwn`
	`21:01:00:e0:8b:2e:80:93`
		Specifies a device name for the device that is identified by its pWWN. Starts writing to the pending database and simultaneously locks the fabric as this is the first-issued device alias configuration command.

Step 4	**no device-alias name** *device-name*
	
	Example:
	`switch(config-device-alias-db)# no`
	`device-alias name mydevice`
		Removes the device name for the device that is identified by its pWWN.

	Command or Action	Purpose
Step 5	**device-alias rename** *old-device-name new-device-name* **Example:** `switch(config-device-alias-db)# device-alias rename mydevice mynewdevice`	Renames an existing device alias with a new name.

EXAMPLES

This example shows how to display the device alias configuration.

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

Device Alias Modes

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 fails in interop mode VSANs if the corresponding zones have native device alias-based members.

Device Alias Mode Guidelines and Limitations for Device Alias Services

Device Alias services have these configuration guidelines and limitations:

- If two fabrics running in different device alias modes are joined together, the device alias merge fails. There is no automatic conversion to one mode or the other during the merge process. In this situation, you must select one mode over the other.
- Before changing from enhanced to basic mode, you must first explicitly remove all native device alias-based configurations from both local and remote switches, or replace all device alias-based configuration members with the corresponding pWWN.

- If you remove a device alias from the device alias database, all applications automatically stop enforcing the corresponding device alias. If that corresponding device alias is part of an active zone set, all the traffic to and from that pWWN is disrupted.
- Renaming the device alias not only changes the device alias name in the device alias database, but also replaces the corresponding device alias configuration in all of the applications.
- When a new device alias is added to the device alias database, and the application configuration is present on that device alias, it automatically takes effect. For example, if the corresponding device alias is part of the active zoneset and the device is online, then zoning is enforced automatically. You do not have to reactivate the zone set.
- If a device alias name is mapped to a new HBA's pWWN, the application's enforcement changes accordingly. In this case, the zone server automatically enforces zoning based on the new HBA's pWWN.

Configuring Device Alias Modes

You can configure device aliases to operate in enhanced mode.

Procedure

	Command or Action	Purpose

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

Step 2	**device-alias mode enhanced** **Example:** `switch(config)# device-alias mode enhanced`	Assigns the device alias to operate in enhanced mode.
Step 3	**no device-alias mode enhance** **Example:** `switch(config)# no device-alias mode enhance`	Assigns the device alias to operate in basic mode.

EXAMPLES

This example shows how to display the current device alias mode setting.

```
switch# show device-alias status
Fabric Distribution: Enabled
Database:- Device Aliases 0 Mode: Basic
Locked By:- User "admin" SWWN 20:00:00:0d:ec:30:90:40
Pending Database:- Device Aliases 0 Mode: Basic
```

Device Alias Distribution

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

If device alias distribution is disabled, database changes are not distributed to the switches in the fabric. The same changes would have to be performed manually on all switches in the fabric to keep the device alias database up-to-date. Database changes immediately take effect, so there would also not be any pending database and commit or abort operations. If you have not committed the changes and you disable distribution, a commit task fails.

This example shows how to display a failed device alias status:

```
switch# show device-alias status
Fabric Distribution: Disabled
Database:- Device Aliases 25
Status of the last CFS operation issued from this switch:
=========================================================
Operation: Commit
Status: Failed (Reason: Operation is not permitted as the fabric distribution is
 currently disabled.)
```

Locking the Fabric

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

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

Committing Changes

You can commit changes.

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

- The pending database content overwrites the effective database content.
- The pending database is distributed to the switches in the fabric and the effective database on those switches is overwritten with the new changes.
- The pending database is emptied of its contents.
- The fabric lock is released for this feature.

Procedure

	Command or Action	Purpose

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

Step 2	**device-alias commit** **Example:** `switch(config)# device-alias commit`	Commits the changes made to the currently active session.

Discarding Changes

You can discard the device alias session changes.

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

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

Procedure

	Command or Action	Purpose

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

Step 2	**device-alias abort** **Example:** `switch(config)# device-alias abort`	Discards the currently active session.

EXAMPLES

This example shows how to display the status of the discard operation:

```
switch# 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
```

Overriding the Fabric Lock

You can use locking operations (clear, commit, abort) only when device alias distribution is enabled. If you have performed a device alias task and have forgotten to release the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your changes to the pending database are discarded and the fabric lock is released.

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

To use administrative privileges and release a locked device alias session, use the **clear device-alias session** command in EXEC mode.

```
switch# clear device-alias session
```

This example shows how to display the status of the clear operation:

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

Disabling and Enabling Device Alias Distribution

You can disable or enable the device alias distribution.

Procedure

	Command or Action	Purpose

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

Step 2	**no device-alias distribute** **Example:** `switch(config)# no device-alias distribute`	Disables the distribution.
Step 3	**device-alias distribute** **Example:** `switch(config)# device-alias distribute`	Enables the distribution (default).

EXAMPLES

This example shows how to display the status of device alias distribution:

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

This 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
```

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

You can import the zone alias for a specific VSAN.

Procedure

	Command or Action	Purpose

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

Step 2	**device-alias import fcalias vsan** *vlan-id* **Example:** `switch(config)# device-alias import fcalias vsan`	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 System Management Configuration Guide for your device.

Verifying the Device Alias Configuration

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

Command	Purpose
show zoneset [**active**]	Displays the device aliases in the zone set information.
show device-alias database [**pending** | **pending-diffs**]	Displays the device alias database.
show device-alias {**pwwn** *pwwn-id* | **name** *device-name* } [**pending**]	Displays the device alias information for the specified pwwn or alias.

Command	Purpose
show flogi database [**pending**]	Displays device alias information in the flogi database.
show fcns database [**pending**]	Displays device alias information in the fcns database.

Default Settings for Device Alias Services

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

Table 16: 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.

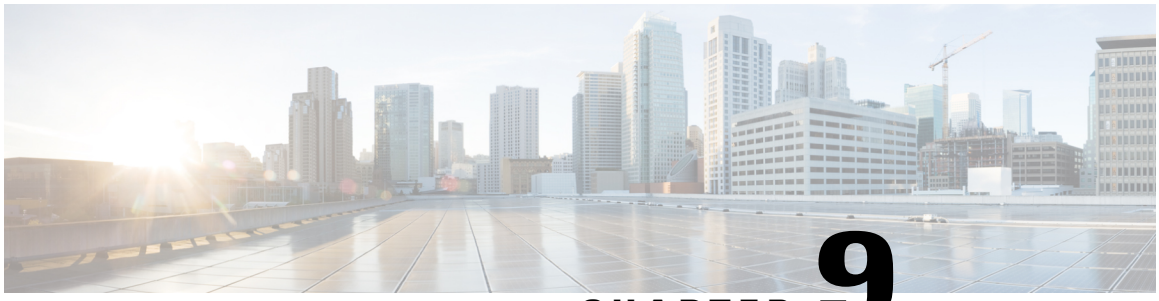

CHAPTER 9

Managing FLOGI, Name Server, FDMI, and RSCN Databases

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

This chapter includes the following sections:

Managing FLOGI, Name Server, FDMI, and RSCN Databases

Fabric Login

In a Fibre Channel fabric, each host or disk requires an FC ID. Use the **show flogi** command to verify if a storage device is displayed in the fabric login (FLOGI) table as in the following examples. If the required device is displayed in the FLOGI table, the fabric login is successful. Examine the FLOGI database on a switch that is directly connected to the host HBA and connected ports.

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

```
switch# show flogi database
--------------------------------------------------------------------------------
INTERFACE  VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------------
vfc23      1     0xb200e2  21:00:00:04:cf:27:25:2c  20:00:00:04:cf:27:25:2c
vfc23      1     0xb200e1  21:00:00:04:cf:4c:18:61  20:00:00:04:cf:4c:18:61
vfc23      1     0xb200d1  21:00:00:04:cf:4c:18:64  20:00:00:04:cf:4c:18:64
vfc23      1     0xb200ce  21:00:00:04:cf:4c:16:fb  20:00:00:04:cf:4c:16:fb
vfc23      1     0xb200cd  21:00:00:04:cf:4c:18:f7  20:00:00:04:cf:4c:18:f7
vfc31      2    0xb30100  10:00:00:05:30:00:49:63  20:00:00:05:30:00:49:5e
Total number of flogi = 6.
```

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

```
switch# show flogi database interface vfc1/1
INTERFACE  VSAN    FCID           PORT NAME               NODE NAME

--------------------------------------------------------------------------------
vfc1/1     1     0x870000  20:00:00:1b:21:06:58:bc  10:00:00:1b:21:06:58:bc
Total number of flogi = 1.
```

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

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

Name Server Proxy

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

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

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

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

About Registering Name Server Proxies

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

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

Registering Name Server Proxies

You can register the name server proxy.

Procedure

	Command or Action	Purpose

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

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

Rejecting Duplicate pWWNs

By FC standard, NX-OS will accept a login on any interface of a pwwn that is already logged in on the same switch, same vsan, same fcdomain. To prevent the same pwwn from logging in the same switch on a different interface, use the port security feature.

By default, any future flogi (with duplicate pwwn) on different switch in the same vsan, will be rejected and earlier FLOGI retained, which does not follow FC standards.

If you disable this option, any future flogi (with duplicate pwwn) on different switch in the same VSAN,will be allowed to succeed by deleting earlier FCNS entry.

Rejecting Duplicate pWWNs

By FC standard, NX-OS will accept a login on any interface of a pwwn that is already logged in on the same switch, same vsan, same fcdomain. To prevent the same pwwn from logging in the same switch on a different interface, use the port security feature.

By default, any future flogi (with duplicate pwwn) on different switch in the same vsan, will be rejected and earlier FLOGI retained, which does not follow FC standards.

If you disable this option, any future flogi (with duplicate pwwn) on different switch in the same VSAN,will be allowed to succeed by deleting earlier FCNS entry.

To reject duplicate pWWNs, follow these steps:

Procedure

	Command or Action	Purpose

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

Step 2	**fcns reject-duplicate-pwwn vsan** *vsan-id* **Example:** `switch(config)# fcns reject-duplicate-pwwn vsan 100`	Any future flogi (with duplicate pwwn) on different switch, will be rejected and earlier FLOGI retained (default).
Step 3	**no fcns reject-duplicate-pwwn vsan** *vsan-id* **Example:** `switch(config)# no fcns reject-duplicate-pwwn vsan 256`	Any future flogi (with duplicate pwwn) on different switch, will be allowed to succeed by deleting earlier FCNS entry. But you can still see the earlier entry in FLOGI database in the other switch.

Name Server Database Entries

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

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

Displaying Name Server Database Entries

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

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

This 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
```

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

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

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

```
switch# show fcns statistics
```

FDMI

Cisco SAN 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

This 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

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.

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.

Multi-pid Option

If the RSCN multi-pid option is enabled, RSCNs generated to the registered Nx ports might contain more than one affected port IDs. In this case, zoning rules are applied before putting the multiple affected port IDs together in a single RSCN. By enabling this option, you can reduce the number of RSCNs. For example, you have two disks (D1, D2) and a host (H) connected to switch 1. Host H is registered to receive RSCNs. D1, D2, and H belong to the same zone. If disks D1 and D2 are online at the same time, one of the following actions applies:

- The multi-pid option is disabled on switch 1— Two RSCNs are generated to host H: one for the disk D1 and another for disk D2.
- The multi-pid option is enabled on switch 1—A single RSCN is generated to host H, and the RSCN payload lists the affected port IDs (in this case, both D1 and D2).

Note Some Nx ports may not support multi-pid RSCN payloads. If so, disable the RSCN multi-pid option.

Configuring the multi-pid Option

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

Procedure

	Command or Action	Purpose

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

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

Suppressing Domain Format SW-RSCNs

A domain format SW-RSCN is sent whenever the local switch name or the local switch management IP address changes. This SW-RSCN is sent to all other domains and switches over the ISLs. The remote switches can issue GMAL and GIELN commands to the switch that initiated the domain format SW-RSCN to determine what changed. Domain format SW-RSCNs can cause problems with some non-Cisco SAN switches.

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

Procedure

	Command or Action	Purpose

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

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

Clearing RSCN Statistics

You can clear the counters and later view the counters for a different set of events. For example, you can keep track of how many RSCNs or SW-RSCNs are generated on a particular event (such as ONLINE or OFFLINE events). You can use these statistics to monitor responses for each event in the VSAN.

This example shows how to clear the RSCN statistics for the specified VSAN:

```
switch# clear rscn statistics vsan 1
```

After clearing the RSCN statistics, you can view the cleared counters by entering the **show rscn statistics** command:

```
switch# show rscn statistics vsan 1
```

Configuring the RSCN Timer

RSCN maintains a per VSAN event list queue, where the RSCN events are queued as they are generated. When the first RSCN event is queued, a per VSAN timer starts. When a timeout occurs, all the events are dequeued and coalesced RSCNs are sent to registered users. The default timer values minimize the number of coalesced RSCNs that are sent to registered users. Some deployments require smaller event timer values to track changes in the fabric.

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

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

You can configure the RSCN timer.

Procedure

	Command or Action	Purpose

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

Step 2	**rscn distribute** **Example:** `switch(config)# rscn distribute`	Enables RSCN timer configuration distribution.
Step 3	**rscn event-tov** *timeout* **vsan** *vsan-id* **Example:** `switch(config)# rscn event-tov 1000 vsan 501`	Sets the event time-out value in milliseconds for the specified VSAN. The range is 0 to 2000 milliseconds. Setting a zero (0) value disables the timer.
Step 4	**no rscn event-tov** *timeout* **vsan** *vsan-id* **Example:** `switch(config)# no rscn event-tov 1100 vsan 245`	Reverts to the default value (2000 milliseconds for Fibre Channel VSANs).
Step 5	**rscn commit vsan** *vsan-id* **Example:** `switch(config)# rscn commit vsan 25`	Commits the RSCN timer configuration to be distributed to the switches in the specified VSAN.

Verifying the RSCN Timer Configuration

You verify the RSCN timer configuration using the **show rscn event-tov vsan** command. This example shows how to clear the RSCN statistics for VSAN 10:

```
switch# show rscn event-tov vsan 10
Event TOV : 1000 ms
```

RSCN Timer Configuration Distribution

Because the timeout value for each switch is configured manually, a misconfiguration occurs when different switches time out at different times. Different N-ports in a network can receive RSCNs at different times. Cisco Fabric Services (CFS) infrastructure alleviates this situation by automatically distributing the RSCN timer configuration information to all switches in a fabric, which also reduces the number of SW-RSCNs.

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

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

Caution Only the RSCN timer configuration is distributed.

The RSCN timer is registered with CFS during initialization and switchover. For high availability, if the RSCN timer distribution crashes and restarts or a switchover occurs, it resumes normal functionality from the state prior to the crash or switchover.

For additional information, refer to Using Cisco Fabric Services in the System Management Configuration Guide for your device.

Enabling RSCN Timer Configuration Distribution

You can enable RSCN timer configuration distribution.

Procedure

	Command or Action	Purpose

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

Step 2	**rscn distribute** **Example:** `switch(config)# rscn distribute`	Enables RSCN timer distribution.
Step 3	**no rscn distribute** **Example:** `switch(config)# no rscn distribute`	Disables (default) RSCN timer distribution.

Locking the Fabric

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

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

Committing RSCN Timer Configuration Changes

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

You can commit RSCN timer configuration changes.

Procedure

	Command or Action	Purpose

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

Step 2	**rscn commit vsan** *timeout* **Example:** `switch(config)# rscn commit vsan 500`	Commits the RSCN timer changes.

Discarding the RSCN Timer Configuration Changes

If you discard (abort) the changes made to the pending database, the configuration database remains unaffected and the lock is released.

You can discard RSCN timer configuration changes.

Procedure

	Command or Action	Purpose

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

Step 2	**rscn abort vsan** *timeout* **Example:** `switch(config)# rscn abort vsan 800`	Discards the RSCN timer changes and clears the pending configuration database.

Clearing a Locked Session

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

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

To use administrative privileges and release a locked DPVM session, use the **clear rscn session vsan** command in EXEC mode. This example shows how to clear the RSCN session for VSAN 10:

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

Displaying RSCN Configuration Distribution Information

This example shows how to display the registration status for RSCN configuration distribution:

```
switch# show cfs application name rscn
 Enabled        : Yes
 Timeout        : 5s
 Merge Capable  : Yes
 Scope          : Logical
```

Note A merge failure results when the RSCN timer values are different on the merging fabrics.

This example shows how to display the set of configuration commands that would take effect when you commit the configuration:

Note The pending database includes both existing and modified configuration.

```
switch# show rscn pending
rscn event-tov 2000 ms vsan 1
rscn event-tov 2000 ms vsan 2
rscn event-tov 300 ms vsan 10
```

This 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 Settings for RSCN

The following table lists the default settings for RSCN.

Table 17: Default RSCN Settings

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

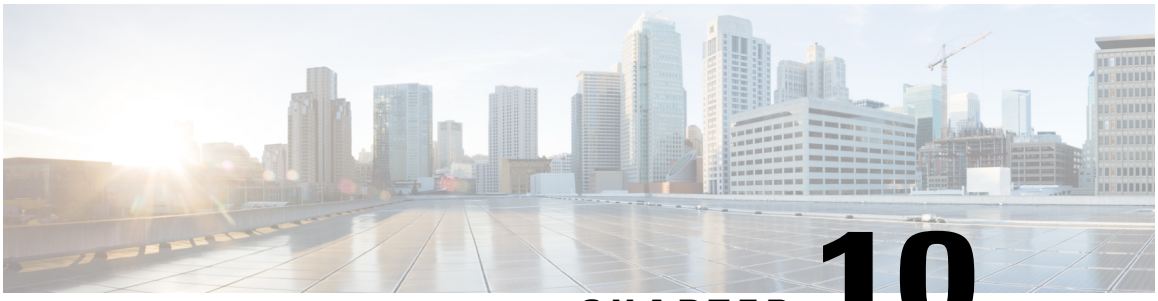

CHAPTER 10

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 are Cisco Nexus devices.

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:

Procedure

	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 shows how to discover SCSI targets for the specified VSAN (1) and FCID (0x9c03d6):

```
switch# discover scsi-target vsan 1 fcid 0x9c03d6 os aix
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...
```

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:

Procedure

	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
```

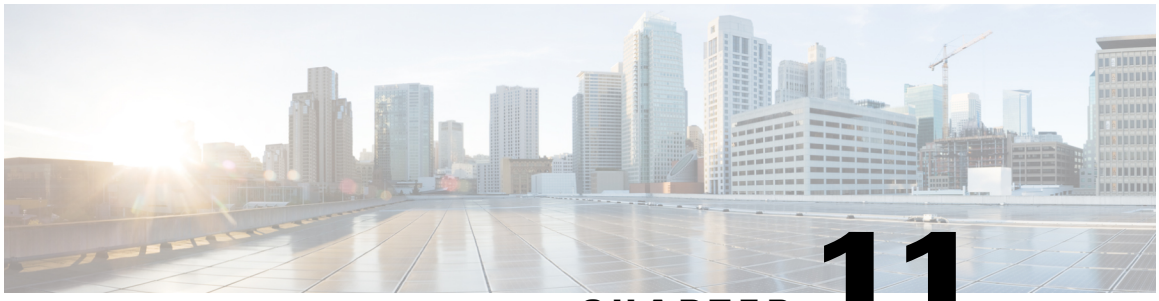

CHAPTER 11

Configuring FC-SP and DHCHAP

This chapter describes how to configure the Fibre Channel Security Protocol (FC-SP) and the Diffie-Hellman Challenge Handshake Authentication Protocol (DHCP).

This chapter includes the following sections:

Information About FC-SP and DHCHAP

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

Fabric Authentication

All Cisco SAN 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 or accidentally interconnect incompatible switches, resulting in Inter-Switch Link (ISL) isolation and link disruption.

Cisco SAN switches support authentication features to address physical security (see the following figure).

Figure 28: Switch and Host Authentication

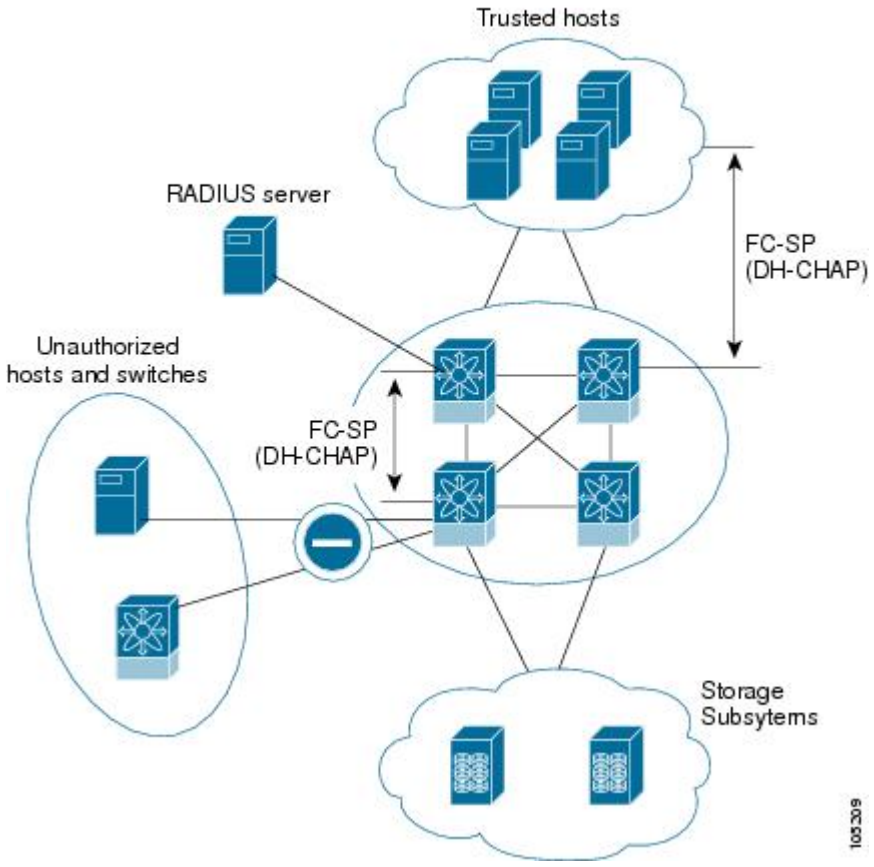

Note Fibre Channel host bus adapters (HBAs) with appropriate firmware and drivers are required for host-switch authentication.

Configuring DHCHAP Authentication

You can configure DHCHAP authentication using the local password database.

Before You Begin

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.

Procedure

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

When configuring the DHCHAP feature along with existing Cisco NX-OS features, consider these compatibility issues:

- 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.

By default, the DHCHAP feature is disabled in all Cisco SAN switches.

About Enabling DHCHAP

By default, the DHCHAP feature is disabled in all Cisco SAN 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

You can enable DHCHAP for a Cisco Nexus device.

Procedure

	Command or Action	Purpose

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

	Command or Action	Purpose
Step 2	**feature fcsp** **Example:** `switch(config)# feature fcsp`	Enables the DHCHAP in this switch.
Step 3	**no feature fcsp** **Example:** `switch(config)# no feature fcsp`	Disables (default) the DHCHAP in this switch.

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 18: DHCHAP Authentication Status Between Two SAN 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

You can configure the DHCHAP mode for a particular interface.

Procedure

	Command or Action	Purpose

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

Step 2	switch(config)# **interface vfc** *vfc-id* - *vfc-id*	Selects a range of interfaces and enters the interface configuration mode.
Step 3	**fcsp on** **Example:** `switch(config-if)# fcsp on`	Sets the DHCHAP mode for the selected interfaces to be in the on state.
Step 4	**no fcsp on** **Example:** `switch(config-if)# no fcsp on`	Reverts to the factory default of auto-passive for these three interfaces.
Step 5	**fcsp auto-active 0** **Example:** `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. **Note** The reauthorization interval configuration is the same as the default behavior.

	Command or Action	Purpose
Step 6	**fcsp auto-active** *timeout-period* **Example:** `switch(config-if)# fcsp auto-active 10`	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	**fcsp auto-active** **Example:** `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).

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

You can configure the hash algorithm.

Procedure

	Command or Action	Purpose

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

Step 2	**fcsp dhchap hash** [**md5**] [**sha1**] **Example:** `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	**no fcsp dhchap hash sha1** **Example:** `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.

DHCHAP Group Settings

All Cisco SAN 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

You can change the DH group settings.

Procedure

	Command or Action	Purpose

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

Step 2	**fcsp dhchap dhgroup** [**0** | **1** | **2** | **3** | **4**] **Example:** `switch(config)# fcsp dhchap dhgroup [0|1|2|3|4]`	Prioritizes the use of DH groups in the configured order.
Step 3	**no fcsp dhchap dhgroup** [**0** | **1** | **2**| **3** |]**4**] **Example:** `switch(config)# no fcsp dhchap dhgroup [0|1|2|3|4]`	Reverts to the DHCHAP factory default order of 0, 1, 2, 3 and 4.

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 Cisco MDS 9000 Family Fabric Manager to manage the password database.

Configuring DHCHAP Passwords for the Local Switch

You can configure the DHCHAP password for the local switch.

Procedure

	Command or Action	Purpose

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

Step 2	**fcsp dhchap password** [**0** | **7**] *password* [**wwn** *wwn-id*]
	
	Example:
	`switch(config)# fcsp dhchap password [0|7]`
	`myword wwn 11:22:11:22:33:44:33:44`
		Configures a clear text password for the local switch.

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

You can locally configure the remote DHCHAP password for another switch in the fabric.

Procedure

	Command or Action	Purpose

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

Step 2	**fcsp dhchap devicename** *switch-wwn* **password** *password* **Example:** `switch(config)# fcsp dhchap devicename 21:00:05:30:23:1a:11:03 password mypassword`	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* **Example:** `switch(config)# no fcsp dhchap devicename 21:00:05:30:23:1a:11:03 password mypassword`	Removes the password entry for this switch from the local authentication database.

DHCHAP Timeout Value

During the DHCHAP protocol exchange, if the 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

You can configure the DHCHAP timeout value.

Procedure

	Command or Action	Purpose

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

Step 2	**fcsp timeout** *timeout* **Example:** `switch(config)# fcsp timeout 60`	Configures the reauthentication timeout to the specified value. The unit is seconds.
Step 3	**no fcsp timeout** *timeout* **Example:** `switch(config)# no fcsp timeout 60`	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 vfc24
vfc24
        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 vfc24 statistics
```

The following example shows how to display the FC-SP WWN of the device connected to the specified interface:

```
switch# show fcsp interface vfc21 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.

Configuration Examples for Fabric Security

This section provides the steps to configure the example illustrated in the following figure.

Figure 29: Sample DHCHAP Authentication

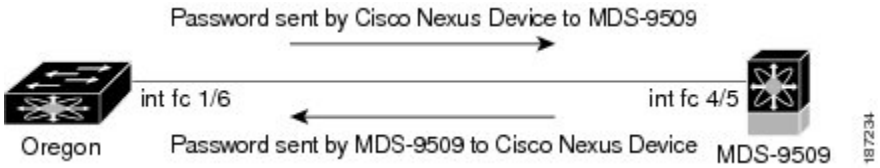

This example shows how to set up authentication:

Procedure

Step 1 Obtain the device name of the Cisco SAN switch in the fabric. The Cisco SAN 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 is 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 interface.

Note Whenever DHCHAP port mode is changed to a mode other than the Off mode, reauthentication is performed.

Example:
```
switch(config)# interface vfc24
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 interface.

Example:

```
switch# show fcsp interface vfc24
vfc24
        fcsp authentication mode:SEC_MODE_ON
        Status:Successfully authenticated
```

Step 8 Repeat these steps on the connecting 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 vfc 45
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 fc24
Fc24
        fcsp authentication mode:SEC_MODE_ON
        Status:Successfully authenticated
```

You have now enabled and configured DHCHAP authentication for the sample setup.

Default Settings for Fabric Security

The following table lists the default settings for all fabric security features in any switch.

Table 19: 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

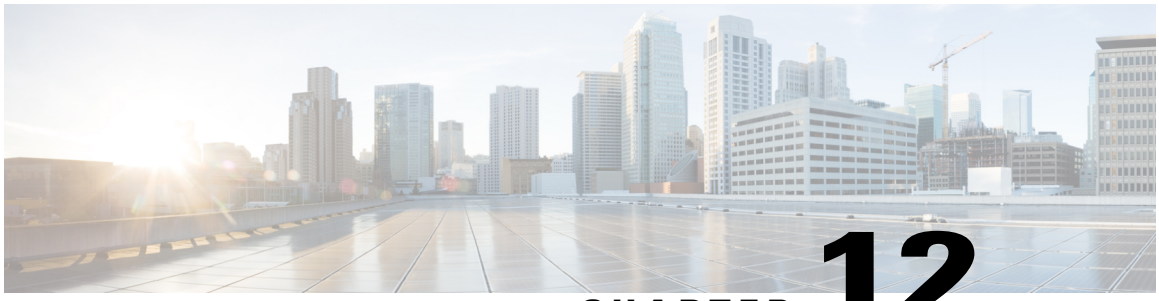

CHAPTER 12

Configuring Port Security

This chapter describes how to configure port security.

This chapter includes the following sections:

Configuring Port Security

Cisco SAN 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, 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.

Note Port security is supported on virtual Fibre Channel ports and physical Fibre Channel ports.

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.

Auto-Learning

You can instruct the switch to automatically learn (auto-learn) the port security configurations over a specified period. This feature allows the 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 beacuse it saves tedious manual configuration for each port. You must configure auto-learning 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 occurs 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, auto-learning does not add a new entry to allow any other pWWN on that interface. All other pWWNs are 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.

When you activate the port security feature, the following operations occur:

- Auto-learning is also automatically enabled, which means the following:

- From this point, auto-learning occurs 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 the **no shutdown** command to bring that port back online.

Configuring Port Security

Configuring Port Security with Auto-Learning and CFS Distribution

You can configure port security using auto-learning and CFS distribution.

Procedure

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-learning 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 action ensures that the configured 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 153
Committing the Changes, on page 162
Copying the Port Security Database, on page 168
Disabling Auto-Learning, on page 157
Enabling Port Security, on page 153
Enabling Port Security Distribution, on page 161

Configuring Port Security with Auto-Learning without CFS

You can configure port security using auto-learning without Cisco Fabric Services (CFS).

Procedure

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-learning on each VSAN.
Step 5 Copy the active database to the configured 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 153
Copying the Port Security Database, on page 168
Disabling Auto-Learning, on page 157
Enabling Port Security, on page 153

Configuring Port Security with Manual Database Configuration

You can configure port security and manually configure the port security database.

Procedure

Step 1 Enable port security.
Step 2 Manually configure all port security entries into the configured database on each VSAN.
Step 3 Activate port security on each VSAN. This action turns on auto-learning by default.
Step 4 Disable auto-learning 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

You can enable port security.

By default, the port security feature is disabled.

Procedure

	Command or Action	Purpose

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

Step 2	**port-security enable** **Example:** `switch(config)# port-security enable`	Enables port security on that switch.
Step 3	**no port-security enable** **Example:** `switch(config)# no port-security enable`	Disables (default) port security on that switch.

Port Security Activation

Activating Port Security

You can activate port security.

Procedure

	Command or Action	Purpose

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

	Command or Action	Purpose
Step 2	**port-security activate vsan** *vsan-id* **Example:** `switch(config)# port-security activate vsan 20`	Activates the port security database for the specified VSAN, and automatically enables auto-learning.
Step 3	**port-security activate vsan** *vsan-id* **no-auto-learn** **Example:** `switch(config)# port-security activate vsan 20 no-auto-learn`	Activates the port security database for the specified VSAN, and disables auto-learning.
Step 4	**no port-security activate vsan** *vsan-id* **Example:** `switch(config)# no port-security activate vsan 20`	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

You can forcefully activate the port security database.

Procedure

	Command or Action	Purpose

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

	Command or Action	Purpose
Step 2	**port-security activate vsan** *vsan-id* **force** **Example:** `switch(config)# port-security activate vsan 210 force`	Forces the port security database to activate for the specified VSAN even if conflicts occur.

Database Reactivation

You can reactivate the port security database.

Procedure

	Command or Action	Purpose

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

Step 2	**no no port-security auto-learn vsan** *vsan-id* **Example:** `switch(config)# no no port-security auto-learn vsan 35`	Disables auto-learning and stops the switch from learning about new devices that access the switch. This command also enforces the database contents based on the devices learned up to this point.
Step 3	**exit** **Example:** `switch(config)# exit`	Exits the configuration mode.
Step 4	**port-security database copy vsan** *vsan-id* **Example:** `switch# port-security database copy vsan 35`	Copies from the active to the configured database.

Step 5	**configure terminal**
	
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Reenters configuration mode.

Step 6	**port-security activate vsan** *vsan-id* **Example:** `switch(config)# port-security activate vsan 35`	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

You can enable 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.

Procedure

	Command or Action	Purpose

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

Step 2	**port-security auto-learn vsan** *vsan-id* **Example:** `switch(config)# port-security auto-learn vsan 1`	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

You can disable auto-learning.

Procedure

	Command or Action	Purpose

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

Step 2	**no port-security auto-learn vsan** *vsan-id* **Example:** `switch(config)# no port-security auto-learn vsan 23`	Disables auto-learning and stops the switch from learning about new devices that access the switch. This command 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 20: 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 if auto-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 vfc21 (F1).
- A pWWN (P2) is allowed access through interface vfc22 (F1).
- A nWWN (N1) is allowed access through interface vfc22 (F2).
- Any WWN is allowed access through interface vfc31 (F3).
- A nWWN (N3) is allowed access through any interface.
- A pWWN (P3) is allowed access through interface vfc24 (F4).
- A sWWN (S1) is allowed access through interface vfc31-33 (F10 to F13).
- A pWWN (P10) is allowed access through interface vfc41 (F11).

The following table summarizes the port security authorization results for this active database.

Table 21: Authorization Results for Scenario

Device Connection Request	Authorization	Condition	Reason
P1, N2, F1	Permitted	1	No conflict.
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.

Device Connection Request	Authorization	Condition	Reason
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 157

Port Security Manual Configuration

You can manually configure port security.

Procedure

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

The WWN Identification has the following configuration guidelines and limitations:

- 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 a SAN switch port F, 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, 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.
- You must configure all port channel xE ports 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, you can modify the configuration database 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, you can 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.

Procedure

	Command or Action	Purpose

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

Step 2	**port-security database vsan** *vsan-id* **Example:** `switch(config)# port-security database vsan 25`	Enters the port security database mode for the specified VSAN.
Step 3	**no port-security database vsan** *vsan-id* **Example:** `switch(config)# no port-security database vsan 25`	Deletes the port security configuration database from the specified VSAN.
Step 4	switch(config-port-security)# **swwn** *swwn-id* **interface san-port-channel 5** **Example:** `switch(config-port-security)# swwn 21:00:05:30:23:1a:11:03 interface san-port-channel 5`	Configures the specified sWWN to only log in through SAN port channel 5.

	Command or Action	Purpose
Step 5	**any-wwn interface vfc** *vfc-id* **- vfc** *vfc-id* **Example:** `switch(config-port-security)# any-wwn interface vfc 32 - vfc 35`	Configures any WWN to log in through the specified interfaces.

EXAMPLES

This example shows how to enter the port security database mode for VSAN 2:

```
switch(config)# port-security database vsan 2
```

This example shows how to configure 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 shows how to configure 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 vfc 32
```

This example shows how to configure any WWN to log in through the specified interface in any switch:

```
switch(config-port-security)# any-wwn interface vfc 32
```

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 System Management Configuration Guide for your device.

Enabling Port Security Distribution

You can enable port security distribution.

Procedure

	Command or Action	Purpose

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

Step 2	**port-security distribute** **Example:** `switch(config)# port-security distribute`	Enables distribution.

	Command or Action	Purpose
Step 3	**no port-security distribute** **Example:** `switch(config)# no port-security distribute`	Disables distribution.

Related Topics

Activation and Auto-Learning Configuration Distribution, on page 163

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

You can commit the port security configuration changes for the specified VSAN.

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.

Procedure

	Command or Action	Purpose

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

Step 2	**port-security commit vsan** *vsan-id* **Example:** `switch(config)# port-security commit vsan 100`	Commits the port security changes in the specified VSAN.

Discarding the Changes

You can discard the port security configuration changes for the specified VSAN.

If you discard (abort) the changes made to the pending database, the configuration remains unaffected and the lock is released.

Procedure

	Command or Action	Purpose

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

Step 2	**port-security abort vsan** *vsan-id* **Example:** `switch(config)# port-security abort vsan 35`	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 22: 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 and 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 and 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.

Merging the Port Security Database

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 forcefully synchronizes the databases and the activation states in the fabric.

For additional information, refer to CFS Merge Support in the Series System Management Configuration Guide for your device.

Database Interaction

The following table lists the differences and interaction between the active and configuration databases.

Table 23: 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 lists the differences between the active database and the configuration database.

The following figure shows various scenarios of the active database and the configuration database status based on port security configurations.

Figure 30: Port Security Database Scenarios

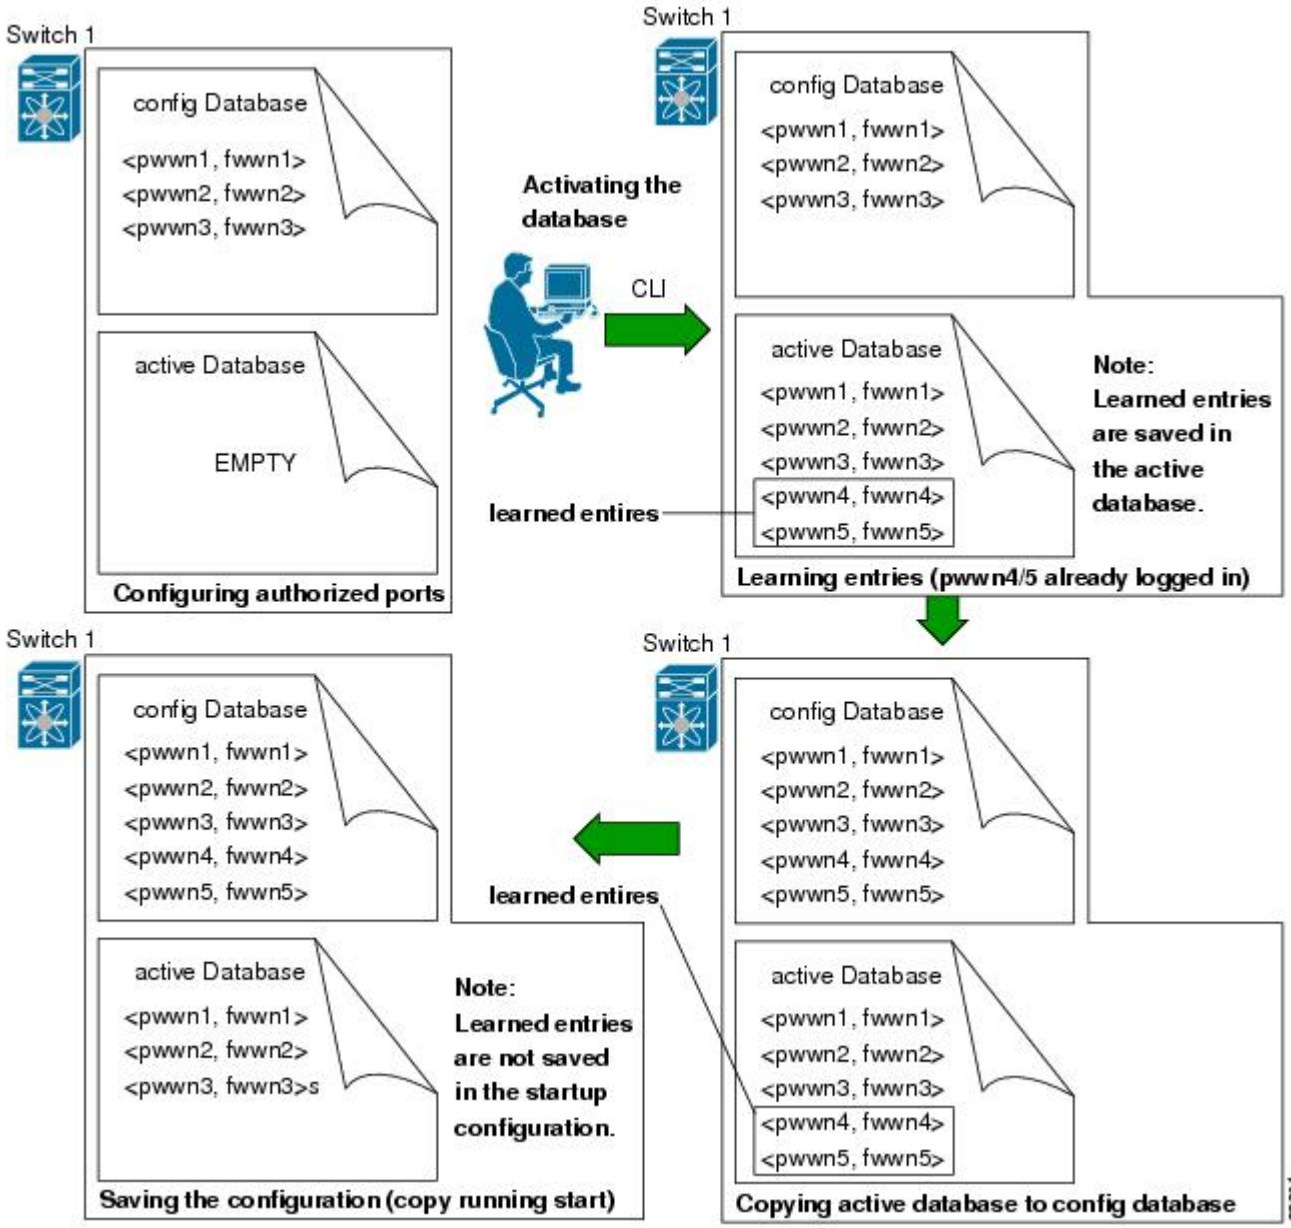

Database Scenarios

the following figure illustrates various scenarios showing the active database and the configuration database status based on port security configurations.

Figure 31: Port Security Database Scenarios

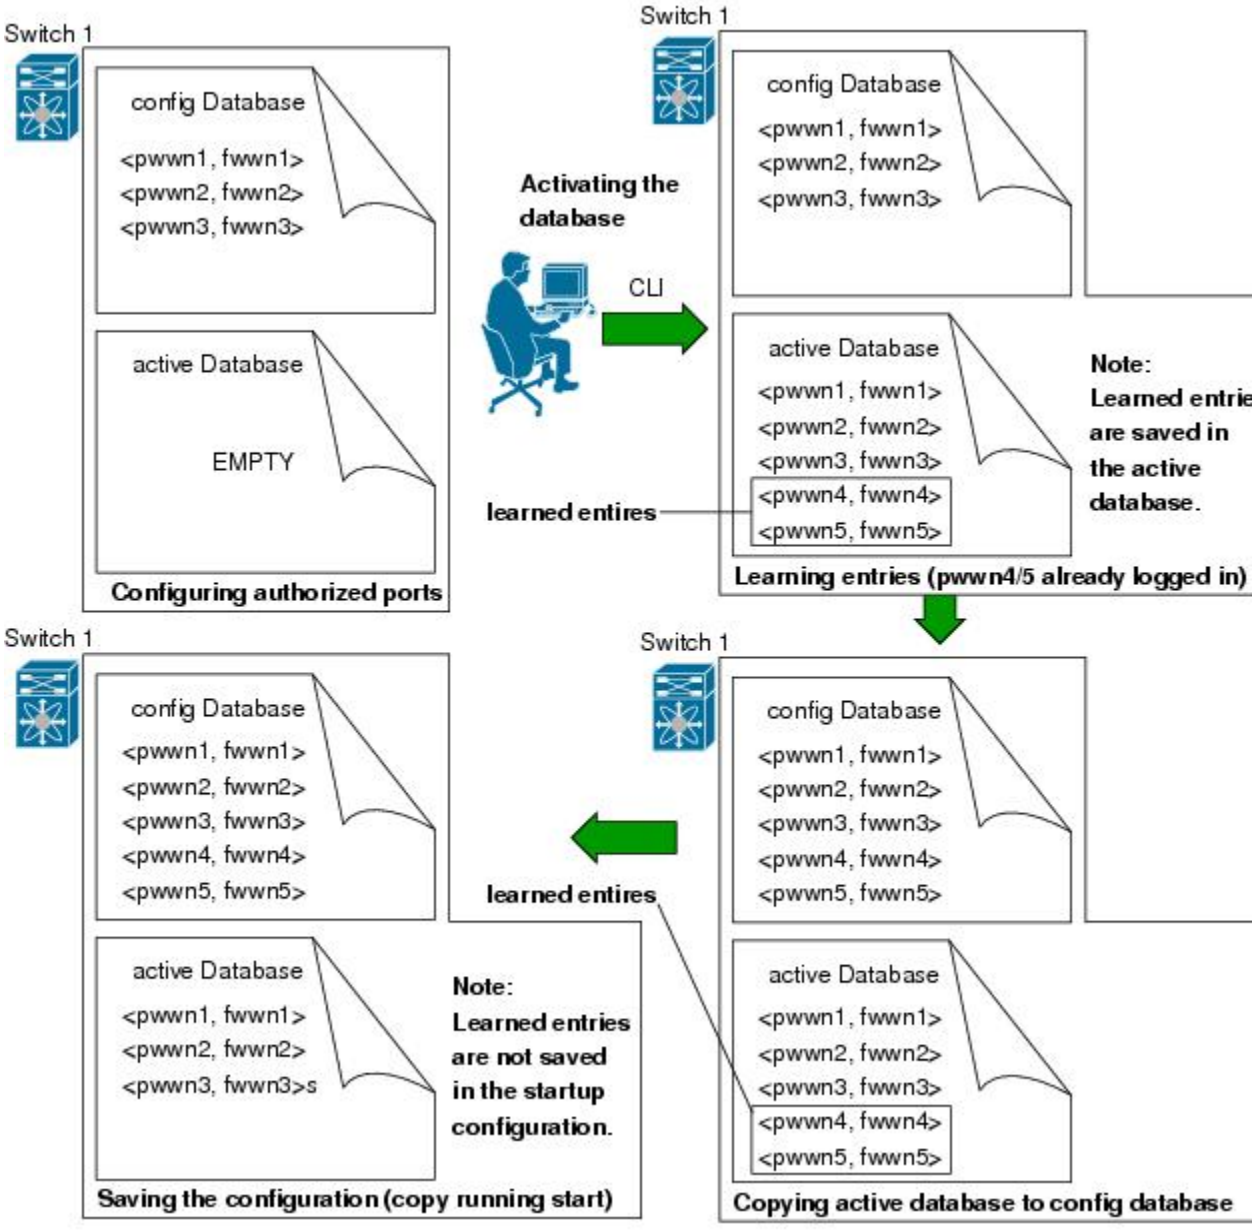

Copying the Port Security Database

Tip We recommend that you copy the active database to the config database after disabling auto-learning. This action ensures that the configuration database is in synchronization with the active database. If distribution is enabled, this command creates a temporary copy (and a fabric lock) of the configuration database. If you lock the fabric, you must 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. You must enter the **port-security commit** command 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 vfc21 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 Settings for Port Security

The following table lists the default settings for all port security features in any switch.

Table 24: 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.

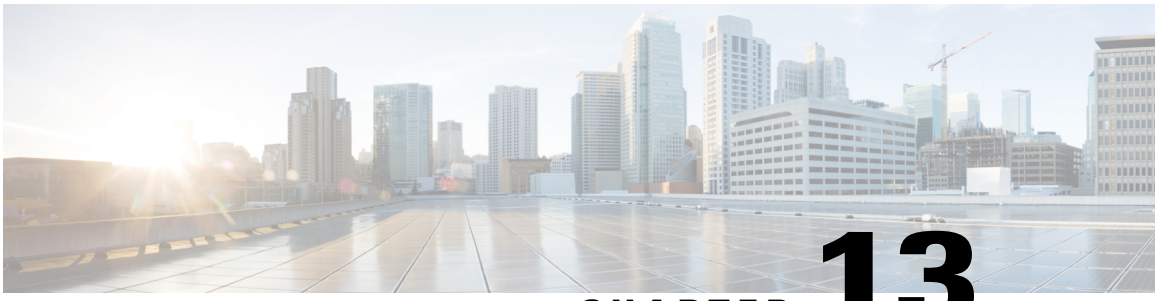

CHAPTER 13

Configuring Fabric Binding

This chapter describes how to configure fabric binding.

This chapter includes the following sections:

Configuring Fabric Binding

Information About Fabric Binding

Fabric binding ensures that Inter-Switch Links (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 25: 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 per VSAN.	Requires activation per VSAN.
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 Cisco Fabric Services (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 the port VSAN
  - TE port security binding check on each allowed VSAN

While port security complements fabric binding, they are independent features and that you can enable or disable separately.

Fabric Binding Enforcement

You must enable fabric binding in each switch in the fabric that participates in the fabric binding. By default, this feature is disabled. 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 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 per VSAN.

Configuring Fabric Binding

You can configure fabric binding in each switch in the fabric.

Procedure

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

You can enable fabric binding on any participating switch.

Procedure

	Command or Action	Purpose

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

Step 2	**feature fabric-binding** **Example:** `switch(config)# feature fabric-binding`	Enables fabric binding on that switch.

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:

Procedure

	Command or Action	Purpose

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

Step 2	**fabric-binding database vsan** *vsan-id* **Example:** `switch(config)# fabric-binding database vsan 35`	Enters the fabric binding submode for the specified VSAN.
Step 3	**no fabric-binding database vsan** *vsan-id* **Example:** `switch(config)# no fabric-binding database vsan 35`	Deletes the fabric binding database for the specified VSAN.
Step 4	**swwn** *swwn-id* **domain** *domain-id* **Example:** `switch(config-fabric-binding)# swwn 21:00:05:30:23:1a:11:03 domain 25`	Adds the sWWN of another switch for a specific domain ID to the configured database list.
Step 5	**no swwn** *swwn-id* **domain** *domain-id* **Example:** `switch(config-fabric-binding)# no swwn 21:00:05:30:23:1a:11:03 domain 25`	Deletes the sWWN and domain ID of a switch from the configured database list.

Fabric Binding Activation and Deactivation

Fabric binding maintains a configuration database (config database) and an active database. The config database is a read-write database that collects the configurations that 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 might 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

You can activate the fabric binding feature.

Procedure

	Command or Action	Purpose

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

Step 2	**fabric-binding activate vsan** *vsan-id* **Example:** `switch(config)# fabric-binding activate vsan 25`	Activates the fabric binding database for the specified VSAN.
Step 3	**no fabric-binding activate vsan** *vsan-id* **Example:** `switch(config)# no fabric-binding activate vsan 25`	Deactivates the fabric binding database for the specified VSAN.

Forcing Fabric Binding Activation

You can forcefully activate the fabric binding database.

If the database activation is rejected due to one or more conflicts listed in the previous section, you might decide to proceed with the activation by using the force option.

Procedure

	Command or Action	Purpose

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

	Command or Action	Purpose
Step 2	**fabric-binding activate vsan** *vsan-id* **force** **Example:** `switch(config)# fabric-binding activate vsan 12 force`	Activates the fabric binding database for the specified VSAN forcefully, even if the configuration is not acceptable.
Step 3	**no fabric-binding activate vsan** *vsan-id* **force** **Example:** `switch(config)# no fabric-binding activate vsan 12 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 the Fabric Binding Configuration

To display fabric binding information, perform one of the following tasks:

Command	
show fabric-binding database [**active**]	Displays the configured fabric binding database. You can add the **active** keyword to display only the active fabric binding database.
show fabric-binding database [**active**] [**vsan** *vsan-id*]	Displays the configured fabric binding database for the specified VSAN.
show fabric-binding statistics	Displays statistics for the fabric binding database.
show fabric-binding status	Displays fabric binding status for all VSANs.
show fabric-binding violations	Displays fabric binding violations.
show fabric-binding efmd [**vsan** *vsan-id*]	Displays the configured fabric binding database for the specified VSAN.

This example shows how to display the active fabric binding information for VSAN 4:

```
switch# show fabric-binding database active vsan 4
```

This example shows how to display 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.

This example shows how to display EFMD Statistics for VSAN 4:

```
switch# show fabric-binding efmd statistics vsan 4
```

Default Settings for Fabric Binding

The following table lists the default settings for the fabric binding feature.

Table 26: Default Fabric Binding Settings

Parameters	Default
Fabric binding	Disabled

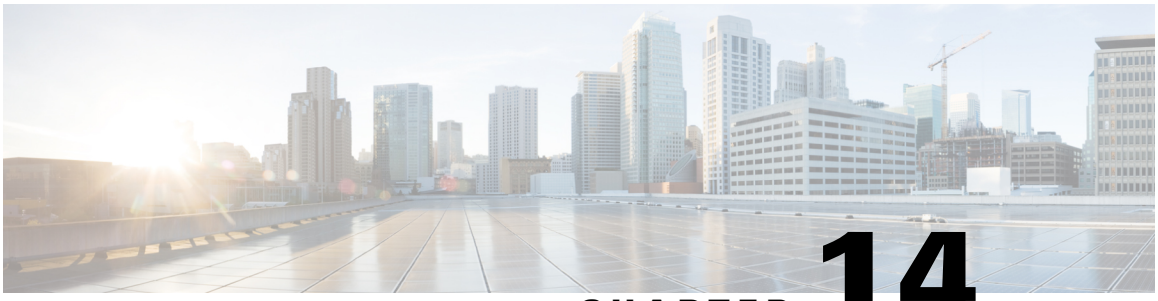

CHAPTER 14

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 device 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 32: FCSs in a VSAN Environment

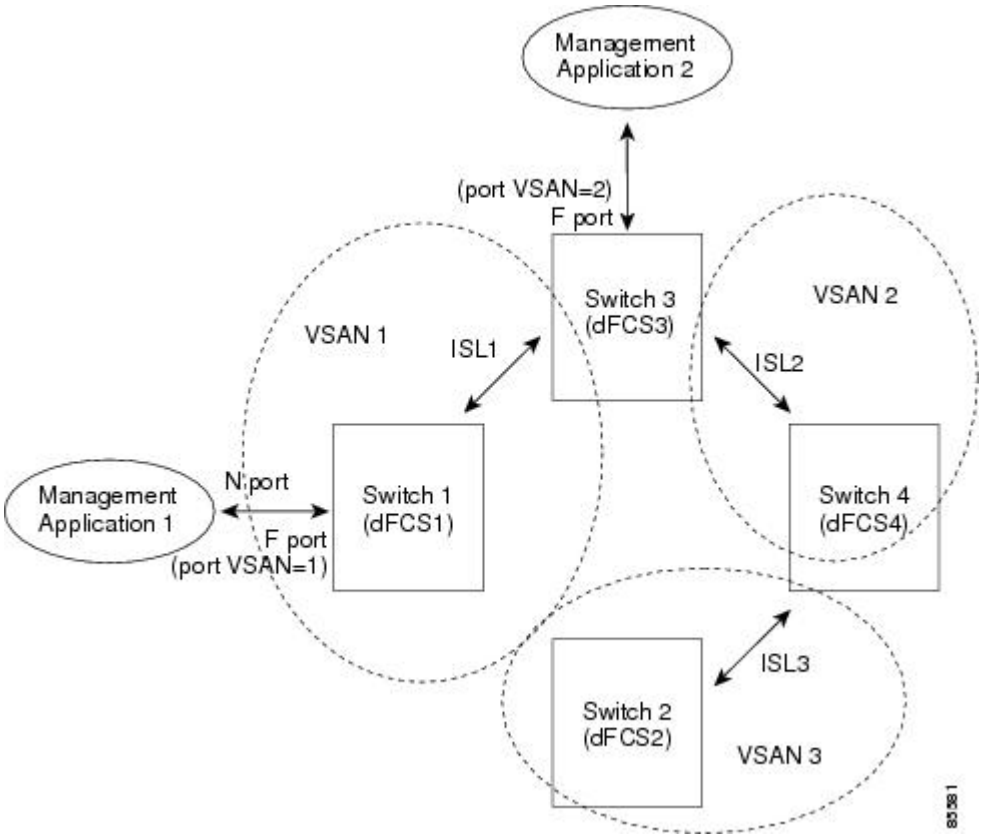

FCS Characteristics

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 devices.

To enable global checking of the platform name, perform this task:

To register platform attributes, perform this task:

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global 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 27: Default FCS Settings

Parameters	Default
Global checking of the platform name	Disabled
Platform node type	Unknown

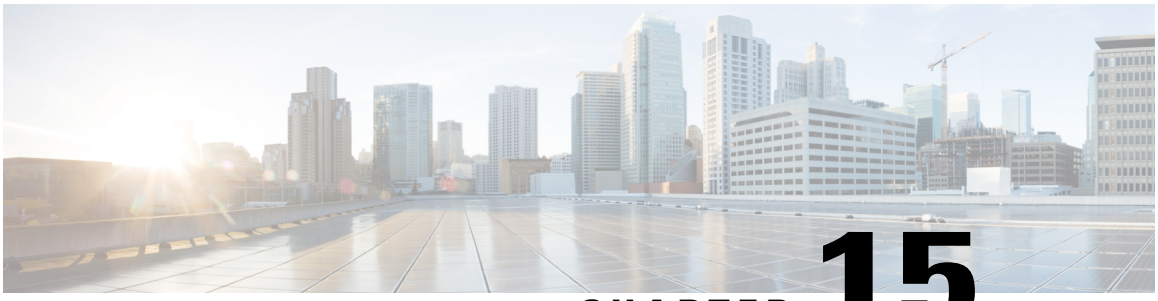

CHAPTER 15

Configuring Port Tracking

This chapter describes how to configure port tracking.

This chapter includes the following sections:

Configuring Port Tracking

Cisco SAN 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

Port tracking allows you to use information about the operational state of the link so that you can initiate a failure in the link that connects the edge device. Converting the indirect failure to a direct failure triggers a faster recovery process towards redundant links. When enabled, port tracking brings down the configured links based on the failed link and forces the traffic to be redirected to another redundant link.

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 33: Traffic Recovery Using Port Tracking

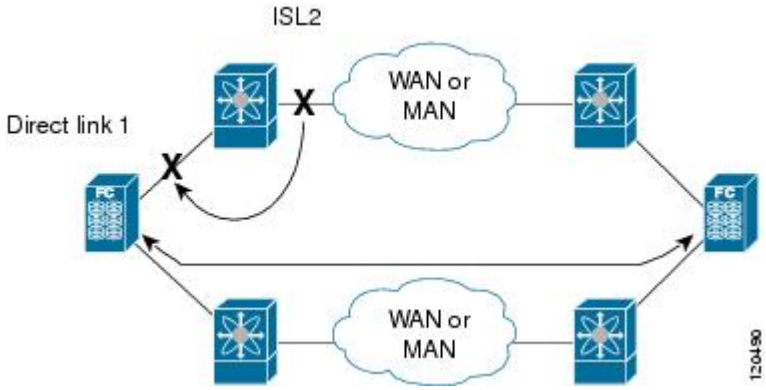

Port tracking monitors and detects failures that cause topology changes and brings down the links that connect the attached devices. When you enable this feature and explicitly configure the linked and tracked ports, the switch software monitors the tracked ports and alters the operational state of the 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.

Related Topics

About RSCN Information, on page 123

Fibre Channel Timeout Values

Default Settings for Port Tracking

The following table lists the default settings for port tracking parameters.

Table 28: Default Port Tracking Parameters

Parameters	Default
Port tracking	Disabled
Operational binding	Enabled along with port tracking

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 vfc22 to Port vfc24 and back to Port vfc22) to avoid recursive dependency.

Enabling Port Tracking

The port tracking feature is disabled by default. 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.

You can enable port tracking

Procedure

	Command or Action	Purpose

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

Step 2	**port-track enable** **Example:** `switch(config)# port-track enable`	Enables port tracking.
Step 3	**no port-track enable** **Example:** `switch(config)# no port-track enable`	Removes the currently applied port tracking configuration and disables port tracking.

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:

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface vfc** *vfc-id*	Enters the interface configuration mode for the linked port. You can now configure the tracked ports.
Step 3	switch(config-if)# **port-track interface vfc** *vfc-id* | **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 vfc** *vfc-id* | **san-port-channel** *port*	Removes the port tracking configuration that is currently applied to the interface.

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 34: Traffic Recovery Using Port Tracking

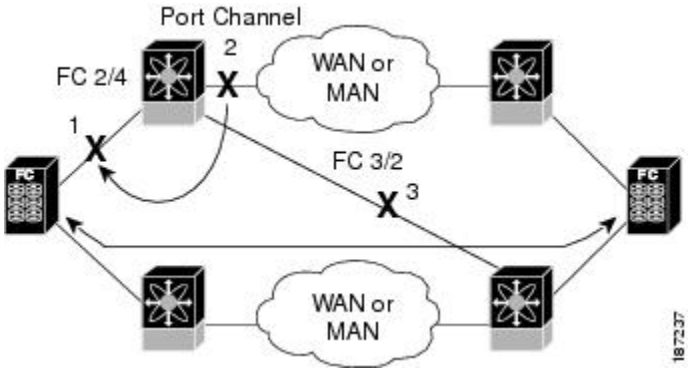

Tracking Multiple Ports

To track multiple ports, perform this task:

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **interface vfc** *vfc-id*	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 vfc** *vfc-id* | **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.

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

You can monitor a tracked port in a specific VSAN.

Procedure

	Command or Action	Purpose

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

Step 2	switch(config)# **interface vfc** *vfc-id*	Configures the specified interface and enters the interface configuration mode. You can now configure tracked ports.
Step 3	**port-track interface san-port-channel 1 vsan 2** **Example:** `switch(config-if)# port-track interface san-port-channel 1 vsan 2`	Enables tracking of the SAN port channel in VSAN 2.
Step 4	**no port-track interface san-port-channel 1 vsan 2** **Example:** `switch(config-if)# 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.

Forcefully Shutting down

If a tracked port flaps frequently, tracking ports using the operational binding feature may cause frequent topology changes. You might 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

You can forcefully shut down a tracked port.

Procedure

	Command or Action	Purpose

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

Step 2	switch(config)# **interface vfc** *vfc-id*	Configures the specified interface and enters the interface configuration mode. You can now configure tracked ports.
Step 3	**port-track force-shut** **Example:** `switch(config-if)# port-track force-shut`	Forcefully shuts down the tracked port.
Step 4	**no port-track force-shut** **Example:** `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 vfc21
vfc21 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 vc22 (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 vfc21
...
```

The following example shows how to display the port track mode:

```
switch# show interface vfc 24
vfc24 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
```

INDEX

E

F

H

I

L

M

N

P

R

S

T

U

V

Z