Inspur

CN12700 Series

SAN Switching Config Guide

(Release 8.x)

Inspur-Cisco Networking Technology Co.,Ltd. provides customers with comprehensive technical support and services.
For any assistance, please contact our local office or company headquarters.
Website: http://www.inspur.com/
Technical Support Tel: 400-691-1766
Technical Support Email:inspur_network@inspur.com
Technical Document Support Email:inspur_network@inspur.com
Address: 1036 Langchao Road, Lixia District, Jinan City, Shandong Province
Postal code: 250101

Notice

Preface

Objectives

This guide describes main functions of the CN12700 Series. To have a quick grasp of the CN12700 Series, please read this manual carefully.

Versions

The following table lists the product versions related to this document.

Product name	Version
CN12700 Series	

Conventions

Symbol conventions

The symbols that may be found in this document are defined as follows.

Symbol	Description
⚡ Warning	Indicates a hazard with a medium or low level of risk which, if not avoided, could result in minor or moderate injury.
⚠ Caution	Indicates a potentially hazardous situation that, if not avoided, could cause equipment damage, data loss, and performance degradation, or unexpected results.
✎ Note	Provides additional information to emphasize or supplement important points of the main text.
🔍 Tip	Indicates a tip that may help you solve a problem or save time.

General conventions

Convention	Description
Boldface	Names of files, directories, folders, and users are in **boldface**. For example, log in as user **root**.
Italic	Book titles are in *italics*.
`Lucida Console`	Terminal display is in `Lucida Console`.

Command conventions

Convention	Description
Boldface	The keywords of a command line are in **boldface**.
Italic	Command arguments are in *italics*.
[]	Items (keywords or arguments) in square brackets [ ] are optional.
{ x	y	... }	Alternative items are grouped in braces and separated by vertical bars. One is selected.
[x	y	...]	Optional alternative items are grouped in square brackets and separated by vertical bars. One or none is selected.
{ x	y	... } *	Alternative items are grouped in braces and separated by vertical bars. A minimum of one or a maximum of all can be selected.
[x	y	...] *	The parameter before the & sign can be repeated 1 to n times.

GUI conventions

Convention	Description
Boldface	Buttons, menus, parameters, tabs, windows, and dialog titles are in **boldface**. For example, click **OK**.
>	Multi-level menus are in boldface and separated by the ">" signs. For example, choose **File** > **Create** > **Folder**.

Keyboard operation

Format	Description
Key	Press the key. For example, press **Enter** and press **Tab**.

Format	Description
Key 1+Key 2	Press the keys concurrently. For example, pressing **Ctrl+C** means the two keys should be pressed concurrently.
Key 1, Key 2	Press the keys in turn. For example, pressing **Alt**, **A** means the two keys should be pressed in turn.

Mouse operation

Action	Description
Click	Select and release the primary mouse button without moving the pointer.
Double-click	Press the primary mouse button twice continuously and quickly without moving the pointer.
Drag	Press and hold the primary mouse button and move the pointer to a certain position.

Change history

Updates between document versions are cumulative. Therefore, the latest document version contains all updates made to previous versions.

Issue 01 (2020-02-24)

Initial commercial release

Contents

Figures

Tables

Chapter 1 Overview

This chapter contains the following sections:
・SAN Switching Overview

1.1 SAN Switching Overview

This chapter provides an overview of SAN switching for Inspur INOS 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.

VSAN Trunking

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

Virtual SANs

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

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 Inspur 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.

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

2.1 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.

2.1.1 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

**Inspur-Cisco Networking
Technology Co.,Ltd.**
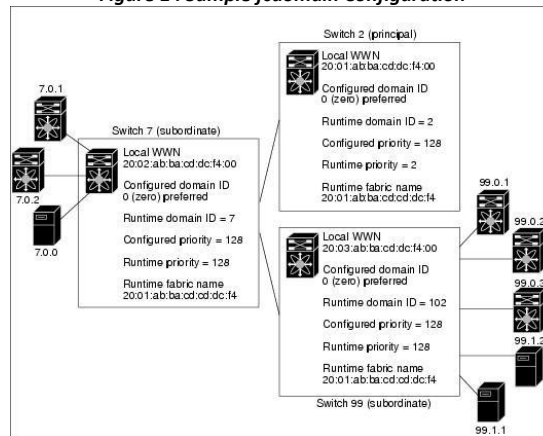
Domain Restarts

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).

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.

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

DETAILED STEPS

	Command or Action	Purpose

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

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

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.

**Inspur-Cisco Networking
Technology Co.,Ltd.**

Enabling Domain Manager Fast Restart

You can enable the domain manager fast restart feature.

SUMMARY STEPS

1. **configure terminal**
2. **fcdomain optimize fast-restart vsan** *vsan-id*
3. **no fcdomain optimize fast-restart vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**fcdomain 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.

SUMMARY STEPS

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

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**fcdomain priority** *number* **vsan** *vsan-id* **Example:** `switch(config)# fcdomain priority 12 vsan`	Configures the specified priority for the local switch in the specified VSAN. The fcdomain priority ranges from 1 to 254. The VSAN ID ranges from 1 to 4093.

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

Configuring Fabric Names

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

SUMMARY STEPS

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

DETAILED STEPS

	Command or Action	**Purpose**

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

Step 2	**fcdomain fabric-name 20:1:ac:16:5e:0:21:01 vsan** *vsan-id* **Example:** `switch(config)# fcdomain fabric-`	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.

Rejecting Incoming RCFs

You can reject incoming RCF request frames.

SUMMARY STEPS

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

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**interface vfc** *vfc-id* **Example:** `switch(config)# interface vfc 20`	Configures the specified interface. The virtual interface ID ranges from 1 to 8192.
Step 3	switch(config)# **interface fc** *slot/port*	Configures the specified interface.
Step 4	**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 5	**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

Inspur-Cisco Networking
Technology Co.,Ltd.

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

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

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**fcdomain auto-reconfigure vsan** *vsan-id*
	
	Example:
	`switch(config)# fcdomain auto-reconfigure`
	`vsan 1`
		Enables the automatic reconfiguration option in the specified VSAN. The VSAN ID ranges from 1 to 4093.

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

2.1.2 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.

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
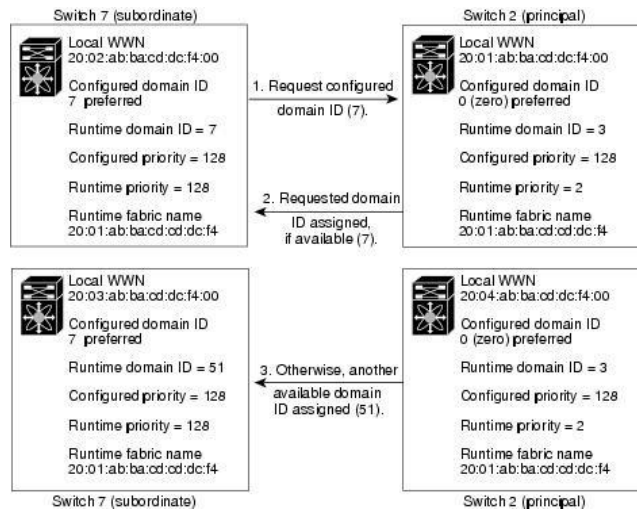
The operation of a subordinate switch changes based on three factors:
· The allowed domain ID lists
· The configured domain ID
· The domain ID that the principal switch has assigned to the requesting switch In specific situations, the changes are as follows:
· When the received domain ID is not within 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.

Related Topics
Allowed Domain ID Lists

Configuring Static or Preferred Domain IDs
You can specify a static or preferred domain ID.

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

Inspur-Cisco Networking
Technology Co.,Ltd.

5.   **no fcdomain domain** *domain-id* **preferred vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	**configure terminal** **Example:** `switch# configure termina l`	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`	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.

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

Allowed Domain ID Lists

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

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

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

Configuring Allowed Domain ID Lists

You can configure the allowed domain ID list.

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

Inspur-Cisco Networking
Technology Co.,Ltd.

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**fcdomain allowed** *domain-id range* **vsan** *vsan-id* **Example:** `switch(config)# fcdomain allowed 3 vsan 10`	Configures the list to allow switches with the domain ID range in the specified VSAN. The domain ID range is from 1 to 239. The VSAN ID range is from 1 to 4093.
Step 3	**no fcdomain allowed** *domain-id range* **vsan** *vsan-id*	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 Inspur SAN switches in the fabric using the Inspur 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.

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.

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

DETAILED STEPS

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

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

Locking the Fabric

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

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

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

Committing Changes

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

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

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

DETAILED STEPS

	Command or Action	**Purpose**

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

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

Discarding Changes

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

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

SUMMARY STEPS

---

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

DETAILED STEPS

	Command or Action	Purpose

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

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

Clearing a Fabric Lock

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

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

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

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

Displaying CFS Distribution Status

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

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

Displaying Pending Changes

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

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

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

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

```
9,24,100,231-239. [User] configured allowed
domain IDs: 10-230.
```

Displaying Session Status

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

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

Contiguous Domain ID Assignments

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

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

• If the contiguous domain assignment is disabled in the principal switch, the principal switch assigns the available domains to the subordinate switch.

Enabling Contiguous Domain ID Assignments

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

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

DETAILED STEPS

	Command or Action	Purpose

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

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.

2.1.3 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:

Inspur-Cisco Networking
Technology Co.,Ltd.

• 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.

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

Enabling the Persistent FC ID Feature

You can enable the persistent FC ID feature.

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

DETAILED STEPS

	Command or Action	Purpose

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

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

Persistent FC ID Configuration Guidelines

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

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

• Ensure that the persistent FC ID feature is enabled in the required VSAN.

• Ensure that the required VSAN is an active VSAN. Persistent FC IDs can only be configured on active VSANs.

• Verify that the domain part of the FC ID is the same as the runtime domain ID in the required VSAN. If the software detects a domain mismatch, the command is rejected.

• Verify that the port field of the FC ID is 0 (zero) when configuring an area.

Configuring Persistent FC IDs

You can configure persistent FC IDs.

SUMMARY STEPS

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

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**fcdomain fcid database** **Example:** `switch(config)# fcdomain fcid database`	Enters FC ID database configuration submode.

Step 3	**vsan** *vsan-id* **wwn 33:e8:00:05:30:00:16:df fcid** *fcid*
	
	Example:
	`switch(config-fcid-db)# vsan 26`
	`wwn 33:e8:00:05:30:00:16:df fcid`
		Configures a device WWN (33:e8:00:05:30:00:16:df) with the FC ID 0x070128 in the specified VSAN.
		
		Note To avoid assigning a duplicate FC ID, use the **show fcdomain address-allocation vsan** command to display the FC IDs in use.

Step 4	**vsan** *vsan-id* **wwn 11:22:11:22:33:44:33:44 fcid** *fcid* **dynamic** **Example:** `switch(config-fcid-db)# vsan 13 wwn`	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`	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

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.

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

Configuring Unique Area FC IDs for an HBA

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

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

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

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

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

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

switch(config-if)# shutdown

switch(config-if)# end
```

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

```
switch# show fcdomain vsan 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
----------------------------------------------------------------
```

Inspur-Cisco Networking
Technology Co.,Ltd.

```
        vfc20 3        0x6fee00        50:05:08:b2:00:71:c8:c2        50:05:08:b2:00:71:c8:c0
```

Persistent FC ID Selective Purging

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

Table 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.

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

DETAILED STEPS

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

2.1.4 Verifying the fcdomain Configuration

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
```

2.1.5 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

**Inspur-Cisco Networking
Technology Co.,Ltd.**

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 Identifier Virtualization

This chapter describes how to configure N Port Identifier Virtualization (NPIV). This chapter includes the following sections:
- Information About N Port Identifier Virtualization
- Enabling N Port Identifier Virtualization

3.1 Information About N Port Identifier Virtualization

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

Figure 3 : NPIV Example
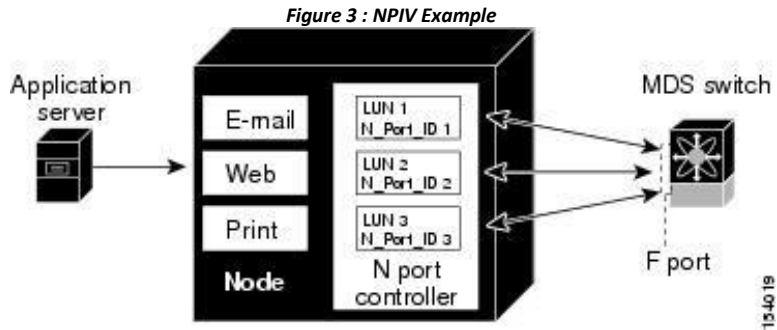
3.2 Enabling N Port Identifier Virtualization

You can enable or disable NPIV on the switch.

Before You Begin

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

SUMMARY STEPS

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

DETAILED STEPS

	Command or Action	Purpose

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

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

CHAPTER 4 Configuring and Managing VSANs

This chapter describes how to configure and manage VSANs. This chapter includes the following sections:
   • Configuring and Managing VSANs

4.1 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.

4.1.1 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.

Inspur-Cisco Networking
Technology Co.,Ltd.

Figure 4 : Logical VSAN Segmentation
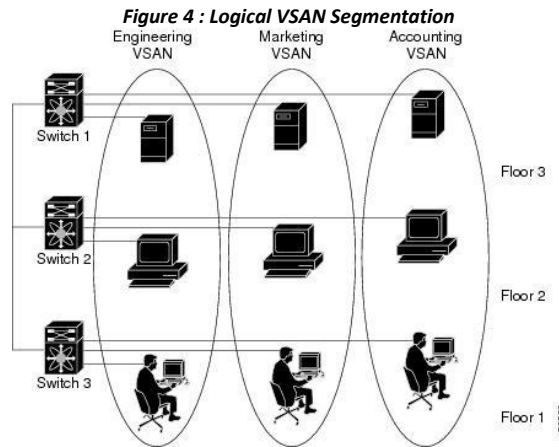
The application servers or storage arrays can be connected to the switch using Fibre Channel or virtual Fibre Channel interfaces. A VSAN can include a mixture of Fibre Channel and virtual Fibre Channel interfaces.

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 5 : 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

**Inspur-Cisco Networking
Technology Co.,Ltd.**

• 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 3 : VSAN and Zone Comparison

VSAN Characteristic	Zone Characteristic
VSANs equal SANs with routing, naming, and zoning protocols.	Routing, naming, and zoning protocols are not available on a per-zone basis.
VSANs limit unicast, multicast, and broadcast traffic.	Zones limit unicast traffic.
Membership is typically defined using the VSAN ID to F ports.	Membership is typically defined by the pWWN.
An HBA or a storage device can belong only to a single VSAN (the VSAN associated with the F port).	An HBA or storage device can belong to multiple zones.
VSANs enforce membership at each E port, source port, and destination port.	Zones enforce membership only at the source and destination ports.
VSANs are defined for larger environments (storage service providers).	Zones are defined for a set of initiators and targets not visible outside the zone.
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

Inspur-Cisco Networking
Technology Co.,Ltd.

defined in VSAN 2 is different and separate from zone A defined in VSAN 7.

Figure 6 : VSANS with Zoning

4.1.2 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.

• 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.

SUMMARY STEPS

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

DETAILED STEPS

	Command or Action	Purpose

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

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).Inspur CN 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
Configuring VSAN Trunking

Assigning Static Port VSAN Membership

You can statically assign VSAN membership for an interface port.

SUMMARY STEPS

1. **configure terminal**
2. **vsan database**
3. **vsan** *vsan-id*
4. **vsan** *vsan-id* **interface vfc** *vfc-id*
5. **vsan** *vsan-id* **vfc** *vfc-id*}

DETAILED STEPS

	Command or Action	Purpose
Step 1	**configure terminal** **Example:** `switch# configure terminal` `switch(config)#`	Enters global configuration mode.
Step 2	**vsan database** **Example:** `switch(config)# vsan database` `switch(config-vsan-db)#`	Configures the database for a VSAN.
Step 3	**vsan** *vsan-id* **Example:** `switch(config-vsan-db)# vsan 50`	Creates a VSAN with the specified ID if that VSAN does not exist already.

| Step 4 | **vsan** *vsan-id* **interface vfc** *vfc-id*

Example:
`switch(config-vsan-db)# vsan 34 interface vfc 5` | Assigns the membership of the specified interface to the VSAN. |
| Step 5 | **vsan** *vsan-id* **vfc** *vfc-id*}

Example:
`switch(config-vsan-db)# vsan 10 vfc 3` | Updates the membership information of the interface to reflect the changed VSAN.

Note   To remove the VSAN membership of a vFC interface, assign the VSAN membership of that interface to another VSAN. Inspur recommends that you assign it to VSAN 1. |

Default VSANs

The factory settings for Inspur 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.

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).

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 7 : 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.

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

Deleting Static VSANs

You can delete a VSAN and its various attributes.

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

DETAILED STEPS

	Command or Action	Purpose

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

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

Step 3	**vsan** *vsan-id* **Example:** `switch(config-vsan-db)# vsan 2`	Places you in VSAN configuration mode.

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

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.

SUMMARY STEPS
1. **configure terminal**
2. **vsan database**
3. **vsan** _vsan-id_
4. **vsan** _vsan-id_ **loadbalancing src-dst-id**
5. **no vsan** _vsan-id_ **loadbalancing src-dst-id**
6. **vsan** _vsan-id_ **loadbalancing src-dst-ox-id**
7. **vsan** _vsan-id_ **suspend**
8. **no vsan** _vsan-id_ **suspend**
9. **end**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**vsan database**
	
	Example:
	`switch(config)# vsan database`
	`switch(config-vsan-db)#`
		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

4.1.3 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
```

4.1.4 Default Settings for VSANs

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

Table 4 : 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 5 DPVM

· Information About DPVM
· Default Settings
· Licensing Requirements
· Guidelines and Limitations for DPVM
· Configuring DPVM
· Verifying the DPVM Configuration
· DPVM Example Configuration
· Feature History

5.1 Information About DPVM

You can use Dynamic Port VSAN Membership (DPVM) to dynamically assign VSAN membership to ports by assigning VSANs based on the device WWN. DPVM eliminates the need to reconfigure the port VSAN membership to maintain fabric topology when a host or storage device connection is moved between two Inspur SAN switches or two ports within a switch. It retains the configured VSAN regardless of where a device is connected or moved.

DPVM assignment is based on the port world wide name (pWWN) and node world wide name (nWWN). A DPVM database contains mapping information for each device pWWN/nWWN assignment and the corresponding VSAN. Inspur INOS checks the database during a device FLOGI and obtains the required VSAN details.

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

DPVM uses the Inspur Fabric Services (CFS) infrastructure to allow efficient database management and distribution.

5.1.1 DPVM Databases

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

DPVM uses the following three databases:

Configuration (config) database

Stores all configuration changes when CFS distribution is disabled. Changes to this database are reflected in the active DPVM database when you activate the DPVM config database.

Active database

Represents the DPVM configuration that is currently active in the fabric.

Pending database

Stores all configuration changes when CFS distribution is enabled. Changes to this database are reflected in the config or active DPVM database when you commit the DPVM pending database.

Related Topics

Activating the DPVM Config Database
Verifying the DPVM Configuration

5.1.2 DPVM Database Distribution

DPVM can use CFS to distribute the database to all switches in the fabric to allow devices to move anywhere and keep the same VSAN membership.

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

When you enable CFS distribution for DPVM, the DPVM configuration database is copied into the DPVM pending database. All changes to the DPVM configuration are now stored in the DPVM pending database and the feature is locked (that is, no other switch can make changes to the DPVM database until you commit the changes or discard the changes and free the CFS lock).

The DPVM pending database includes the following changes:

 • Adding, deleting, or modifying database entries.

 • Activating, deactivating, or deleting the configuration database.

 • Enabling or disabling autolearning.

CFS distributes these changes to all switches in a fabric when you commit the changes. You can also discard (abort) the changes at this point.

5.1.3 Database Merge

When you merge two independent fabrics into one fabric, DPVM attempts to merge the DPVM database (the configuration database and static (unlearned) entries in the active DPVM database). To ensure a successful database merge, follow these guidelines:

 • Verify that the activation status and the auto-learn status is the same for both fabrics.

 • Verify that the combined number of device entries in each database does not exceed 16000 entries.

Related Topics
Displaying DPVM Database Merge Results

5.2 Default Settings

Table 5 : Default DPVM Parameter Settings

Parameters	Default
DPVM feature	Disabled
DPVM CFS distribution	Enabled
Autolearning	Disabled

5.3 Licensing Requirements

Product	License
Inspur CN12700 Series	DPVM requires the Fibre Channel over Ethernet (FCoE) license for each F-series module. FCoE enabled in a storage virtual device context (VDC) does not require the Advanced Services License. For a complete explanation of the Inspur INOS licensing scheme and how to obtain and apply licenses, see the Inspur INOS Licensing

5.4 Guidelines and Limitations for DPVM

DPVM has the following guidelines and limitations:
  • You should enable DPVM CFS distribution for all switches in your fabric.
  • Connect the dynamic device to an F-port on the switch.
  • Verify that the static port VSAN of the F port is valid (not isolated, not suspended, and in existence).
  • Verify that the dynamic VSAN configured for the device in the DPVM database is valid (not isolated, not suspended, and in existence).
  • DPVM supports MAC-based device mapping for FCoE devices. DPVM does not support pWWN mapping for FCoE devices.

5.5 Configuring DPVM

5.5.1 Enabling the DPVM Feature

You must enable the DPVM feature before you can configure DPVM.

SUMMARY STEPS
1. **config t**
2. **feature dpvm**
3. (Optional) **show feature**
4. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**feature dpvm** **Example:** `switch(config)# feature dpvm`	Enables the DPVM feature.
Step 3	**show feature** **Example:** `switch(config)# show feature`	(Optional) Displays the enabled or disabled state for each feature.
Step 4	**copy running-config startup-config** **Example:** `switch(config)# copy running-config startup-config`	(Optional) Copies the running configuration to the startup configuration.

5.5.2 Adding Entries into the DPVM Database

You can manually add entries into the config and pending DPVM databases.

Before You Begin

• Ensure that you have enabled the DPVM feature.

• Ensure that you have configured device aliases in enhanced mode if you want to configure device aliases in the DPVM database.

SUMMARY STEPS

1.   **config t**
2.   **dpvm database**
3.   **pwwn** *pwwn* **vsan** *vsan-id*
4.   **nwwn** *nwwn* **vsan** *vsan-id*
5.   **device-alias** *alias* **vsan** *vsan-id*
6.   **exit**
7.   (Optional) **show dpvm pending-diff**
8.   **dpvm commit**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**dpvm database**
	
	Example:
	`switch(config)# dpvm database`
	`switch(config-dpvm-db)#`
		Creates the DPVM config database and enters database configuration mode.

Step 3	**pwwn** *pwwn* **vsan** *vsan-id*
	
	Example:
	`switch(config-dpvm-db)# pwwn`
	`12:33:56:78:90:12:34:56 vsan`
	`100`
		Maps the configured pWWN to the VSAN. The *pwwn* is in pWWN dotted notation. The *vsan-id* range is from 1 to 4093.

Step 4	**nwwn** *nwwn* **vsan** *vsan-id*
	
	Example:
	`switch(config-dpvm-db)# nwwn`
	`14:21:30:12:63:39:72:81 vsan`
	`101`
		Maps the configured nWWN to the VSAN. The *nwwn* is in nWWN dotted notation. The *vsan-id* range is from 1 to 4093.

Step 5	**device-alias** *alias* **vsan** *vsan-id*
	
	Example:
	`switch(config-dpvm-db)# device-alias`
	`device1 vsan 102`
		Maps the configured device alias to the VSAN. The *alias* is any case-sensitive alphanumeric string up to 64 characters. The *vsan-id* range is from 1 to 4093.

Step 6	**exit** **Example:** `switch(config-dpvm-db)# exit`	Exits DPVM database configuration mode.

Inspur-Cisco Networking
Technology Co.,Ltd.

Step 7	(Optional) **show dpvm pending-diff** **Example:** `switch(config)# show dpvm pending`	(Optional) Displays the differences between the pending database and the config database. You can optionally discard these changes using the **dpvm abort** command.
Step 8	**dpvm commit** **Example:** `switch(config)# dpvm commit`	Commits the DPVM pending database to the config database. This step is required to release the CFS lock on the DPVM configuration and to distribute this change across the fabric. You can optionally use the **dpvm abort** command to discard these changes and release the CFS lock.

What to Do Next

You should compare the DPVM config database to the active database and activate these changes.

5.5.3 Activating the DPVM Config Database

You can activate the DPVM config database to make it the active database. Activation might fail if conflicting entries are found between the DPVM config database and the currently active DPVM database. However, you can force activation to override conflicting entries.

To disable DPVM, you must explicitly deactivate the currently active DPVM database by entering the **no dpvm activate** command.

Before You Begin

• Ensure that you have enabled the DPVM feature.

SUMMARY STEPS

1. **config t**
2. (Optional) **dpvm database diff config**
3. **dpvm activate**

DETAILED STEPS

	Command or Action	**Purpose**

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

Step 2	**dpvm database diff config** **Example:** `switch(config)# dpvm database diff config`	(Optional) Compares the DPVM config database to the current DPVM active database. Use this output to verify your config database changes before activating them. You can optionally use the **dpvm database copy active** command to copy the active database into the config database to discard the old config database.
Step 3	**dpvm activate** **Example:** `switch(config)# dpvm activate`	Copies the DPVM config database into the DPVM active database.

Related Topics
DPVM Databases
Verifying the DPVM Configuration

5.5.4 Clearing the DPVM CFS Session Lock

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

Before You Begin

SUMMARY STEPS
1.   clear dpvm session

DETAILED STEPS

	Command or Action	Purpose
Step 1	**clear dpvm session** **Example:** `switch# clear dpvm session`	Discards the DPVM pending database and releases the CFS lock.

5.5.5 Enabling Autolearning

You can configure the DPVM database to automatically learn (autolearn) about new devices within each VSAN. Autolearning is a two-part process. When you enable autolearning, DPVM creates learned entries by populating device pWWNs and VSANs in the active DPVM database. DPVM learns currently logged in devices as well as any new devices that log in while autolearning is enabled. These learned entries become permanent in the active DPVM database when you disable autolearning.

The following conditions apply to autolearning:

• If a device logs out while autolearn is enabled, that entry is automatically deleted from the active DPVM database.

• If the same device logs multiple times into the switch through different ports, the VSAN that corresponds to last login is remembered

• Learned entries do not override previously configured and activated entries.

Before You Begin
• Ensure that the active DPVM database is already available.

SUMMARY STEPS

1.   **configure terminal**
2.   **dpvm auto-learn**
3.   (Optional)   **show dpvm ports** [**vsan** *vsan-id*]
4.   **no dpvm auto-learn**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**dpvm auto-learn** **Example:** `switch(config)# dpvm auto-learn`	Enables autolearning. Disable autolearning after a period of time to make the learned entries permanent in the active DPVM database.
Step 3	**show dpvm ports** [**vsan** *vsan-id*] **Example:** `switch(config)# show dpvm ports vsan 3`	(Optional) Displays dynamic (autolearned) entries.
Step 4	**no dpvm auto-learn** **Example:** `switch(config)# no dpvm auto-learn`	Disables autolearning. Any learned entries become permanent in the active DPVM database.

5.5.6 Clearing Autolearned Entries

If DPVM autolearning is enabled, you can clear any or all learned entries from the active DPVM database.

Before You Begin
  • Ensure that DPVM autolearning is enabled.

SUMMARY STEPS

1.   **clear dpvm auto-learn pwwn** *pwwn*
2.   **clear dpvm auto-learn**

DETAILED STEPS

	Command or Action	**Purpose**

Step 1	**clear dpvm auto-learn pwwn** *pwwn*
	
	Example:
	`switch# clear dpvm auto-learn pwwn`
	`55:22:33:44:55:66:77:88`
		Clears an individual autolearned entry.

Step 2	**clear dpvm auto-learn** **Example:** `switch# clear dpvm auto-learn 8`	Clears all autolearned entries.

5.5.7 Displaying DPVM Database Merge Results

When you merge two independent fabrics, DPVM attempts to merge the associated DPVM databases. You can review the results of this database merge to determine if it succeeded or failed.

SUMMARY STEPS

Inspur-Cisco Networking
Technology Co.,Ltd.

1. **show dpvm merge status**
2. **show dpvm merge statistics**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**show dpvm merge status** **Example:** `switch(config)# show dpvm merge status`	Displays information about the last DPVM database merge event.
Step 2	**show dpvm merge statistics** **Example:** `switch(config)# show dpvm merge statistics`	Displays statistics about the last DPVM database merge.

The following example shows conflicts in the DPVM database merge:

```
switch# show dpvm merge status
Last Merge Time Stamp     : Fri March 25
15:46:36 2011 Last Merge State  : Fail
Last Merge Result        : Fail
Last Merge Failure Reason : DPVM DB conflict found during merge [cfs_status:
76]
              Last Merge Failure Details: DPVM merge failed due to database
conflict Local Switch WWN : 20:00:00:0d:ec:24:e5:00
Remote Switch WWN        : 20:00:00:0d:ec:09:d5:c0


----------------------------------------------------------------------
        Conflicting DPVM member(s)              Loc VSAN   Rem VSAN
----------------------------------------------------------------------
dev-alias dpvm dev alias 1                      1313     1414
dev-alias dpvm dev alias 2                      1313     1414
dev-alias dpvm_dev_alias_3                      1313     1414
[21:00:00:04:cf:cf:45:bc] [Total 3 conflict(s)]
switch#
```

Related Topics
Database Merge

5.6 Verifying the DPVM Configuration

To display the DPVM configuration, perform one of the following tasks:

Command	Purpose
show dpvm status	Displays the status for the DPVM configuration.
show dpvm database [**active**]	Displays information about DPVM databases.
show dpvm merge {**status**	**statistics**}	Displays information the last DPVM merge event.

show dpvm pending [activation]	Displays information about the DPVM pending
show dpvm pending-diff	Displays the differences between the pending database and the config database.
show dpvm ports [vsan *vsan-id*]	Displays information about the dynamic ports associated with a VSAN.
show dpvm session status	Displays information about DPVM CFS session.

Related Topics
DPVM Databases
Activating the DPVM Config Database

5.7 DPVM Example Configuration

This example shows how to configure a basic DPVM configuration.

SUMMARY STEPS

1. Enable DPVM and DPVM CFS distribution.
2. Activate the DPVM database.
3. Enable autolearning.
4. Access other switches in the fabric to verify the DPVM configuration.
5. Disable autolearning.
6. Access other switches in the fabric to verify the DPVM configuration.

DETAILED STEPS

Step 1   Enable DPVM and DPVM CFS distribution.
Example:
```
switch1# config t
Enter configuration commands, one per line. End with CNTL/Z.
switch1(config)# feature dpvm
switch1(config)# show dpvm status
No active DB, auto-learn is off, distribution is enabled,
Duplicated pwwn will be  Rejected.
```
DPVM is enabled but the active database is empty.

Step 2   Activate the DPVM database.
Example:
```
switch1(config)# dpvm activate
switch1(config)# dpvm commit
switch1(config)# show dpvm database active
switch1(config)#
```
DPVM is enabled but the active database is empty.

Step 3   Enable autolearning.
Example:
```
switch1(config)# dpvm auto-learn
switch1(config)# dpvm commit
```

Inspur-Cisco Networking
Technology Co.,Ltd.

```
switch1(config)# show dpvm database active pwwn
21:00:00:e0:8b:0e:74:8a vsan 4(*) pwwn
21:01:00:e0:8b:2e:87:8a vsan 5(*) [Total 2 entries]
* is auto-learnt entry

switch1# show dpvm ports
--------------------------------------------------------------
Interface    Vsan          Device pWWN        Device nWWN
--------------------------------------------------------------
------ fc1/24    4 21:00:00:e0:8b:0e:74:8a
20:00:00:e0:8b:0e:74:8a fc1/27    5
            21:01:00:e0:8b:2e:87:8a
20:01:00:e0:8b:2e:87:8a
switch1# show flogi database
--------------------------------------------------------------------
--------- INTERFACE  VSAN          FCID                      PORT
NAME             NODE NAME
--------------------------------------------------------------------
--------- fc1/24    4       0xe70100      21:00:00:e0:8b:0e:74:8a
20:00:00:e0:8b:0e:74:8a   fc1/27        5               0xe80100
21:01:00:e0:8b:2e:87:8a   20:01:00:e0:8b:2e:87:8a

Total number of flogi = 2.


switch1# show dpvm status
DB is activated successfully, auto-learn is on
```

The currently logged in devices (and their current VSAN assignment) populate the active DPVM database. However, these autolearned entries are not permanent in the active DPVM database. The output of the **show dpvm ports** and the **show flogi database** commands display two other devices that have logged in (referred to as switch9 and switch3 in this sample configuration).

Step 4   Access other switches in the fabric to verify the DPVM configuration.
          Example:
```
switch9# show dpvm database active pwwn
21:00:00:e0:8b:0e:87:8a vsan 1(*) pwwn
21:01:00:e0:8b:2e:74:8a vsan 1(*) [Total 2 entries]
* is auto-learnt entry

switch9# show dpvm status
DB is activated successfully, auto-learn is on


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

switch3# show dpvm status
DB is activated successfully, auto-learn is on
```

The autolearned entries show up in the active database for other switches in the fabric.

Step 5   Disable autolearning.
          Example:
```
 switch1(config)# no dpvm auto-learn
switch1(config)# dpvm commit

switch1# show dpvm status
DB is activated successfully, auto-
learn is off switch1# show dpvm
database active
pwwn
21:00:00:e0:8b:0e:74:8a
vsan        4        pwwn
21:01:00:e0:8b:2e:87:8a
vsan        5        pwwn
```

```
21:00:00:e0:8b:0e:87:8a
vsan          1          pwwn
21:01:00:e0:8b:2e:74:8a
vsan          1          pwwn
21:00:00:e0:8b:0e:76:8a
vsan          1          pwwn
21:01:00:e0:8b:2e:76:8a
vsan 1 [Total 6 entries]
* is auto-learnt entry

switch1# show dpvm status
DB is activated successfully, auto-learn is off
```

The autolearned entries are now permanent in the active DPVM database.

Step 6   Access other switches in the fabric to verify the DPVM configuration.
            Example:
```
 switch9#      show      dpvm
database    active   pwwn
21:00:00:e0:8b:0e:87:8a
vsan          1          pwwn
21:01:00:e0:8b:2e:74:8a
vsan          1          pwwn
21:00:00:e0:8b:0e:76:8a
vsan          1          pwwn
21:01:00:e0:8b:2e:76:8a
vsan          1          pwwn
21:00:00:e0:8b:0e:74:8a
vsan          4          pwwn
21:01:00:e0:8b:2e:87:8a
vsan 5 [Total 6 entries]
* is auto-learnt entry

switch9# show dpvm status
DB is activated successfully, auto-learn is
off

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

switch3# show dpvm status
DB is activated successfully, auto-learn is off
```

The autolearned entries show up in the active database for other switches in the fabric.

5.8 Feature History

Table 6 : Feature History for DPVM

Feature Name	Releases	Feature Information
DPVM	5.2(1)	This feature was introduced.

CHAPTER 6 Configuring VSAN Trunking

This chapter describes how to configure VSAN trunking. This chapter includes the following sections:
• Configuring VSAN Trunking

6.1 Configuring VSAN Trunking

6.1.1 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.

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 8 : 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 Inspur SAN switches (see the following figure).

Figure 9 : Third-Party Switch VSAN Mismatch

VSAN 2 and VSAN 3 are effectively merged with overlapping entries in the name server and the zone applications. Inspur DCNM for SAN 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.

6.1.2 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.

SUMMARY STEPS
1. **configure terminal**
2. **no trunk protocol enable**
3. **trunk protocol enable**

DETAILED STEPS

	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 virtual 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 7 : Trunk Mode Status Between Switches

Your Trunk Mode	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 Inspur SAN switches is that one side of the trunk is set to auto and the other is set to on.

Configuring Trunk Mode

You can configure trunk mode.

SUMMARY STEPS
1. **configure terminal**
2. switch(config)# **interface fc** *slot*/*port*
3. **interface vfc** *vfc-id*
4. **switchport trunk mode on**
5. **switchport trunk mode auto**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	switch(config)# **interface fc** *slot*/*port*	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.

Step 5	**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 10 : 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 11 : Operational and Allowed VSAN Configuration
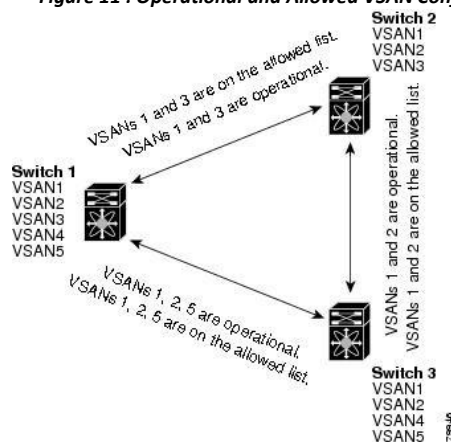
Configuring an Allowed-Active List of VSANs

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

SUMMARY STEPS
1. **configure terminal**
2. **interface fc** *slot/port*
3. **switchport trunk allowed vsan** *vsan-id* **-** *vsan-id*
4. **switchport trunk allowed vsan add** *vsan-id*
5. **switchport trunk allowed vsan all**
6. **no switchport trunk allowed vsan** *vsan-id* **-** *vsan-id*
7. **no switchport trunk allowed vsan add** *vsan-id*

DETAILED STEPS

		Command or Action	Purpose

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

Step 2		**interface fc** *slot/port* **Example:** `switch(config)# interface fc 3`	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.
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		**switchport trunk allowed vsan all** **Example:** `switch(config-if)# switchport trunk allowed vsan all`	Adds all VSANs to the new allowed list.
Step 6		**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 7		**no switchport trunk allowed vsan add** *vsan-id* **Example:** `switch(config-if)# no switchport trunk allowed vsan add 40`	Deletes the expanded allowed list.

6.1.3 Default Settings for VSAN Trunks

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

Table 8 : 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 7 Configuring and Managing Zones

This chapter describes how to configure and manage zones. This chapter contains the following sections:
• Information About Zones

7.1 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.

7.1.1 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 Inspur switch domain and additionally specifies a port belonging to a non-Inspur switch.

- 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.

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 12 : 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 13 : Fabric with Three Zones
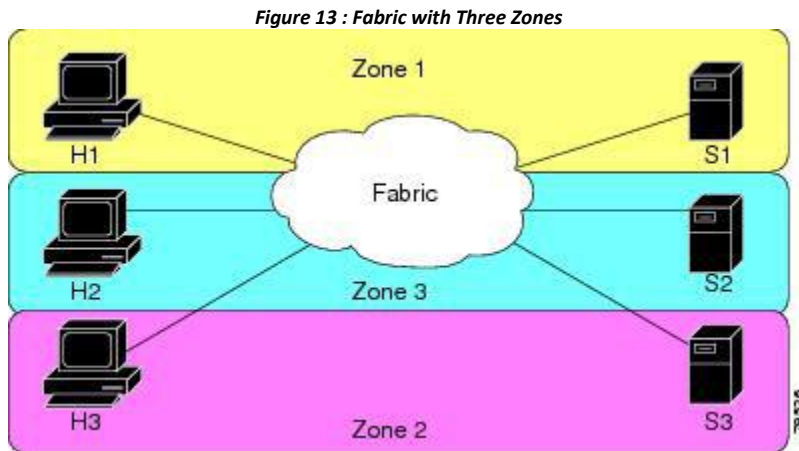
Zone Implementation

Inspur 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.

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

Figure 14 : Active and Full Zone Sets
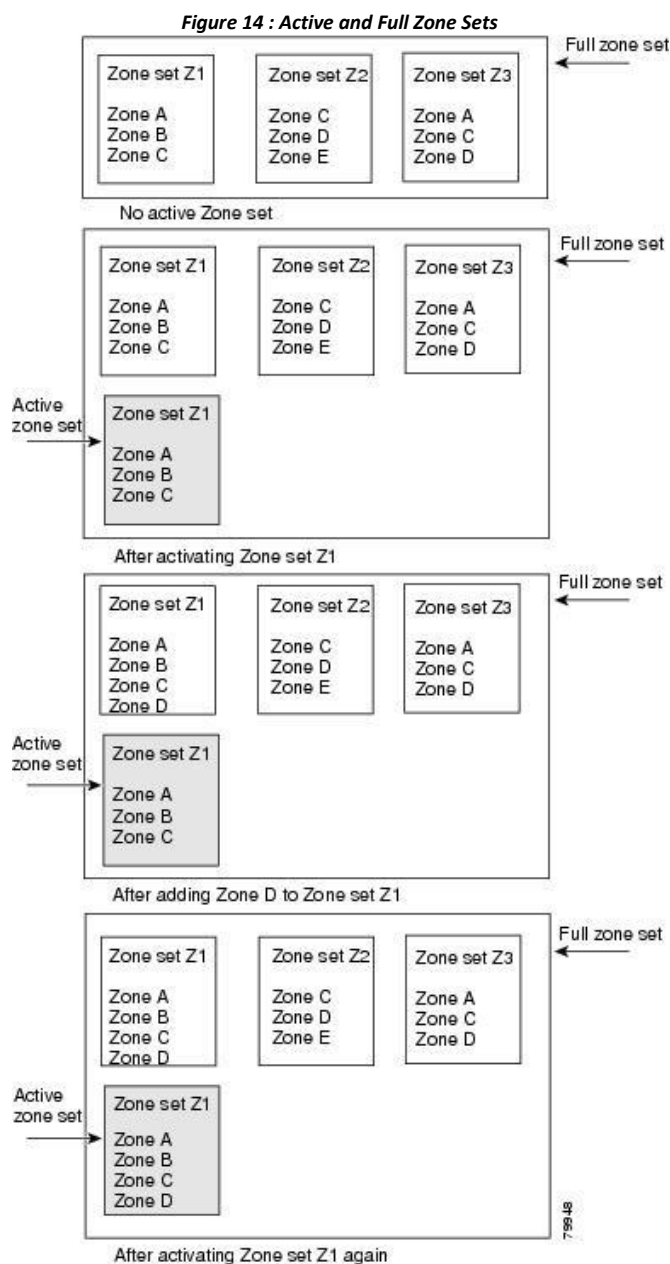
7.1.2 Configuring a Zone

You can configure a zone and assign a zone name.

SUMMARY STEPS
1. **configure terminal**
2. **zone name** *zone-name* **vsan** *vsan-id*
3. **member** *type value*

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**zone name** *zone-name* **vsan** *vsan-id*
	
	Example:
	`switch(config)# zone name test`
	`vsan 5`
		Configures a zone in the specified VSAN.
		
		Note All alphanumeric characters or one of the following symbols ($,
		
		-, ^, _) are supported.

Step 3	**member** *type value*
	
	Example:
	`switch(config-zone)# member`
	` interface 4`
		Configures a member for the specified zone based on the type (pWWN, fabric pWWN, FC ID, fcalias, domain ID, or interface) and value specified.
		
		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
switch(config-zone)# member interface fc 2/1
switch(config-zone)# member interface fc 2/1 swwn 20:00:00:05:30:00:4a:de
switch(config-zone)# member interface fc 2/1 domain-id 25
```
The following example shows how to configure different types of member alias:

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

Inspur-Cisco Networking
Technology Co.,Ltd.

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
```

7.1.3 Zone Sets

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

Figure 15 : Hierarchy of Zone Sets, Zones, and Zone Members

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

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

Activating a Zone Set

You can activate or deactivate an existing zone set.
Changes to a zone set do not take effect in a full zone set until you activate it.

SUMMARY STEPS
1. **configure terminal**
2. **zoneset activate name** *zoneset-name* **vsan** *vsan-id*
3. **no zoneset activate name** *zoneset-name* **vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**zoneset activate name** *zoneset-name* **vsan** *vsan-id* **Example:** `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.

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.

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.

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

Configuring the Default Zone Access Permission

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

SUMMARY STEPS

1. **configure terminal**
2. **zone default-zone permit vsan** *vsan-id*
3. **no zone default-zone permit vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

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

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

FC Alias Creation

You can assign an alias name and configure an alias member using the following values:
• pWWN—The WWN of the N port is in hex format (for example, 10:00:00:23:45:67:89:ab).
• fWWN—The WWN of the fabric port name is in hex format (for example, 10:00:00:23:45:67:89:ab).
• FC ID—The N port ID is in 0xhhhhhh format (for example, 0xce00d1).
• Domain ID—The domain ID is an integer from 1 to 239. A mandatory port number of a non-Inspur 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.

SUMMARY STEPS
1. **configure terminal**
2. **fcalias name** *alias-name***vsan** *vsan-id*
3. **member** *type value*

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**fcalias name** *alias-name***vsan** *vsan-id*
	
	Example:
	`switch(config)# fcalias name testname`
	` vsan 50`
		Configures an alias name. The alias name can be any case-sensitive, alphanumeric string up to 64 characters.

Step 3	**member** *type value* **Example:** `switch(config-fcalias)# member pwwn 4`	Configures a member for the specified fcalias based on the type (pWWN, fabric pWWN, FC ID, domain ID, or interface) and value specified. **Note** Multiple members can be specified on multiple lines.

Creating FC Aliases Example

Table 9 : Type and Value Syntax for the member Command

Device alias	**member device-alias** *device-alias*
Domain ID	**member domain-id** *domain-id* **portnumber** *number*
FC ID	**member fcid** *fcid*
Fabric pWWN	**member fwwn** *fwwn-id*
Local sWWN interface	**member interface** *type slot/port*
Domain ID interface	**member interface** *type slot/port* **domain-id** domain-id
Remote sWWN interface	**member interface** *type slot/port* **swwn** *swwn-id*
pWWN	**member pwwn** *pwwn-id*

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

```
switch(config)# fcalias name AliasSample vsan 3
```
pWWN example:

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

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

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

```
switch(config-fcalias)# member domain-id 2
portnumber 23 switch(config-fcalias)# member
interface fc 2/1 switch(config-fcalias)# member
interface fc2/1 domain-id 25
```
Device alias example:

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

Creating Zone Sets and Adding Member Zones

You can create a zone set to include several zones.

SUMMARY STEPS

1. **configure terminal**
2. **zone set name** *zoneset-name* **vsan** *vsan-id*
3. **member** *name*

Inspur-Cisco Networking
Technology Co.,Ltd.

4. **zone name** *zone-name*
5. **member fcid** *fcid*

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**zone set name** *zoneset-name* **vsan** *vsan-id*
	Example:
	`switch(config)# zone set name new vsan`
	`23`
		Configures a zone set with the configured zoneset-name.
		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
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`
	`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.

Inspur SAN switches support both hard and soft zoning.

7.1.4 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 10 : 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 Inspur SAN switches distribute active zone sets when new E port links come up or when a new zone set is activated in a VSAN. The zone set distribution takes effect while sending merge requests to the adjacent switch or while activating a zone set.

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

SUMMARY STEPS
1. **configure terminal**
2. **zoneset distribute full vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

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

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

Enabling a One-Time Distribution

You can perform a one-time distribution of inactive, unmodified zone sets throughout the fabric. Use the **zoneset distribute vsan** *vsan-id c*ommand 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

Inspur-Cisco Networking
Technology Co.,Ltd.

information to the startup configuration.

Use the **show zone status vsan** *vsan-id c*ommand 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 16 : Importing and Exporting the Database
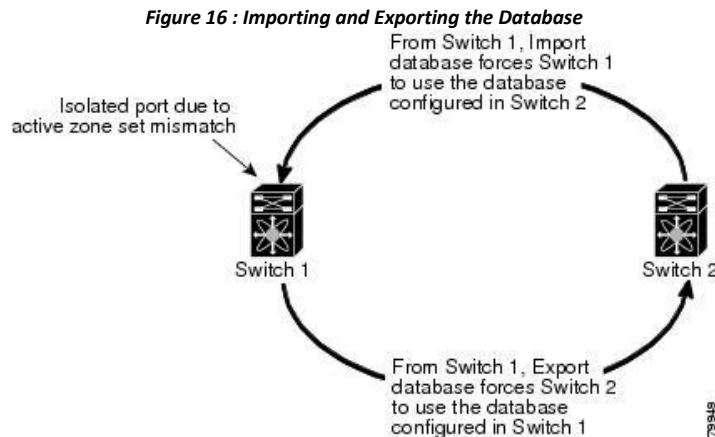
Importing and Exporting Zone Sets

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

SUMMARY STEPS
1.   **zoneset import interface** {**vfc** | **vfc-port-channel**} *if-number* **vsan** *vsan-id*
2.   **zoneset export vsan** *vsan-id*

DETAILED STEPS

	Command or Action	**Purpose**
Step 1	**zoneset import interface** {**vfc**	**vfc-port-channel**} *if-number* **vsan** *vsan-id* **Example:** `switch# zoneset import interface 6 vsan 10`	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.

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

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

DETAILED STEPS

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

Renaming Zones, Zone Sets, and Aliases

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

SUMMARY STEPS
1. **configure terminal**
2. **zoneset rename** *oldname newname* **vsan** *vsan-id*
3. **zone rename** *oldname newname* **vsan** *vsan-id*
4. **fcalias rename** *oldname newname* **vsan** *vsan-id*
5. **zone-attribute-group rename** *oldname newname* **vsan** *vsan-id*
6. **zoneset activate name** *newname* **vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**zoneset rename** *oldname newname* **vsan** *vsan-id* **Example:** `switch(config)# zoneset rename test myzoneset vsan 60`	Renames a zone set in the specified VSAN.
Step 3	**zone rename** *oldname newname* **vsan** *vsan-id* **Example:** `switch(config)# zone rename test myzone vsan 50`	Renames a zone in the specified VSAN.
Step 4	**fcalias rename** *oldname newname* **vsan** *vsan-id* **Example:** `switch(config)# fcalias rename test myfc vsan 200`	Renames a fcalias in the specified VSAN.
Step 5	**zone-attribute-group rename** *oldname newname* **vsan** *vsan-id* **Example:** `switch(config)# zone-attribute-group rename test`	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.

SUMMARY STEPS

1. **configure terminal**
2. **zoneset clone** *oldname newname* **vsan** *vsan-id*
3. **zone clone** *oldname newname* **vsan** *number*

---

4. **fcalias clone** *oldname newname* **vsan** *vsan-id*
5. **zone-attribute-group clone** *oldname newname* **vsan** *vsan-id*
6. **zoneset activate name** *newname* **vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**zoneset clone** *oldname newname* **vsan** *vsan-id* Example: `switch(config)# zoneset clone test myzoneset2 vsan 2`	Clones a zone set in the specified VSAN.
Step 3	**zone clone** *oldname newname* **vsan** *number* Example: `switch(config)# zone clone test myzone3 vsan 3`	Clones a zone in the specified VSAN.
Step 4	**fcalias clone** *oldname newname* **vsan** *vsan-id* Example: `switch(config)# fcalias clone test myfcalias vsan 30`	Clones a fcalias in the specified VSAN.
Step 5	**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
```

7.1.6 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.
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.

7.1.7 Information About Smart Zoning

Smart Zoning supports zoning among more devices by reducing the number of zoning entries that needs to be programmed by considering device type information without increasing the size of the zone set. Smart zoning enables you to select the host, target, or both as the end device type.

7.1.8 Smart Zoning Member Configuration

The following table displays the supported smart zoning member configurations.

Feature	Supported
PWWN	Yes
FCID	Yes
FC Alias	Yes
Device-alias	Yes
Interface	No
IP Address	No
Symbolic nodename	No
FWNN	No
Domain ID	No

7.1.9 Enabling Smart Zoning on a VSAN

To enable smart zoning for a VSAN, perform these tasks:

SUMMARY STEPS
1. switch# **configure terminal**
2. switch(config)# **system default zone smart-zone enable**

**Inspur-Cisco Networking
Technology Co.,Ltd.**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **system default zone smart-zone enable**	Enables smart zoning on a VSAN that are created based on the specified default value.

7.1.10 Setting Default Value for Smart Zoning

To set the default value, perform these tasks:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **zone smart-zone enable vsan1**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **zone smart-zone enable vsan1**	Smart Zoning is enabled for the selected zone.

7.1.11 Configuring Device Types for Zone Members

To configure the device types for zone members, perform these tasks:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch (config)# **zoneset name zoneset1**
3. switch(config-zoneset-zone)# **member device-alias mds9000 both**
4. switch(config-zoneset-zone)# **member pwnn 10:00:00:23:45:67:89:ab**
5. switch(config-zoneste-zone)# **member fcid 2020:dbc0:80::4076**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch (config)# **zoneset name zoneset1**	Enters global configuration mode.
Step 3	switch(config-zoneset-zone)# **member device-alias mds9000 both**	Configures the device type for the device-alias member as both. For every supported member-type, init, target, and both are supported.
Step 4	switch(config-zoneset-zone)# **member pwnn 10:00:00:23:45:67:89:ab**	Configures the device type for the pwnn member as target. For every supported member-type, init, target, and both are

| Step 5 | switch(config-zoneste-zone)# **member fcid 2020:dbc0:80::4076** | Configures the device type for the FCID member. There is no specific device type that is configured. For every supported member-type, init, target, and both are supported. |
| | | **Note** When there is no specific device type configured for a zone member, at the backend, zone entries that are generated are created as device type both. |

7.1.12 Converting Zones Automatically to Smart Zoning

To fetch the device-type information from the name server and to add that information to the member, perform these tasks:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **zone convert smart-zoning zone fcalias name fc_1 vsan 1**
3. switch(config)# **zone convert smart-zoning zone sz_1 vsan 1**
4. switch(config)# **zone convert smart-zoning zone szs_1 vsan 1**
5. switch(config)# **zone convert smart-zoning vsan 1**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **zone convert smart-zoning zone fcalias name fc_1 vsan 1**	Fetches the device type information from the name server for the fcalias members.
Step 3	switch(config)# **zone convert smart-zoning zone sz_1 vsan 1**	Fetches the device type information from the name server for the zone members.
Step 4	switch(config)# **zone convert smart-zoning zone szs_1 vsan 1**	Fetches the device type information from the name server for all the zones and the fcalias members in the specified zoneset.
Step 5	switch(config)# **zone convert smart-zoning vsan 1**	Fetches the device type information from the nameserver for all the zones and fcalias members for all the zonesets present in the VSAN.

7.1.13 Removing Smart Zoning Configuration

To remove the smart zoning configuration, perform these tasks:

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **clear zone smart-zoning fcalias name zone 1 vsan vsan1**
3. switch(config)# **clear zone smart-zoning zone name zone 1 vsan 1**
4. switch(config)# **clear zone smart-zoning zoneset zoneset1 vsan1**
5. switch(config)# **clear zone smart-zoning vsan1**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **clear zone smart-zoning fcalias name zone 1 vsan vsan1**	Removes the device type configuration for all the members of the specified fcalias.
Step 3	switch(config)# **clear zone smart-zoning zone name zone 1 vsan 1**	Removes the device type configuration for all the members of the specified zone .
Step 4	switch(config)# **clear zone smart-zoning zoneset zoneset1 vsan1**	Removes the device type configuration for all the members of the zone and fcalias for the specified zoneset.
Step 5	switch(config)# **clear zone smart-zoning vsan1**	Removes the device type configuration for all the members of the zone and fcalias of all the specified zonesets in the VSAN.

7.1.14 Disabling Smart Zoning at Zone Level

To disable smart zoning at the zone level, perform these tasks:

SUMMARY STEPS
1.   switch# **configure terminal**
2.   switch(config)# **zone name zone1 vsan 1**
3.   switch(config)# **no attribute disable-smart-zoning**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters the global configuration mode.
Step 2	switch(config)# **zone name zone1 vsan 1**	Configures as zone name.
Step 3	switch(config)# **no attribute disable-smart-zoning**	Smart Zoning is disabled for the selected zone. **Note** This command only disables the smart zoning for the selected zone and does not remove the device type configurations.

7.1.15 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 Inspur SAN switches.

<div align="center">Table 11 : Advantages of Enhanced Zoning</div>

Basic Zoning	Enhanced Zoning	Enhanced ZoningAdvantages

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

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

Inspur SAN switches allow you to change from enhanced zoning to basic zoning to enable you to downgrade and upgrade to other Inspur INOS releases.

Step 1   Verify that the active and full zone set do not contain any configuration that is specific to the enhanced zoning mode.
Step 2   If such configurations exist, delete them before proceeding with this procedure. If you do not delete the existing configuration, the switch software automatically removes them.
Step 3   Set the operation mode to basic zoning mode.

Enabling Enhanced Zoning

You can enable enhanced zoning in a VSAN.

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

SUMMARY STEPS
1.   **configure terminal**
2.   **zone mode enhanced vsan** *vsan-id*
3.   **no zone mode enhanced vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

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

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

Modifying the Zone Database

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.

Inspur-Cisco Networking
Technology Co.,Ltd.

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

DETAILED STEPS

	Command or Action	Purpose

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

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

Releasing Zone Database Locks

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

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

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

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

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 12 : 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.

SUMMARY STEPS
1. **configure terminal**
2. **zone merge-control restrict vsan** *vsan-id*
3. **no zone merge-control restrict vsan** *vsan-id*
4. **zone commit vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

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

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

Default Zone Policies

You can permit or deny traffic in the default zone.

SUMMARY STEPS
1. **configure terminal**
2. **zone default-zone permit vsan** *vsan-id*
3. **no zone default-zone permit vsan** *vsan-id*
4. **zone commit vsan** *vsan-id*

DETAILED STEPS

	Command or Action	**Purpose**
Step 1	**configure terminal**	Enters global configuration mode.
	Example: `switch# configure terminal switch(config)#`	
Step 2	**zone default-zone permit vsan** *vsan-id*	Permits traffic flow to default zone members.
	Example: `switch(config)# zone default-zone permit vsan 12`	
Step 3	**no zone default-zone permit vsan** *vsan-id*	Denies traffic flow to default zone members and reverts to factory default.
	Example: `switch(config)# no zone default-zone permit vsan 10`	
Step 4	**zone commit vsan** *vsan-id*	Commits the changes made to the specified VSAN.
	Example: `switch(config)# zone commit vsan 340`	

Configuring System Default Zoning Settings

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

switch.

SUMMARY STEPS

1. **configure terminal**
2. **system default zone default-zone permit**
3. **no system default zone default-zone permit**
4. **system default zone distribute full**
5. **no system default zone distribute full**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**system default zone default-zone permit**
	
	Example:
	`switch(config)# system default zone default-`
	`zone permit`
		Configures permit as the default zoning policy for new VSANs on the switch.

Step 3	**no system default zone default-zone permit**
	
	Example:
	`switch(config)# no system default`
	`zone default-zone permit`
		Configures deny (default) as the default zoning policy for new VSANs on the switch.

Step 4	**system default zone distribute full**
	
	Example:
	`switch(config)# system default zone`
	`distribute full`
		Enables full zone database distribution as the default for new VSANs on the switch.

Step 5	**no system default zone distribute full**
	
	Example:
	`switch(config)# no system default zone`
	`distribute full`
		Disables (default) full zone database distribution as the default for new VSANs on the switch. Only the active zone database is distributed.

Verifying Enhanced Zone Information

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

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

7.1.16 Compacting the Zone Database

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

SUMMARY STEPS

1. **configure terminal**
2. **no zone name** *zone-name* **vsan** *vsan-id*

---

3. **zone compact vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**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.

7.1.17 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**

7.1.18 Default Settings for Zones

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

Table 13 : 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

8.1 Distributing Device Alias Services

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

8.1.1 Information About Device Aliases

Inspur 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 Inspur 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 Inspur 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.

Related Topics
Device Alias Modes

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 14 : Comparison Between Zone Aliases and Device Aliases

Zone-Based Aliases	Device Aliases
Aliases are limited to the specified VSAN.	You can define device aliases without specifying the VSAN number. You can also use the same definition in one or more VSANs without any restrictions.
Zone aliases are part of the zoning configuration. The alias mapping cannot be used to configure other features.	Device aliases can be used with any feature that uses the pWWN.
You can use any zone member type to specify the end devices.	Only pWWNs are supported.
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.

8.1.2 Device Alias Databases

The device alias feature uses two databases to accept and implement device alias configurations.
- Effective database—The database currently used by the fabric.
- Pending database—Your subsequent device alias configuration changes are stored in the pending database.

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

Device alias database changes are validated with the applications. If any of the applications cannot accept the device alias database changes, then those changes are rejected; this applies to device alias database changes resulting from either a commit or merge operation.

Creating Device Aliases

You can create a device alias in the pending database.

SUMMARY STEPS
1.   **configure terminal**
2.   **device-alias database**
3.   **device-alias name** *device-name* **pwwn** *pwwn-id*
4.   **no device-alias name** *device-name*
5.   **device-alias rename** *old-device-name new-device-name*

DETAILED STEPS

	Command or Action	Purpose

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

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

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

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

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

EXAMPLES

This example shows how to display the device alias configuration.

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

Device Alias Modes

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.

Device Alias Mode Guidelines and Limitations for Device Alias Services

Device Alias services have these configuration guidelines and limitations:

• If two fabrics running in different device alias modes are joined together, the device alias merge fails. There is no automatic conversion to one mode or the other during the merge process. In this situation, you must select one mode over the other.

• Before changing from enhanced to basic mode, you must first explicitly remove all native device alias-based configurations from both local and remote switches, or replace all device alias-based configuration members with the corresponding pWWN.

• If you remove a device alias from the device alias database, all applications automatically stop enforcing the corresponding device alias. If that corresponding device alias is part of an active zone set, all the traffic to and from that pWWN is disrupted.

• Renaming the device alias not only changes the device alias name in the device alias database, but also replaces the corresponding device alias configuration in all of the applications.

• When a new device alias is added to the device alias database, and the application configuration is present on that device alias, it automatically takes effect. For example, if the corresponding device alias is part of the active zoneset and the device is online, then zoning is enforced automatically. You do not have to reactivate the zone set.

• If a device alias name is mapped to a new HBA's pWWN, the application's enforcement changes accordingly. In this case, the zone server automatically enforces zoning based on the new HBA's pWWN.

Configuring Device Alias Modes

You can configure device aliases to operate in enhanced mode.

SUMMARY STEPS
1. **configure terminal**
2. **device-alias mode enhanced**
3. **no device-alias mode enhance**

DETAILED STEPS

	Command or Action	Purpose

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

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

EXAMPLES

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

```
switch# show device-alias status
Fabric Distribution: Enabled
Database:- Device Aliases 0 Mode: Basic
Locked By:- User "admin" SWWN
```

**Inspur-Cisco Networking
Technology Co.,Ltd.**

```
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.

SUMMARY STEPS
1. **configure terminal**
2. **device-alias commit**

DETAILED STEPS

	Command or Action	Purpose

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

| Step 2 | **device-alias commit**

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

Discarding Changes

You can discard the device alias session changes.

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

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

SUMMARY STEPS
1. **configure terminal**
2. **device-alias abort**

DETAILED STEPS

	Command or Action	Purpose

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

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

EXAMPLES

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

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

SUMMARY STEPS
1. **configure terminal**
2. **no device-alias distribute**
3. **device-alias distribute**

DETAILED STEPS

	Command or Action	Purpose

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

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

8.1.3 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.

SUMMARY STEPS
1. **configure terminal**
2. **device-alias import fcalias vsan** *vlan-id*

DETAILED STEPS

	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.

8.1.4 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 20K.

If the combined number of device entries in both databases exceeds the supported configuration limit, then the merge will fail.

8.1.5 Verifying the Device Alias Configuration

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

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

8.1.6 Default Settings for Device Alias Services

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

Table 15 : 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 Configuring Fibre Channel Routing Services and Protocols

This chapter contains the following sections:
• Information About Fibre Channel Routing Services and Protocols

9.1 InformationAboutFibreChannelRoutingServicesandProtocols

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

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

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

• Bases path status on a link state protocol.

• Routes hop by hop, based only on the domain ID.

• Runs only on E ports or TE ports and provides a loop free topology.

• Runs on a per VSAN basis. Connectivity in a given VSAN in a fabric is guaranteed only for the switches configured in that VSAN.

• Uses a topology database to keep track of the state of the links on all switches in the fabric and associates a cost with each link.

• Guarantees a fast reconvergence time in case of a topology change. Uses the standard Dijkstra algorithm, but there is a static dynamic option for a more robust, efficient, and incremental Dijkstra algorithm. The reconvergence time is fast and efficient as the route computation is done on a per VSAN basis.

9.1.1 Information About FSPF

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

• Supports multipath routing.

• Bases path status on a link state protocol.

• Routes hop by hop, based only on the domain ID.

• Runs only on E ports or TE ports and provides a loop free topology.

• Runs on a per VSAN basis. Connectivity in a given VSAN in a fabric is guaranteed only for the switches configured in that VSAN.

• Uses a topology database to keep track of the state of the links on all switches in the fabric and associates a cost with each link.

• Guarantees a fast reconvergence time in case of a topology change. Uses the standard Dijkstra algorithm, but there is a static dynamic option for a more robust, efficient, and incremental Dijkstra algorithm. The reconvergence time is fast and efficient as the route computation is done on a per VSAN basis.

FSPF Examples

Fault Tolerant Fabric Example

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

Figure 17 : Fault Tolerant Fabric

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

Redundant Link Example

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

By bundling pairs of physical links, FSPF efficiency is considerably improved by the reduced database size and the frequency of link updates. Once physical links are aggregated, failures are not attached to a single link but to the entire SAN port channel. This configuration also improves the resiliency of the network. The failure of a link in a SAN port channel does not trigger a route change, which reduces the risks of routing loops, traffic loss, or fabric downtime for route reconfiguration.

Figure 18 : Fault Tolerant Fabric with Redundant Links

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

9.1.2 FSPF Global Configuration

By default, FSPF is enabled on Inspur SAN switches.

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

Caution

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

SPF Computational Hold Times

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

Link State Records

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

The following table displays the default settings for switch responses.

Table 16 : LSR Default Settings

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

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

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

Configuring FSPF on a VSAN

You can configure an FSPF feature for the entire VSAN.

SUMMARY STEPS
1. **configure terminal**
2. **fspf config vsan** *vsan-id*
3. **spf static**
4. **spf hold-time** *value*
5. **region** *region-id*

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**fspf config vsan** *vsan-id* **Example:** `switch(config)# fspf config vsan 14`	Enters FSPF global configuration mode for the specified VSAN. **Note** User needs to configure the VSAN on which FSPF is being configured.
Step 3	**spf static** **Example:** `switch-config-(fspf-config)# spf static`	Forces static SPF computation for the dynamic (default) incremental VSAN.
Step 4	**spf hold-time** *value* **Example:** `switch-config-(fspf-config)# spf hold-time 10`	Configures the hold time between two route computations in milliseconds (msec) for the entire VSAN. The default value is 0. **Note** If the specified time is shorter, the routing is faster. However, the processor consumption increases accordingly.
Step 5	**region** *region-id* **Example:** `switch-config-(fspf-config)# region 1`	Configures the autonomous region for this VSAN and specifies the region ID.

Resetting FSPF to the Default Configuration

You can return the FSPF VSAN global configuration to its factory default.

SUMMARY STEPS

1. **configure terminal**
2. **no fspf config vsan** *vsan-id*

DETAILED STEPS

	Command or Action	**Purpose**

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

Step 2	**no fspf config vsan** *vsan-id* **Example:** `switch(config)# no fspf config vsan 24`	Deletes the FSPF configuration for the specified VSAN.

Enabling or Disabling FSPF

You can enable or disable FSPF routing protocols.

SUMMARY STEPS

1. **configure terminal**
2. **fspf enable vsan** *vsan-id*
3. **no fspf enable vsan** *vsan-id*

DETAILED STEPS

Inspur-Cisco Networking
Technology Co.,Ltd.

		Command or Action	Purpose

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

Step 2		**fspf enable vsan** *vsan-id* **Example:** `switch(config)# fspf enable vsan 567`	Enables the FSPF routing protocol in the specified VSAN.
Step 3		**no fspf enable vsan** *vsan-id* **Example:** `switch(config)# no fspf enable vsan 567`	Disables the FSPF routing protocol in the specified VSAN.

Clearing FSPF Counters for the VSAN

You can clear the FSPF statistics counters for the entire VSAN.

SUMMARY STEPS
1.   **clear fspf counters vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose
Step 1	**clear fspf counters vsan** *vsan-id* **Example:** `switch# clear fspf counters vsan 345`	Clears the FSPF statistics counters for the specified VSAN. If an interface reference is not specified, all counters are cleared.

9.1.3 FSPF Interface Configuration

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

FSPF Link Cost

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

Configuring FSPF Link Cost

You can configure FSPF link cost.

SUMMARY STEPS
1.   **configure terminal**
2.   **interface vfc** *if-number*

3.   **fspf cost** *value* **vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**interface vfc** *if-number* **Example:** `switch(config)# interface vfc 2`	Configures the specified interface, or if already configured, enters configuration mode for the specified interface.
Step 3	**fspf cost** *value* **vsan** *vsan-id* **Example:** `switch(config-if)# fspf cost 500 vsan 38`	Configures the cost for the selected interface in the specified VSAN.

Hello Time Intervals

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

Configuring Hello Time Intervals

You can configure the FSPF Hello time interval.

SUMMARY STEPS
1.   **configure terminal**
2.   **interface vfc** *if-number*
3.   **fspf hello-interval** *value* **vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**interface vfc** *if-number* **Example:** `switch(config)# interface vfc 21`	Configures the specified interface, or if already configured, enters configuration mode for the specified interface.

Step 3	**fspf hello-interval** *value* **vsan** *vsan-id*
	
	Example:
	`switch(config-if)# fspf hello-interval 25`
	` vsan 10`
		Specifies the hello message interval to verify the health of the link in the VSAN. The default is 20 seconds.

Dead Time Intervals

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

Caution

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

Configuring Dead Time Intervals

You can configure the FSPF dead time interval.

SUMMARY STEPS
1. **configure terminal**
2. **interface vfc** *if-number*
3. **fspf dead-interval** *value* **vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**interface vfc** *if-number* **Example:** `switch(config)# interface vfc 10`	Configures the specified interface, or if already configured, enters configuration mode for the specified interface.

Step 3	**fspf dead-interval** *value* **vsan** *vsan-id*
	
	Example:
	`switch(config-if)# fspf dead-interval`
	`60 vsan 101`
		Specifies the maximum interval for the specified VSAN before which a hello message must be received on the selected interface before the neighbor is considered lost. The default is 80 seconds.

Retransmitting Intervals

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

Configuring Retransmitting Intervals

You can configure the FSPF retransmit time interval.

SUMMARY STEPS
1. **configure terminal**
2. **interface vfc** *if-number*
3. **fspf retransmit-interval** *value* **vsan** *vsan-id*

Inspur-Cisco Networking
Technology Co.,Ltd.

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**interface vfc** *if-number* **Example:** `switch(config)# interface vfc 23`	Configures the specified interface, or if already configured, enters configuration mode for the specified interface.

Step 3	**fspf retransmit-interval** *value* **vsan** *vsan-id*
	
	Example:
	`switch(config-if)# fspf retransmit-interval`
	`10 vsan 25`
		Specifies the retransmit time interval for unacknowledged link state updates in the specified VSAN. The default is 5 seconds.

About Disabling FSPF for Specific Interfaces

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

Disabling FSPF for Specific Interfaces

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

SUMMARY STEPS

1. **configure terminal**
2. **interface vfc** *if-number*
3. **fspf passive vsan** *vsan-id*
4. **no fspf passive vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**interface vfc** *if-number* **Example:** `switch(config)# interface vfc 5`	Configures a specified interface, or if already configured, enters configuration mode for the specified interface.

Inspur-Cisco Networking
Technology Co.,Ltd.

Step 3	fspf passive vsan *vsan-id* **Example:** `switch(config-if)# fspf passive vsan 24`	Disables FSPF for the specified interface in the specified VSAN.
Step 4	no fspf passive vsan *vsan-id* **Example:** `switch(config-if)# no fspf passive vsan 23`	Reenables FSPF for the specified interface in the specified VSAN.

Clearing FSPF Counters for an Interface

You can clear the FSPF statistics counters for an interface.

SUMMARY STEPS

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

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **clear fspf counters vsan** *vsan-id* **interface fc** *slot*/*port*.	Clears the FSPF statistics counters for the specified interface in the specified VSAN.
Step 2	**clear fspf counters vsan** *vsan-id* **intrface** *type_if-number*.	`switch# clear fspf counters vsan 12` `interface pwwn 9`

9.1.4 FSPF Routes

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

Fibre Channel Routes

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

Figure 19 : Fibre Channel Routes
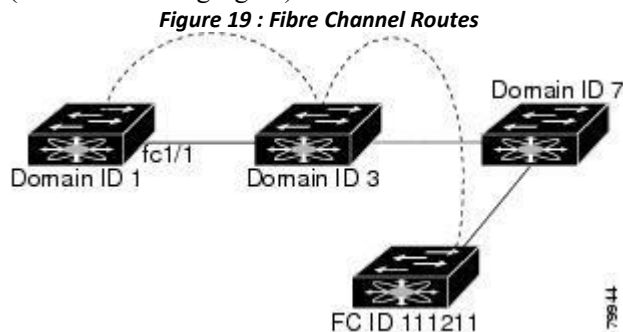
9.1.5 In-Order Delivery

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

Some Fibre Channel protocols or applications cannot handle out-of-order frame delivery. In these cases, Inspur

SAN switches preserve frame ordering in the frame flow. The source ID (SID), destination ID (DID), and optionally, the originator exchange ID (OX ID) identify the flow of the frame.

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

Use IOD only if your environment cannot support out-of-order frame delivery. If you enable IOD, the graceful shutdown feature is not implemented.

Reordering Network Frames

When you experience a route change in the network, the new selected path might be faster or less congested than the old route (See the following figure).

Figure 20 : Route Change Delivery
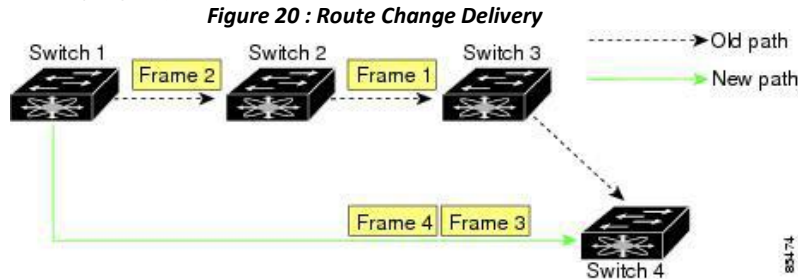
In the figure above, the new path from Switch 1 to Switch 4 is faster. In this scenario, Frame 3 and Frame 4 might be delivered before Frame 1 and Frame 2.

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

• Frames in the network are delivered in the order in which they are transmitted.

• Frames that cannot be delivered in order within the network latency drop period are dropped inside the network.

Reordering SAN Port Channel Frames

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

Figure 21 : Link Congestion Delivery

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

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

• Frames using the old path are delivered before new frames are accepted.

• The new frames are delivered through the new path after the network latency drop period has elapsed and all old frames are flushed.

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

Related Topics

Configuring the Drop Latency Time

About Enabling In-Order Delivery

You can enable IOD for a specific VSAN or for the entire switch. By default, IOD is disabled on Inspur SAN switches.

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

Enabling In-Order Delivery

You can enable in-order delivery for the switch.

SUMMARY STEPS

1. **configuration terminal**
2. **in-order-guarantee**
3. **no in-order-guarantee**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**in-order-guarantee** **Example:** `switch(config)# in-order-guarantee`	Enables in-order delivery in the switch.
Step 3	**no in-order-guarantee** **Example:** `switch(config)# no in-order-guarantee`	Reverts the switch to the factory defaults and disables the in-order delivery feature.

Enabling In-Order Delivery for a VSAN

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

SUMMARY STEPS

1. **configuration terminal**
2. **in-order-guarantee vsan** *vsan-id*
3. **no in-order-guarantee vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**in-order-guarantee vsan** *vsan-id* **Example:** `switch(config)# in-order-guarantee vsan 30`	Enables in-order delivery in the specified VSAN.
Step 3	**no in-order-guarantee vsan** *vsan-id* **Example:** `switch(config)# no in-order-guarantee vsan 30`	Reverts the switch to the factory defaults and disables the in-order delivery feature in the specified VSAN.

Displaying the In-Order Delivery Status

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

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

Configuring the Drop Latency Time

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

SUMMARY STEPS
1. **configure terminal**
2. **fcdroplatency network** *value*
3. **fcdroplatency network** *value* **vsan** *vsan-id*
4. **no fcdroplatency network** *value*

DETAILED STEPS

	Command or Action	Purpose

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

	Command or Action	Purpose
Step 2	**fcdroplatency network** *value* **Example:** `switch(config)# fcdroplatency network 1000`	Configures network drop latency time for the network. The valid range is from 0 to 60000 msec. The default is 2000 msec. **Note** The network drop latency must be computed as the sum of all switch latencies of the longest path in the network.
Step 3	**fcdroplatency network** *value* **vsan** *vsan-id* **Example:** `switch(config)# fcdroplatency network 1000 vsan 12`	Configures network drop latency time for the specified VSAN.
Step 4	**no fcdroplatency network** *value* **Example:** `switch(config)# no fcdroplatency network 1000`	Removes the current fcdroplatency network configuration and reverts the switch to the factory defaults.

Displaying Latency Information

You can view the configured latency parameters by using the **show fcdroplatency** command:

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

9.1.6 Flow Statistics Configuration

Flow statistics count the ingress traffic in the aggregated statistics table. You can collect two kinds of statistics:
- Aggregated flow statistics to count the traffic for a VSAN.
- Flow statistics to count the traffic for a source and destination ID pair in a VSAN.

Flow Statistics

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

Counting Aggregated Flow Statistics

You can count the aggregated flow statistics for a VSAN.

SUMMARY STEPS

1. **configure terminal**
2. **fcflow stats aggregated index** *value* **vsan** *vsan-id*

---

3. **no fcflow stats aggregated index** *value* **vsan** *vsan-id*

DETAILED STEPS

	Command or Action	**Purpose**

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

Step 2	**fcflow stats aggregated index** *value* **vsan** *vsan-id*
	
	Example:
	`switch(config)# fcflow stats aggregated index 20 vsan`
	`10`
		Enables the aggregated flow counter.

Step 3	**no fcflow stats aggregated index** *value* **vsan** *vsan-id*
	
	Example:
	`switch(config)# no fcflow stats aggregated index 20`
	`vsan 12`
		Disables the aggregated flow counter.

Counting Individual Flow Statistics

You can count the flow statistics for a source and destination FC ID in a VSAN.

SUMMARY STEPS
1. **configure terminal**
2. **fcflow stats index** *value dest-fcid source-fcid netmask* **vsan** *vsan-id*
3. **no fcflow stats aggregated index** *value* **vsan** *vsan-id*

DETAILED STEPS

	Command or Action	**Purpose**

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

Step 2	**fcflow stats index** *value dest-fcid source-fcid netmask* **vsan** *vsan-id* **Example:** `switch(config)# fcflow stats index 10`	Enables the flow counter. **Note** The source ID and the destination ID are specified in FC ID hex format (for example, 0x123aff). The mask can be one of 0xff0000 or 0xffffff.
Step 3	**no fcflow stats aggregated index** *value* **vsan** *vsan-id* **Example:** `switch(config)# no fcflow stats aggregated index 11 vsan 200`	Disables the flow counter.

Clearing FIB Statistics

Use the **clear fcflow stats** command to clear the aggregated flow counter:

```
switch# clear fcflow stats aggregated index 1
```
The following example shows how to clear the flow counters for source and destination FC IDs:

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

Displaying Flow Statistics

Use the **show fcflow stats** commands to view flow statistics:

```
switch# show fcflow stats aggregated
Idx        VSAN     frames
---------- ---------- --
           -------- 6 1     42871
```
The following example shows how to display flow statistics:

```
switch# show fcflow stats
```
The following example shows how to display flow index usage:

```
switch# show fcflow stats usage
2 flows
configured
Configured
flows : 3,7
```
The following example shows how to display global FSPF information for a specific VSAN:

```
switch# show fspf vsan 1
```
The following example shows how to display a summary of the FSPF database for a specified VSAN. If no additional parameters are specified, all LSRs in the database are displayed:

```
switch# show fspf database vsan 1
```
The following example shows how to display FSPF interface information:

```
switch# show fspf vsan 1 interface vfc 1
```

9.1.7 Default Settings for FSFP

The following table lists the default settings for FSPF features.

Table 17 : Default FSPF Settings

Parameters	Default
FSPF	Enabled on all E ports and TE ports
SPF computation	Dynamic
SPF hold time	0
Backbone region	0
Acknowledgment interval (RxmtInterval)	5 seconds
Refresh time (LSRefreshTime)	30 minutes
Maximum age (MaxAge)	60 minutes

**Inspur-Cisco Networking
Technology Co.,Ltd.**

Hello interval	20 seconds
Dead interval	80 seconds
Distribution tree information	Derived from the principal switch (root node)
Routing table	FSPF stores up to 16 equal cost paths to a given destination
Load balancing	Based on destination ID and source ID on different, equal cost paths
In-order delivery	Disabled
Drop latency	Disabled
Static route cost	If the cost (metric) of the route is not specified, the default is 10
Remote destination switch	If the remote destination switch is not specified, the default is direct
Multicast routing	Uses the principal switch to compute the multicast tree

CHAPTER 10 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

10.1 Managing FLOGI, Name Server, FDMI, and RSCN Databases

10.1.1 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
--------------
-------------------------------------------------
              fc2/3  1    0xb200e2  21:00:00:04:cf:27:25:2  20:00:00:04:cf:27:25:
              fc2/3  1    0xb200e1  21:00:00:04:cf:4c:18:6  20:00:00:04:cf:4c:18:
              fc2/3  1    0xb200d1  21:00:00:04:cf:4c:18:6  20:00:00:04:cf:4c:18:
              fc2/3  1    0xb200ce  21:00:00:04:cf:4c:16:f  20:00:00:04:cf:4c:16:
              fc2/3  1    0xb200cd  21:00:00:04:cf:4c:18:f  20:00:00:04:cf:4c:18:
              vfc3/1 2    0xb30100  10:00:00:05:30:00:49:6  20:00:00:05:30:00:49:
              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
```

10.1.2 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

**Inspur-Cisco Networking
Technology Co.,Ltd.**

does not, then the request is rejected.

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

Registering Name Server Proxies

You can register the name server proxy.

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

DETAILED STEPS

	Command or Action	Purpose

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

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

Rejecting Duplicate pWWNs

By FC standard, INOS 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, INOS 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:

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

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**fcns reject-duplicate-pwwn vsan** *vsan-id*
	
	Example:
	`switch(config)# fcns reject-duplicate-pwwn`
	` vsan 100`
		Any future flogi (with duplicate pwwn) on different switch, will be rejected and earlier FLOGI retained (default).

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

Name Server Database Entries

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

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

Displaying Name Server Database Entries

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

```
switch# show fcns database
-------------------------------------------------------------
-------- FCID                         TYPE  PWWN       (VENDOR)
        FC4-TYPE:FEATURE
-----------------------------------------------------------------------
0x01000  N      50:06:0b:00:00:10:a7:80                  scsi-fcp fc-
0                                                        ~~
0x01000  N      10:00:00:05:30:00:24:63 (Inspur)         ipfc
1
0x01000  N      50:06:04:82:c3:a0:98:52 (Company 1)      scsi-fcp 250
2
0x01010  N      21:00:00:e0:8b:02:99:36 (Company A)      scsi-fcp
0
0x02000   N     21:00:00:e0:8b:08:4b:20 (Company A)
0               10:00:00:05:30:00:24:23 (Inspur)         ipfc
0x02010   N
0x02020  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
-----------------------------------------------------------------------
0x03000   N     10:00:00:05:30:00:25:a3 (Inspur)         ipfc
```

```
0x03010    N    10:00:00:00:77:99:60:2c
0x03020    N    10:00:00:49:c9:28:c7:01
0xec000    N    21:00:00:20:37:a6:be:14 (Seagate)       scsi-
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
```

10.1.3 FDMI

Inspur 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.

10.1.4 Displaying FDMI

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

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

10.1.5 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.

About RSCN Information

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

Configuring the Port-Address Format

The zone server on Inspur SAN switches allows you to switch between the fabric-address format and port-address format. You can configure this feature on a per VSAN basis. This configuration remains unchanged even after an In-Service Software Upgrade (ISSU) or a switchover. By default, the Registered State Change Notification (RSCN) format is fabric address

You can configure the port-address format.

SUMMARY STEPS
1. **configure terminal**
2. **zone rscn address-format port vsan** *vsan-id*
3. **no zone rscn address-format port vsan** *vsan-id*
4. **show zone status vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**zone rscn address-format port vsan** *vsan-id*
	
	Example:
	`switch(config)# zone rscn address-format port vsan`
	`10`
		Switches to the port-address format.

Step 3	**no zone rscn address-format port vsan** *vsan-id*
	
	Example:
	`switch(config)# no zone rscn address-format port`
	`vsan 10`
		Reverts to the default format that is fabric-address format.

Step 4	**show zone status vsan** *vsan-id* **Example:** `switch(config)# show zone status vsan 10`	Displays the active the RSCN address format.

Displaying RSCN Information

The following example shows how to display registered device information:

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

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).

Configuring the multi-pid Option

You can configure the multi-pid option.

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

DETAILED STEPS

	Command or Action	Purpose

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

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

Suppressing Domain Format SW-RSCNs

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

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

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

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**rscn suppress domain-swrscn vsan** *vsan-id* **Example:** `switch(config)# rscn suppress domain-swrscn vsan`	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.

You can configure the RSCN timer.

SUMMARY STEPS
1. **configure terminal**
2. **rscn distribute**
3. **rscn event-tov** *timeout* **vsan** *vsan-id*
4. **no rscn event-tov** *timeout* **vsan** *vsan-id*
5. **rscn commit vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

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

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

Verifying the RSCN Timer Configuration

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

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

RSCN Timer Configuration Distribution

Because the timeout value for each switch is configured manually, a misconfiguration occurs when different switches time out at different times. Different N-ports in a network can receive RSCNs at different times.

Inspur 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 Inspur 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.

Caution   Only the RSCN timer configuration is distributed.

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

Enabling RSCN Timer Configuration Distribution
You can enable RSCN timer configuration distribution.

SUMMARY STEPS
1. **configure terminal**
2. **rscn distribute**
3. **no rscn distribute**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**rscn distribute** **Example:** `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.

SUMMARY STEPS
1. **configure terminal**
2. **rscn commit vsan** *timeout*

DETAILED STEPS

	Command or Action	Purpose

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

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

Discarding the RSCN Timer Configuration Changes

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

You can discard RSCN timer configuration changes.

SUMMARY STEPS
1. **configure terminal**
2. **rscn abort vsan** *timeout*

DETAILED STEPS

	Command or Action	Purpose

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

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

Clearing a Locked Session

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

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

To use administrative privileges and release a locked 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
```

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

10.1.6 Default Settings for RSCN

The following table lists the default settings for RSCN.

Table 18 : Default RSCN Settings

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

CHAPTER 11 Configuring iSCSI TLV

This chapter contains the following sections:
  • Overview of iSCSI TLV
  • iSCSI TLV and FCoE TLV Configuration

11.1 Overview of iSCSI TLV

iSCSI is an IP-based storage networking standard for linking data storage facilities. IP and Small Computer System Interface over IP (iSCSI) storage refers to block access of storage disks across devices connected using traditional Ethernet and TCP/IP networks. iSCSI protocol enables the transport of Small Computers Systems Interface (SCSI) commands over TCP/IP networks. By transmitting SCSI commands over IP networks, iSCSI facilitates block-level transfers over the Internet.

Ethernet networks are highly susceptible to broadcast and multicast storms, leading to congested networks. Data Center Bridging (DCB) extends lossless capabilities to Ethernet networks, thereby providing an option to define networks that are well suited for storage traffic. When used with traditional Ethernet networks to provide lossless iSCSI networks, DCB provides the following features:

  • Priority Flow Control—Priority Flow Control enables eight virtual queues on a single wire and helps send pause frames to a single type of traffic instead of all the traffic on the wire. This feature helps prevent head-of-the-line blocking while maintaining lossless capabilities. In consolidated networks, where a mix of traffic is sent on the same wire, PFC helps prioritize the traffic and assign it to either the drop or no-drop class. PFC is useful for iSCSI networks when they are designed for lossless, oversubscribed networks.

  • Enhanced Transmission Selection (ETS)—ETS provides the capability to allocate bandwidth to each traffic class on the same wire. ETS also helps prioritize and optimize the throughput for iSCSI and IP storage traffic, both of which shares a medium with the other traffic on the same link. The guaranteed bandwidth also aids in performance calculations to tune applications during peak traffic.

  • Data Center Bridging eXchange (DCBX) protocol—DCBX protocol is used to exchange all the DCB features across the devices and maintain consistency. DCBX protocol helps ensure consistent quality-of-service (QoS) parameters across the network and servers. The features are advertised to the servers in type-length-value (TLV) format using the Link Layer Discovery Protocol (LLDP). iSCSI TLV can be used to separate iSCSI traffic from other traffic. It can also be used to extend lossless behavior to the Ethernet infrastructure.

The iSCSI TLV over Data Center Bridging eXchange (DCBX) protocol feature lowers the cost of the lossless Ethernet deployment solution. iSCSI targets that can perform end-to-end iSCSI with initiators, are present.

DCBX negotiates the configuration and settings between the switch and the adapter through a variety of TLV and sub-TLVs. This allows the switch to distribute configuration values to all the attached adapters from a centralized location instead of having to manually program the class of service (CoS) markings on each individual server and adapter. For flexibility, ETS and Priority Flow Control parameters are coded in the TLV format. However, the use of Priority Flow Contro or ETS for lossless protocol behavior is not a requirement for iSCSI TLV operations. TLV can be leveraged for both traditional TCP or drop-behavior iSCSI networks as well as for a complete end-to-end lossless iSCSI fabric. Enabling ETS and Priority Flow Control separates storage traffic from other IP traffic and enables accurate and error-free configuration information to be transmitted from the switch to the adapter.

11.1.1 Guidelines and Limitations
  • We recommend that you use iSCSI TLV-supported CNA on both, the host and the target.
  • We recommend that you use CoS 4 for iSCSI traffic and create a custom 6e QoS policy based on the *default-nq-6e-policy* template, since most of the CNAs use CoS 4 for iSCSI.

11.2 iSCSI TLV and FCoE TLV Configuration

11.2.1 Identifying iSCSI and FCoE Traffic

To identify iSCSI, define a class map for iSCSI traffic. If a packet matches the iSCSI criteria configured for the corresponding class map using the **match** command, this class map is applied to the packet. If no execution strategy is specified (**match-any** or **match-all**), the default value of **match-any** is applied to the iSCSI traffic class.

Step 1	Enter global configuration mode: switch# **configure terminal**
Step 2	Create a named object that represents a class of traffic, and enter class-map mode. (Class-map names can contain alphabetic, hyphen, or underscore characters, are case sensitive, and can be up to 40 characters). switch(config)# **class-map type network-qos match-any** {*class-map-name*}
Step 3	Exit class-map configuration mode and enter global configuration mode: switch(config-cmap-qos)# **exit**
Step 4	Create a class map, provide conditions for applying this class map to a packet, and enter class-map configuration mode: switch(config)# **class-map type network-qos match-any** {*class-map-name*}
Step 5	Specify the CoS value to match and specify which protocol has to be mapped to a given CoS value: switch(config-cmap-qos)# **match protocol [fcoe	iscsi]**

Note
To enable TLV enter iscsi as the match protocol.

Step 6	Specify the CoS value as CoS 4 to match. The range is from 0 to 7: switch(config-cmap-qos)# **match cos** *cos value*
Step 7	Exit class-map configuration mode and enter global configuration mode: switch(config-cmap-qos)# **exit**

This example shows how to identify iSCSI traffic. Replace the placeholders with relevant values for your setup.

```
switch# configure terminal
switch(config)# class-map type network-qos match-any class-fcoe
switch(config-cmap-nqos)# exit
switch(config)# class-map type network-qos
match-any c1 switch(config-cmap-nqos)# match
protocol iscsi switch(config-cmap-nqos)# match
cos 4
switch(config-cmap-nqos)# exit
```

11.2.2 Configuring iSCSI Network QoS Policies

You can configure a network QoS policy by following one of these methods:

• Copy the predefined templates—You can copy a network QoS policy template and modify it as needed. Copying a network QoS policy trims the default policy name by stripping the default and policy substrings from it.

• Create a user-defined policy—You can create a network QoS policy that conforms to the **default-nq-6e-policy** template.

Copying a Predefined Network QoS Policy Template
Copy a predefined network QoS policy template by performing this procedure:

Step 1   Enter global configuration mode:
 switch# configure terminal

Step 2   Copy a predefined 6e network QoS policy and add a suffix or prefix to its name. (A prefix or suffix
 name can contain alphabetic, hyphen(-), or underscore (_) characters, is case sensitive, and can be up
 to 40 characters):
 switch(config)# **qos copy policy-map type network-qos default-nq-6e-4q8q-policy {prefix** *prefix*
 | suffix *suffix***}**

Step 3   Display the network QoS policy map type:
 switch(config)# **show policy-map type network-qos** [*my_template*]

11.2.3 Configuring a No-Drop Policy Map

Step 1   Enter global configuration mode:
 switch# **configure terminal**

Step 2   Remove a class map reference from the network QoS policy map type:
 switch(config-pmap-nqos)# **no class type network-qos** [*my_template]*

Step 3   Configure the class map of the network QoS type and specify the class map name:
 switch(config-pmap-nqos)# **class type network-qos** [*my_template]*

Step 4   Specify no drop. Default is no pause.
 switch(config-pmap-nqos-c)# **pause**

Step 5   Specify the MTU or the payload length. The range is from 576 to 9216.
 switch(config-pmap-nqos-c)# **mtu** [*mtu_size*]

This example shows the policy map configuration for the *default-nq-6e-policy* template. Replace the placeholders
with relevant values for your setup.

```
switch(config-pmap-nqos-c)# show policy-map type network-qos iscsi-
nodropnq-6e-4q8q policy-map type network-qos iscsi-nodropnq-6e-4q8q
template 6e-4q8q
    class type network-qos c-nq-6e-
      4q8q-drop match cos 0-2,5-7
      congestion-control tail-drop threshold burst-
      optimized mtu 1500
    class type network-qos c-nq-6e-4q8q-
      ndrop-fcoe match cos 3
      match
      protocol
      fcoe pause

      mtu 2112
    class type network-qos c-nq-6e-4q8q-
      ndrop-iscsi match protocol iscsi
    match cos 4 pause

      mtu 1500
```
This example shows how you can remove an old class map and add the newly created iSCSI class map:
```
switch# configure terminal
switch(config)# policy-map type network-qos iscsi-nq-6e-4q8q
```

```
switch(config-pmap-nqos)# no class type network-qos c-nq-6e-
4q8q-ndrop
switch(config-pmap-nqos)# class type network-qos c-nqc-nq-4e-4q8q-
ndrop-iscsi switch(config-pmap-nqos-c)#
```

11.2.4 Applying System Service Policies

Across VDCs, you can apply a network QoS policy on a system only globally. Applying a network QoS policy is applicable to the corresponding queuing policies.

To apply a network QoS policy to a target, use the **service-policy** command.

Step 1   Enter global configuration mode:
switch# **configure terminal**

Step 2   Enter system qos mode:
switch(config)# **system qos**

Step 3   Add the policy map to the input or output packets of system:
switch(config-sys-qos)# **service-policy type network-qos** {*my_template*}

Step 4   Exit config-sys-qos mode and enter the configuration mode:
switch(config-sys-qos)#**exit**

This example shows you how to apply system service policies. Replace the placeholders with relevant values for your setup.

```
switch# configure
terminal
switch(config)#
system qos
switch(config-sys-qos)# service-policy type network-qos iscsi-
nodropnq-6e-4q8q switch(config-sys-qos)# exit
```

CHAPTER12 Advanced Fibre Channel Features

This chapter describes how to configure advanced Fibre Channel features. This chapter includes the following sections:

   • Advanced Fibre Channel Features and Concepts

12.1 Advanced Fibre Channel Features and Concepts

12.1.1 Fibre Channel Timeout Values

You can modify Fibre Channel protocol-related timer values for the switch by configuring the following timeout values (TOVs):

   • Distributed services TOV (D_S_TOV)—The valid range is from 5,000 to 10,000 milliseconds. The default is 5,000 milliseconds.

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

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

Timer Configuration Across All VSANs

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

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

You can configure Fibre Channel timers across all VSANs.

SUMMARY STEPS
1.   **configure terminal**
2.   **fctimer R_A_TOV** *timeout*

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**fctimer R_A_TOV** *timeout* **Example:** `switch(config)# fctimer R_A_TOV 800`	Configures the R_A_TOV timeout value for all VSANs. The unit is milliseconds. This type of configuration is not permitted unless all VSANs are suspended.

Timer Configuration Per-VSAN

You can also issue the fctimer for a specified VSAN to configure different TOV values for VSANs with special links such as Fibre Channel. You can configure different E_D_TOV, R_A_TOV, and D_S_TOV values for individual VSANs. Active VSANs are suspended and activated when their timer values are changed.

You can configure per-VSAN Fibre Channel timers.

SUMMARY STEPS
1. **configure terminal**
2. **fctimer D_S_TOV** *timeout* **vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**fctimer D_S_TOV** *timeout* **vsan** *vsan-id*
	
	Example:
	`switch(config#)# fctimer D_S_TOV 900 vsan`
	`15`
		Configures the D_S_TOV timeout value (in milliseconds) for the specified VSAN. Suspends the VSAN temporarily. You have the option to end this command, if required.

EXAMPLES

This example shows how to configure the timer value for VSAN 2:

```
switch(config#)# fctimer D_S_TOV 6000 vsan 2
Warning: The vsan will be temporarily suspended when updating the timer
value This configuration would impact whole fabric. Do you want to
continue? (y/n) y
Since this configuration is not propagated to other switches, please configure
the same value in all the switches
```

fctimer Distribution

You can enable per-VSAN fctimer fabric distribution for all Inspur SAN switches in the fabric. When you perform fctimer configurations, and distribution is enabled, that configuration is distributed to all the switches in the fabric.

You automatically acquire a fabric-wide lock when you enter the first configuration command after you enabled distribution in a switch. The fctimer application uses the effective and pending database model to store or commit the commands based on your configuration.

Enabling or Disabling fctimer Distribution

You can enable or disable fctimer fabric distribution.

SUMMARY STEPS
1. **configure terminal**
2. **fctimer distribute**
3. **no fctimer distribute**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**fctimer distribute** **Example:** `switch(config)# fctimer distribute`	Enables fctimer configuration distribution to all switches in the fabric. Acquires a fabric lock and stores all future configuration changes in the pending database.
Step 3	**no fctimer distribute** **Example:** `switch(config)# no fctimer distribute`	Disables (default) fctimer configuration distribution to all switches in the fabric.

Committing fctimer Changes

When you commit the fctimer configuration changes, the effective database is overwritten by the configuration changes in the pending database and all the switches in the fabric receive the same configuration. When you commit the fctimer configuration changes without implementing the session feature, the fctimer configurations are distributed to all the switches in the physical fabric.

SUMMARY STEPS
1.   **configure terminal**
2.   **fctimer commit**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**fctimer commit** **Example:** `switch(config)# fctimer commit`	Distributes the fctimer configuration changes to all switches in the fabric and releases the lock. Overwrites the effective database with the changes made to the pending database.

Discarding fctimer Changes

After making the configuration changes, you can choose to discard the changes by discarding the changes instead of committing them. In either case, the lock is released.

SUMMARY STEPS
1.   **configure terminal**
2.   **fctimer abort**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**fctimer abort** **Example:** ```switch(config)# fctimer abort```	Discards the fctimer configuration changes in the pending database and releases the fabric lock.

Overriding the Fabric Lock

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

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

To use administrative privileges and release a locked fctimer session, use the **clear fctimer session** command.

```
switch# clear fctimer session
```

Fabric Database Merge Guidelines

When merging two fabrics, follow these guidelines:
 • Be aware of the following merge conditions:
 ◦The merge protocol is not implemented for distribution of the fctimer values. You must manually merge the fctimer values when a fabric is merged.
 ◦The per-VSAN fctimer configuration is distributed in the physical fabric.
 ◦The fctimer configuration is only applied to those switches containing the VSAN with a modified fctimer value.
 ◦The global fctimer values are not distributed.
 • Do not configure global timer values when distribution is enabled.

Verifying Configured fctimer Values

Use the **show** fctimer command to display the configured fctimer values. The following example displays the configured global TOVs:

```
switch# show fctimer
F_S_TOV    D_S_TOV    E_D_TOV    R_A_TOV
-------------------------------------
5000 ms    5000 ms    2000 ms    10000 ms
```

The following example displays the configured TOV for VSAN 10:

```
switch# show fctimer vsan 10
vsan no.  F_S_TOV    D_S_TOV    E_D_TOV    R_A_TOV
-------------------------------------
-------- 10          5000 ms    5000 ms
          3000 ms    10000 ms
```

12.1.2 World Wide Names

The world wide name (WWN) in the switch is equivalent to the Ethernet MAC address. As with the MAC address, you must uniquely associate the WWN to a single device. The principal switch selection and the allocation of domain IDs rely on the WWN.

Inspur SAN switches support three network address authority (NAA) address formats. (see the following table).

Table 19 : Standardized NAA WWN Formats

NAA Address	NAA Type	WWN Format	
IEEE 48-bit address	Type 1 = 0001b	000 0000 0000b	48-bit MAC address
IEEE extended	Type 2 = 0010b	Locally assigned	48-bit MAC address
IEEE registered	Type 5 = 0101b	IEEE company ID: 24 bits	VSID: 36 bits

Caution   Changes to the world-wide names should be made by an administrator or individual who is completely familiar with switch operations.

Verifying the WWN Configuration

Use the **show wwn** commands to display the status of the WWN configuration. This example shows how to display the status of all WWNs:

```
switch# show wwn status
Type    Configured    Available    Resvd.  Alarm State
----    ----------    --------------  ------  ---
       -------- 1              64            48
   ( 75%)    16        NONE
 2,5       524288    442368 ( 84%)   73728    NONE
```

This example shows how to display the information for block ID 51:

```
switch# show wwn status block-id 51
WWNs in this block: 21:00:ac:16:5e:52:00:03 to
21:ff:ac:16:5e:52:00:03 Num. of WWNs:: Configured: 256
Allocated:                                0 Available: 256

Block Allocation Status: FREE
```

This example shows how to display the WWN for a specific switch:

```
switch# show wwn switch
Switch WWN is 20:00:ac:16:5e:52:00:00
```

Link Initialization WWN Usage

Exchange Link Protocol (ELP) and Exchange Fabric Protocol (EFP) use WWNs during link initialization. ELPs and EFPs both use the VSAN WWN by default during link initialization. However, the ELP usage changes based on the peer switch's usage:

· If the peer switch ELP uses the switch WWN, then the local switch also uses the switch WWN.
· If the peer switch ELP uses the VSAN WWN, then the local switch also uses the VSAN WWN.

Configuring a Secondary MAC Address
You can allocate secondary MAC addresses.

SUMMARY STEPS
1. **configure terminal**
2. **wwn secondary-mac** *wwn-id* **range** *value*

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**wwn secondary-mac** *wwn-id* **range** *value*
	
	Example:
	`switch(config)# wwn secondary-mac`
	`33:e8:00:05:30:00:16:df`
	`range 55`
		Configures the secondary MAC address.
		This command cannot be undone.

EXAMPLES
This example shows how to configure the secondary MAC address:

```
switch(config)# wwn secondary-mac 00:99:55:77:55:55 range 64
This command CANNOT be undone.
Please enter the BASE MAC ADDRESS again: 00:99:55:77:55:55
Please enter the mac address RANGE again: 64
From now on WWN allocation would be based on new MACs. Are you sure? (yes/no)
no
You entered: no. Secondary MAC NOT programmed
```

12.1.3 FC ID Allocation for HBAs

Fibre Channel standards require a unique FC ID to be allocated to an N port attached to an F port in any switch. To conserve the number of FC IDs used, Inspur SAN switches use a special allocation scheme.

Some HBAs do not discover targets that have FC IDs with the same domain and area. The switch software maintains a list of tested company IDs that do not exhibit this behavior. These HBAs are allocated with single FC IDs. If the HBA can discover targets within the same domain and area, a full area is allocated.

To allow further scalability for switches with numerous ports, the switch software maintains a list of HBAs that can discover targets within the same domain and area. Each HBA is identified by its company ID (also known as Organizational Unique Identifier, or OUI) used in the pWWN during a fabric log in. A full area is allocated to the N ports with company IDs that are listed and for the others, a single FC ID is allocated.

Regardless of the type (whole area or single) of FC ID allocated, the FC ID entries remain persistent.

Default Company ID List

All Inspur SAN switches contain a default list of company IDs that require area allocation. Using the company ID reduces the number of configured persistent FC ID entries. You can configure or modify these entries using the CLI.

Caution   Persistent entries take precedence over company ID configuration. If the HBA fails to discover a target, verify that the HBA and the target are connected to the same switch and have the same area in their FC IDs, then perform the following procedure:

1. Shut down the port connected to the HBA.
2. Clear the persistent FC ID entry.
3. Get the company ID from the port WWN.
4. Add the company ID to the list that requires area allocation.
5. Bring up the port.

The list of company IDs have the following characteristics:

• A persistent FC ID configuration always takes precedence over the list of company IDs. Even if the company ID is configured to receive an area, the persistent FC ID configuration results in the allocation of a single FC ID.

• New company IDs added to subsequent releases are automatically added to existing company IDs.

• The list of company IDs is saved as part of the running and saved configuration.

• The list of company IDs is used only when the fcinterop FC ID allocation scheme is in auto mode. By default, the interop FC ID allocation is set to auto, unless changed.

Tip   We recommend that you set the fcinterop FC ID allocation scheme to auto and use the company ID list and persistent FC ID configuration to manipulate the FC ID device allocation.

• Use the **fcinterop FCID allocation auto** command to change the FC ID allocation and the **show running-config** command to view the currently allocated mode.

• When you enter **a write erase**, the list inherits the default list of company IDs shipped with a relevant release.

Verifying the Company ID Configuration

You can view the configured company IDs by entering the show fcid-allocation area command. Default entries are listed first and the user-added entries are listed next. Entries are listed even if they were part of the default list and you later removed them.

This example shows how to display the list of default and configured company IDs:

```
switch# show fcid-allocation area
FCID area allocation company id
info: 00:50:2E <--------------
Default entry 00:50:8B
00:60:B0
00:A0:B8
00:E0:69
00:30:AE + <------------ User-added entry
00:32:23 +
00:E0:8B * <------------ Explicitly deleted entry (from the original default
list)
```

```
                 Total company ids: 7
                 + - Additional user configured company ids.
                 * - Explicitly deleted company ids from default list.
```

You can implicitly derive the default entries shipped with a specific release by combining the list of Company IDs displayed without any identification with the list of deleted entries.

You can also view or obtain the company IDs in a specific WWN by entering the show fcid-allocation company-id-from-wwn command. Some WWN formats do not support company IDs. In these cases, you many need to configure the FC ID persistent entry.

This example shows how to display the company ID for the specified WWN:

```
                 switch# show fcid-allocation company-id-from-wwn 20:00:00:05:30:00:21:60
                 Extracted Company ID: 0x000530
```

12.1.4 Switch Interoperability

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

Not all vendors follow the standards in the same way, which results in the need for interoperability modes. This section briefly explains the basic concepts of these modes.

Each vendor has a regular mode and an equivalent interoperability mode, which specifically turns off advanced or proprietary features and provides the product with a standards-compliant implementation.

About Interop Mode

The software supports the following four interop modes:
 • Mode 1— Standards-based interop mode that requires all other vendors in the fabric to be in interop mode.
 • Mode 2—Brocade native mode (Core PID 0).
 • Mode 3—Brocade native mode (Core PID 1).
 • Mode 4—McData native mode.

The following table lists the changes in switch operation when you enable interoperability mode.

Table 20 : Changes in Switch Operation When Interoperability Is Enabled

Switch Feature	Changes if Interoperability Is Enabled
Domain IDs	Some vendors cannot use the full range of 239 domains within a fabric. Domain IDs are restricted to the range 97 to 127, to accommodate McData's nominal restriction to this same range. Domain IDs can either be static or preferred, which operate as follows: • Static: Inspur switches accept only one domain ID; if a switch does not get that domain ID it isolates itself from the fabric. • Preferred: If the switch does not get its requested domain ID, it accepts any assigned domain ID.

Timers	All Fibre Channel timers must be the same on all switches as these values are exchanged by E ports when establishing an ISL. The timers are F_S_TOV, D_S_TOV, E_D_TOV, and R_A_TOV.
F_S_TOV	Verify that the Fabric Stability Time Out Value timers match exactly.
D_S_TOV	Verify that the Distributed Services Time Out Value timers match exactly.
E_D_TOV	Verify that the Error Detect Time Out Value timers match exactly.
R_A_TOV	Verify that the Resource Allocation Time Out Value timers match exactly.
Trunking	Trunking is not supported between two different vendor's switches. This feature may be disabled per port or per switch.
Default zone	The default zone operation of permit (all nodes can see all other nodes) or deny (all nodes are isolated when not explicitly placed in a zone) may change.
Zoning attributes	Zones may be limited to the pWWN and other proprietary zoning methods (physical port number) may be eliminated. **Note** On a Brocade switch, use the **cfgsave** command to save fabric-wide zoning configuration. This command does not have any effect on Inspur SAN switches if they are part of the same fabric. You must explicitly save the configuration on each Inspur SAN switch.
Zone propagation	Some vendors do not pass the full zone configuration to other switches, only the active zone set gets passed. Verify that the active zone set or zone configuration has correctly propagated to the other switches in the fabric.
VSAN	Interop mode only affects the specified VSAN.
TE ports and SAN port channels	TE ports and SAN port channels cannot be used to connect Inspur switches to non-Inspur SAN switches. Only E ports can be used to connect to non-Inspur SAN switches. TE ports and SAN port channels can still be used to connect a Inspur switch to other Inspur SAN switches even when in interop mode.

FSPF	The routing of frames within the fabric is not changed by the introduction of interop mode. The switch continues to use src-id, dst-id, and ox-id to load balance across multiple ISL links.
Domain reconfiguration disruptive	This is a switch-wide impacting event. Brocade and McData require the entire switch to be placed in offline mode and/or rebooted when changing domain IDs.
Domain reconfiguration nondisruptive	This event is limited to the affected VSAN. Inspur SAN switches have the capability to restart only the domain manager process for the affected VSAN and not the entire switch.
Name server	Verify that all vendors have the correct values in their respective name server database.

Configuring Interop Mode 1

You can interop mode1 in Inspur SAN switches disruptively or nondisruptively.

DETAILED STEPS

	Command or Action	Purpose
Step 1	Place the VSAN of the E ports that connect to the OEM switch in interoperability mode.	`switch# ` **`configuration terminal`** `switch(config)# ` **`vsan database`** `switch(config-vsan-db)# ` **`vsan 1`** **`interop 1`** `switch(config-vsan-db)# ` **`exit`**
Step 2	Assign a domain ID in the range of 97 (0x61) through 127 (0x7F).	**Note** This is an limitation imposed by the McData switches. In Inspur SAN switches, the default is to request an ID from the principal switch. If the preferred option is used, Inspur SAN switches request a specific ID, but still join the fabric if the principal switch assigns a different ID. If the static option is used, the Inspur SAN switches do not join the fabric unless the principal switch agrees and assigns the requested ID. **Note** When changing the domain ID, the FC IDs assigned to N ports also change.

Step 3	Change the Fibre Channel timers (if they have been changed from the system defaults).	**Note**	The Inspur SAN switches, Brocade, and McData FC Error Detect (ED_TOV) and Resource Allocation (RA_TOV) timers default to the same values. They can be changed if needed. The RA_TOV default is 10 seconds, and the ED_TOV default is 2 seconds. Per the FC-SW2 standard, these values must be the same on each switch within the fabric.

```
switch(config)# fctimer e_d_tov ?
  <1000-100000> E_D_TOV in milliseconds(1000-100000)


switch(config)# fctimer r_a_tov ?
  <5000-100000> R_A_TOV in milliseconds(5000-100000)
```

Step 4	When making changes to the domain, you may or may not need to restart the Domain Manager function for the altered VSAN.	• Force a fabric reconfiguration with the **disruptive** option.
		`switch(config)# fcdomain restart disruptive vsan 1`
		
		or
		
		• Do not force a fabric reconfiguration.
		`switch(config# fcdomain restart vsan 1`

12.1.5 Default Settings for Advanced Fibre Channel Features

The following table lists the default settings for the features included in this chapter.

Table 21: Default Settings for Advanced Features

Parameters	Default
CIM server	Disabled
CIM server security protocol	HTTP
D_S_TOV	5,000 milliseconds
E_D_TOV	2,000 milliseconds
R_A_TOV	10,000 milliseconds
Timeout period to invoke fctrace	5 seconds
Number of frame sent by the fcping feature	5 frames
Remote capture connection protocol	TCP
Remote capture connection mode	Passive
Local capture frame limits	10 frames
FC ID allocation mode	Auto mode
Loop monitoring	Disabled

Interop mode	Disabled

CHAPTER 13 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

13.1 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 Inspur SAN switches and other devices. DHCHAP consists of the CHAP protocol combined with the Diffie-Hellman exchange.

13.1.1 Fabric Authentication

All Inspur 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.

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

Figure 22 : Switch and Host Authentication

Inspur-Cisco Networking
Technology Co.,Ltd.

13.1.2 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.

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 Inspur INOS 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 Inspur SAN switches.

About Enabling DHCHAP

By default, the DHCHAP feature is disabled in all Inspur 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

To enable DHCHAP, perform this task:

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **feature fcsp**	Enables the DHCHAP in this switch.
Step 3	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.

The following table identifies switch-to-switch authentication between two Inspur switches in various modes.

Table 22 : 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

To configure the DHCHAP mode for a particular interface, perform this task:

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **interface vfc** *if-number* - *if-number*	Selects a range of interfaces and enters the interface configuration mode.
Step 3	switch(config-if)# **fcsp on**	Sets the DHCHAP mode for the selected interfaces to be in the on state.
Step 4	switch(config-if)# **fcsp off**	Reverts to the factory default of auto-passive for these three interfaces.
Step 5	switch(config-if)# **fcsp auto-active** *timeout-period*	Changes the DHCHAP authentication mode to auto-active for the selected interfaces. The timeout period value (in minutes) sets how often reauthentication occurs after the initial authentication. Zero (0) indicates that the port does not perform reauthentication.
Step 6	switch(config-if)# **fcsp auto-passive**	Changes the DHCHAP authentication mode to auto-passive for the selected interfaces. Reauthentication is disabled (default).

DHCHAP Hash Algorithm

Inspur 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.

DETAILED STEPS

	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.
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 Inspur 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.

DETAILED STEPS

	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.

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 Inspur DCNM for SAN to manage the password database.

Configuring DHCHAP Passwords for the Local Switch

You can configure the DHCHAP password for the local switch.

DETAILED STEPS

	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).

Configuring DHCHAP Passwords for Remote Devices

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

DETAILED STEPS

	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`
		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.

DETAILED STEPS

	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.

13.1.3 Configuration Examples for Fabric Security

This example shows how to set up authentication:

Step 1   Obtain the device name of the Inspur SAN switch in the fabric. The Inspur 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)# feature fcsp
```

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 fc2/4
switch(config-if)# fcsp on
```

Step 6   Verify the protocol security information configured in this switch by displaying the DHCHAP local password database.

     Example:
```
switch# show fcsp dhchap database
DHCHAP Local Password:
     Non-device specific
password:******* Other Devices'
Passwords:
```

```
      Password for device with WWN:20:00:00:05:30:00:38:5e is *******
```

Step 7 Display the DHCHAP configuration in the interface.
    Example:
```
switch# show fcsp interface fc2/4
fc2/4
       fcsp authentication
       mode:SEC_MODE_ON
       Status:Successfully
       authenticated
```
Step 8   Repeat these steps on the connecting switch.
    Example:
```
MDS-9509# show wwn switch
Switch WWN is 20:00:00:05:30:00:38:5e
MDS-9509(config)# fcsp enable
MDS-9509(config)# fcsp dhchap password rtp9509
MDS-9509(config)# fcsp dhchap devicename 20:00:00:05:30:00:54:de password rtp9216
MDS-9509(config)# interface fc 4/5
MDS-9509(config-if)# fcsp on
MDS-9509# show fcsp dhchap
database DHCHAP Local
Password:
       Non-device specific
password:******* Other Devices'
Passwords:
       Password for device with WWN:20:00:00:05:30:00:54:de
is ******* MDS-9509# show fcsp interface fc2/4
Fc2/4
       fcsp authentication
       mode:SEC_MODE_ON
       Status:Successfully
       authenticated
```
You have now enabled and configured DHCHAP authentication for the sample setup.

13.1.4 Default Settings for Fabric Security

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

Table 23 : 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 14 Configuring Port Security

This chapter describes how to configure port security. This chapter includes the following sections:
  • Configuring Port Security

14.1 Configuring Port Security

Inspur SAN switches provide port security features that reject intrusion attempts and report these intrusions to the administrator.

14.1.1 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.

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.

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.

14.1.2 Configuring Port Security

Configuring Port Security with Auto-Learning and CFS Distribution

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

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
Committing the Changes
Copying the Port Security Database
Disabling Auto-Learning
Enabling Port Security
Enabling Port Security Distribution

Configuring Port Security with Auto-Learning without CFS

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

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
Copying the Port Security Database
Disabling Auto-Learning
 Enabling Port Security

Configuring Port Security with Manual Database Configuration

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

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.

14.1.3 Enabling Port Security

You can enable port security.
By default, the port security feature is disabled.

SUMMARY STEPS

1. **configure terminal**
2. **feature port-security**

DETAILED STEPS

	Command or Action	Purpose

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

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

14.1.4 Port Security Activation

Activating Port Security

You can activate port security.

SUMMARY STEPS
1. **configure terminal**
2. **fc-port-security activate vsan** *vsan-id* **no-auto-learn**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**configure terminal** **Example:** ``` switch# configure terminal switch(config)# ```	Enters global configuration mode.
Step 2	**fc-port-security activate vsan** *vsan-id* **no-auto-learn** **Example:** ``` switch(config)# fc-port-security activate vsan 30 no-auto-learn ```	Activates the port security database for the specified VSAN. Use the **no-auto-learn** keyword to disable 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.

SUMMARY STEPS
1. **configure terminal**
2. **fc-port-security activate vsan** *vsan-id* **force**

DETAILED STEPS

	Command or Action	Purpose

Step 1	**configure terminal**	Enters global configuration mode.

	Example:
	`switch# configure terminal`
	`switch(config)#`	

Step 2	**fc-port-security activate vsan** *vsan-id* **force**	Forces the port security database to activate for the specified VSAN even if conflicts occur.

	Example:
	`switch(config)# fc-port-security activate vsan`
	`212 force`	

Database Reactivation

You can reactivate the port security database.

SUMMARY STEPS
1. **configure terminal**
2. **no fc-no port-security auto-learn vsan** *vsan-id*
3. **exit**
4. **fc-port-security database copy vsan** *vsan-id*
5. **configure terminal**
6. **fc-port-security activate vsan** *vsan-id*

DETAILED STEPS

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

	Example:
	`switch# configure terminal`
	`switch(config)#`	

Step 2	**no fc-no port-security auto-learn vsan** *vsan-id*	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.

	Example:
	`switch(config)# no fc-no port-security auto-`
	`learn vsan 35`	

Step 3	**exit**	Exits the configuration mode.
	Example: `switch(config)# exit`	
Step 4	**fc-port-security database copy vsan** *vsan-id*	Copies from the active to the configured database.

	Example:
	`switch# fc-port-security database copy vsan`
	`35`	

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

Step 6	**fc-port-security activate vsan** *vsan-id*
	
	Example:
	`switch(config)# fc-port-security activate`
	`vsan 35`
		Activates the port security database for the specified VSAN, and automatically enables auto-learning.

14.1.5 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.

SUMMARY STEPS
1. **configure terminal**
2. **fc-port-security auto-learn vsan** *vsan-id*

DETAILED STEPS

Command or Action	Purpose

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

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

SUMMARY STEPS

1. **configure terminal**
2. **no fc-port-security auto-learn vsan** *vsan-id*

DETAILED STEPS

	Command or Action	**Purpose**

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

Step 2	**no fc-port-security auto-learn vsan** *vsan-id*
	
	Example:
	`switch(config)# no fc-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 24 : 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 vfc 1 (F1).
- A pWWN (P2) is allowed access through interface vfc 2 (F1).
- A nWWN (N1) is allowed access through interface vfc 3 (F2).
- Any WWN is allowed access through interface vfc 4 (F3).
- A nWWN (N3) is allowed access through any interface.
- A pWWN (P3) is allowed access through interface vfc 5 (F4).
- A sWWN (S1) is allowed access through interface vfc 6 -8 (F10 to F13).
- A pWWN (P10) is allowed access through interface vfc 9 (F11).

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

Table 25 : 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.
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.

14.1.6 Port Security Manual Configuration

You can manually configure port security.

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, add those pairs to the port security database.

Tip   Remote switch binding can be specified at the local switch. To specify the remote interfaces, you can use either the fWWN or sWWN-interface combination.

To add authorized port pairs for port security, perform this task:

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configuration terminal**	Enters configuration mode.
Step 2	switch(config)# **fc-port-security database vsan** *vsan-id*	Enters the port security database mode for the specified VSAN.
Step 3	switch(config)# **no fc-port-security database vsan** *vsan-id*	Deletes the port security configuration database from the specified VSAN.
Step 4	switch(fc-config-port-security)# **swwn** *swwn-id* **interface san-port-channel 5**	Configures the specified sWWN to only log in through SAN port channel 5.
Step 5	switch(fc-config-port-security)# **any-wwn interface vfc** *if-number* - **vfc** *if-number*	Configures any WWN to log in through the specified interfaces.

This example enters the port security database mode for VSAN 2:

```
switch(config)# fc-port-security database vsan 2
```
This example configures the specified sWWN to only log in through SAN port channel 5:

```
switch(fc-config-port-security)#
swwn 20:01:33:11:00:2a:4a:66 interface san-port-channel 5
```
This example configures the specified pWWN to log in through the specified interface in the specified switch:

```
switch(fc-config-port-security)#
pwwn 20:11:33:11:00:2a:4a:66 swwn 20:00:00:0c:85:90:3e:80
 interface vfc 2
```
This example configures any WWN to log in through the specified interface in any switch:

```
switch(fc-config-port-security)# any-wwn interface vfc 2
```

14.1.7 Port Security Configuration Distribution

The port security feature uses the Inspur 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.

Enabling Port Security Distribution

You can enable port security distribution.

SUMMARY STEPS

**Inspur-Cisco Networking
Technology Co.,Ltd.**

1. **configure terminal**
2. **fc-port-security distribute**
3. **no fc-port-security distribute**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**fc-port-security distribute** **Example:** `switch(config)# fc-port-security distribute`	Enables distribution.
Step 3	**no fc-port-security distribute** **Example:** `switch(config)# no fc-port-security distribute`	Disables distribution.

Related Topics

Activation and Auto-Learning Configuration Distribution

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.

SUMMARY STEPS

1. **configure terminal**
2. **fc-port-security commit vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**fc-port-security commit vsan** *vsan-id* **Example:** `switch(config)# fc-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.

SUMMARY STEPS

1. **configure terminal**
2. **fc-port-security abort vsan vsan-id**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**fc-port-security abort vsan** *vsan-id* **Example:** `switch(config)# fc-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 26 : 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

	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 =
	3. You issue a commit.	Not applicable	configuration database = {A,B, E} active database = {A,B, E, C*, D*}
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
	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
	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.

[1] The * (asterisk) indicates learned entries.

14.1.8 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.

14.1.9 Database Interaction

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

Table 27 : 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	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	You can overwrite the configuration database with the active database.

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

Figure 23 : Port Security Database Scenarios

**Inspur-Cisco Networking
Technology Co.,Ltd.**
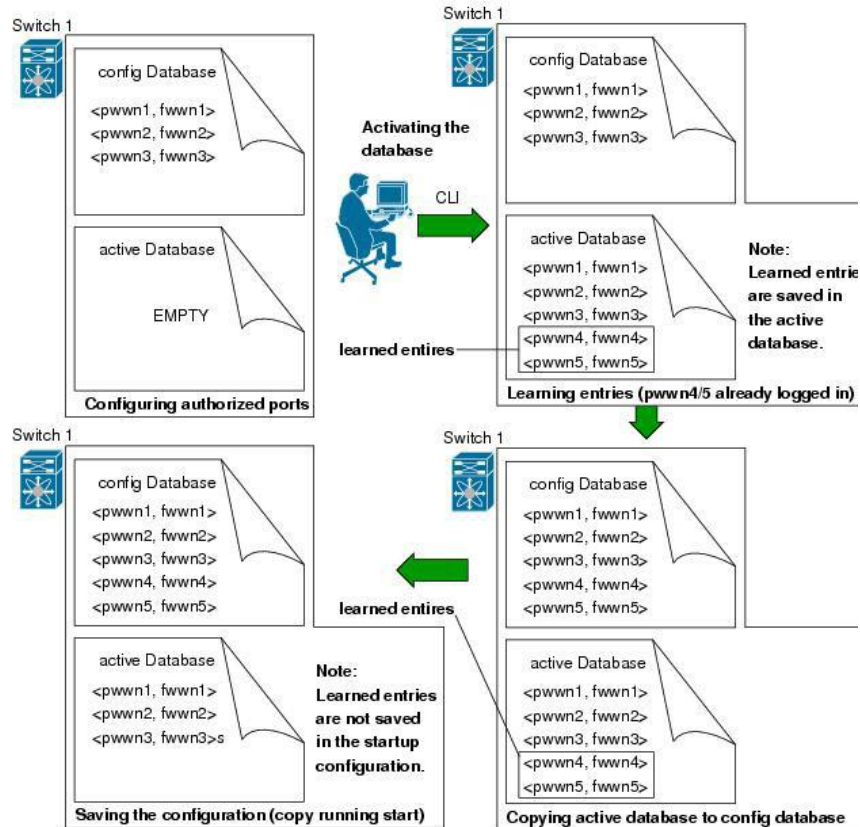
Database Scenarios

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

Figure 24 : Port Security Database Scenarios
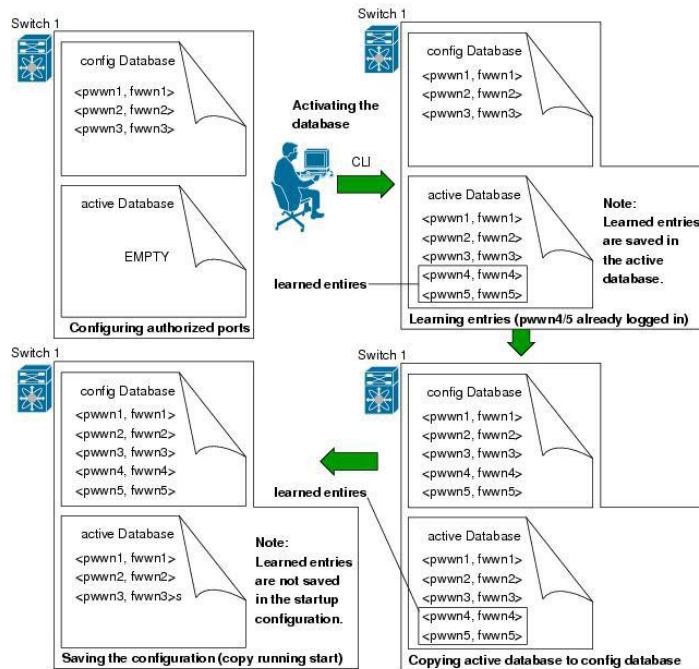
Copying the Port Security Database

| 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 **fc- 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# fc-port-security database copy vsan 1
```

Use the **fc- 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# fc-port-security database diff active vsan 1
```

Use the **fc-port-security database diff config vsan** command to obtain information on the differences between the configuration database and the active database:

```
switch# fc-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 **fc-port-security commit** command to actually delete the database.

Use the **no fc- port-security database vsan** command in configuration mode to delete the configured database for a specified VSAN:

```
switch(config)# no fc-port-security database vsan 1
```

Clearing the Port Security Database

Use the **clear fc-port-security statistics vsan** command to clear all existing statistics from the port security database for a specified VSAN.

```
switch# clear fc-port-security statistics vsan 1
```

Use the **clear fc-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 fc-port-security database auto-learn interface fc2/1 vsan 1
```

Use the **clear fc-port-security database auto-learn vsan** command to clear any learned entries in the active database for the entire VSAN.

```
switch# clear fc-port-security database auto-learn vsan 1
```

Use the **fc-port-security clear vsan** command to clear the pending session in the VSAN from any switch in the VSAN.

```
switch# clear fc-port-security session vsan 5
```

14.1.10 Verifying the Port Security Configuration

The **show fc-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 fc-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.

Command	Purpose
show fc-port-security database	Displays the port security configuration database.
show fc-port-security database vsan 1	Displays the port security configuration database for VSAN 1.
show fc-port-security database active	Displays the activated database.

show fc-port-security pending-diff vsan 1	Displays the difference between the temporary configuration database and the configuration database.
show fc-port-security database fwwn 20:01:00:05:30:00:95:de vsan 1	Displays the configured fWWN port security in VSAN 1.
show fc-port-security statistics	Displays the port security statistics.
show fc-port-security status	Displays the status of the active database and the auto-learning configuration.

14.1.11 Default Settings for Port Security

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

Table 28 : 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 15 Configuring Fabric Binding

This chapter describes how to configure fabric binding. This chapter includes the following sections:
 • Configuring Fabric Binding

15.1 Configuring Fabric Binding

15.1.1 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 29 : 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.
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 Inspur 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.

15.1.2 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.

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.

SUMMARY STEPS
1. **configure terminal**
2. **feature fabric-binding**

DETAILED STEPS

	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:

SUMMARY STEPS
1. **configure terminal**
2. **fabric-binding database vsan** *vsan-id*
3. **no fabric-binding database vsan** *vsan-id*
4. **swwn** *swwn-id* **domain** *domain-id*
5. **no swwn** *swwn-id* **domain** *domain-id*

DETAILED STEPS

	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.

Activating Fabric Binding

You can activate the fabric binding feature.

SUMMARY STEPS
1. **configure terminal**
2. **fabric-binding activate vsan** *vsan-id*
3. **no fabric-binding activate vsan** *vsan-id*

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**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.

SUMMARY STEPS

1. **configure terminal**
2. **fabric-binding activate vsan** *vsan-id* **force**
3. **no fabric-binding activate vsan** *vsan-id* **force**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**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.

<div align="center">switch# fabric-binding database copy vsan 1</div>

• 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.

<div align="center">switch# fabric-binding database diff active vsan 1</div>

• Use the **fabric-binding database diff config vsan** command to obtain information on the differences between the config database and the active database.

<div align="center">switch# fabric-binding database diff config vsan 1</div>

• 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.

<div align="center">switch# copy running-config startup-config</div>

<div align="center">Inspur-Cisco Networking
Technology Co.,Ltd.</div>

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 th**e no fabric-binding** command in configuration mode to delete the configured database for a specified VSAN:

```
switch(config)# no fabric-binding database vsan 10
```

15.1.3 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      [2]      Domain mismatch
     2003
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
```
This example shows how to display EFMD Statistics for VSAN 4:

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

15.1.4 Default Settings for Fabric Binding

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

Table 30 : Default Fabric Binding Settings

Parameters	Default
Fabric binding	Disabled

CHAPTER 16 Configuring Port Tracking

This chapter describes how to configure port tracking. This chapter includes the following sections:
  • Configuring Port Tracking

16.1 Configuring Port Tracking

Inspur SAN switches offer the port tracking feature 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.

16.1.1 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 25 : Traffic Recovery Using Port Tracking
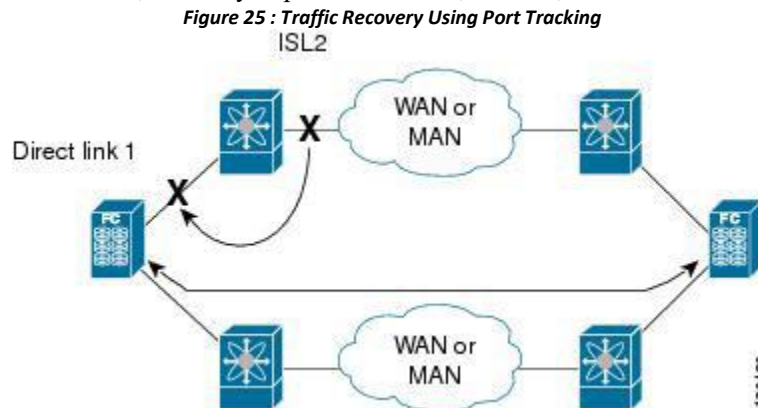
Port tracking monitors and detects failures that cause topology changes and brings down the links that connect the attached devices. When you enable this feature and explicitly configure the linked and tracked ports, the switch software monitors the tracked ports and alters the operational state of the 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. Virtual 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 virtual E or VE 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
   Fibre Channel Timeout Values

16.1.2 Guidelines and Limitations for Port Tracking

Port tracking has the following configuration guidelines and limitations:
   • Port tracking is supported only on virtual fibre channel (VFC) interfaces.
   • Verify that the tracked ports and the linked ports are on the same Inspur switch.
   • Be aware that the linked port is automatically brought down when the tracked port goes down.
   • Do not track a linked port back to itself (for example, Port fc2/2 to Port fc2/4 and back to Port fc2/2) to avoid recursive dependency.

16.1.3 Default Settings for Port Tracking

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

Table 31 : Default Port Tracking Parameters

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

16.1.4 Configuring Port Tracking

Before configuring port tracking, consider the following guidelines:
   • Verify that the tracked ports and the linked ports are on the same Inspur switch.
   • Be aware that the linked port is automatically brought down when the tracked port goes down.
   • Do not track a linked port back to itself (for example, Port fc2/2 to Port fc2/4 and back to Port fc2/2) to avoid recursive dependency.

Enabling the Port Tracking Feature

You can enable port tracking.

Before You Begin
You must enable the port track feature from the storage VDC.

SUMMARY STEPS
1. **configuration terminal**
2. **switchto vdc** *vdc-name*
3. **configuration terminal**
4. **feature port-track**

5.   (Optional) **show feature | port-track**
6.   (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**switchto vdc** *vdc-name*
	
	Example:
	`switch(config)# switchto vdc`
	`storage switch-storage#`
		Switches to the storage VDC. The *vdc-name* can be any case-sensitive, alphanumeric string up to 32 characters.

Step 3	**configuration terminal**
	
	Example:
	`switch-storage# configure terminal`
	`switch-storage(config)#`
		Enters configuration mode.

Step 4	**feature port-track** **Example:** `switch-storage(config)# feature port-track`	Enables the port tracking feature.
Step 5	**show feature	port-track** **Example:** `switch-storage(config)# show feature| port-track`	(Optional) Displays information about the port tracking feature.

Step 6	**copy running-config startup-config**
	
	Example:
	`switch-storage(config)# copy running-config`
	`startup-config`
		(Optional)
		Copies the running configuration to the startup configuration.

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.

Binding a Tracked Port

You can bind a track port head.

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.

Before You Begin
• Ensure you are in the storage VDC.

SUMMARY STEPS

1. **configure terminal**
2. **interface vfc** *if-number*
3. **port-track interface vfc** *if-number* | **san-port-channel** *port* | **vfc-port-channel** *port*

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**interface vfc** *if-number*
	
	Example:
	`switch(config)# interface vfc 2`
	`switch(config-if)#`
		Enters the interface configuration mode for the linked port. You can now configure the tracked ports.

Step 3	**port-track interface vfc** *if-number*	**san-port-channel** *port* | **vfc-port-channel** *port* **Example:** `switch(config-if)# port-track`	Specifies the tracked port. When the tracked port goes down, the linked port is also brought down.

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 26 : Traffic Recovery Using Port Tracking

Inspur-Cisco Networking
Technology Co.,Ltd.
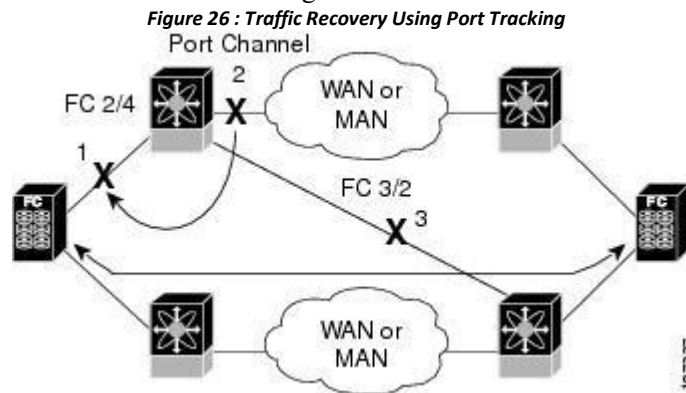
Monitoring Ports in a VSAN

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.

SUMMARY STEPS
1. **configure terminal**
2. **interface vfc** *if-number*
3. **port-track interface san-port-channel 1 vsan 2**
4. **no port-track interface san-port-channel 1 vsan 2**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**interface vfc** *if-number* **Example:** `switch(config)# interface vfc 3`	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.

SUMMARY STEPS

1. **configure terminal**
2. **interface vfc** *if-number*
3. **port-track force-shut**
4. **no port-track force-shut**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**interface vfc** *if-number* **Example:** `switch(config)# interface vfc 5`	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.

**Inspur-Cisco Networking
Technology Co.,Ltd.**

CHAPTER 17 IVR

  • Information About IVR
  • Default Settings
  • Licensing Requirements
  • Guidelines and Limitations
  • Configuring IVR
  • Verifying IVR Configuration
  • Feature History

17.1 Information About IVR

Virtual SANs (VSANs) improve storage area network (SAN) scalability, availability, and security by allowing multiple Fibre Channel SANs to share a common physical infrastructure of switches and ISLs. These benefits are derived from the separation of Fibre Channel services in each VSAN and the isolation of traffic between VSANs. Data traffic isolation between the VSANs also inherently prevents sharing of resources attached to a VSAN, such as robotic tape libraries. Using IVR, you can access resources across VSANs without compromising other VSAN benefits.

IVR supports the following features:
  • Accesses resources across VSANs without compromising other VSAN benefits
  • Transports data traffic between specific initiators and targets on different VSANs without merging VSANs into a single logical fabric.
  • Establishes proper interconnected routes that traverse one or more VSANs across multiple switches. IVR is not limited to VSANs present on a common switch.
  • Shares valuable resources (such as tape libraries) across VSANs without compromise. Fibre Channel traffic does not flow between VSANs, nor can initiators access resources across VSANs other than the designated VSAN.
  • Provides efficient business continuity or disaster recovery solutions when used in conjunction with FCIP.
  • Is in compliance with Fibre Channel standards.
  • Incorporates third-party switches, however, IVR-enabled VSANs may need to be configured in one of the interop modes.

17.1.1 IVR Terminology

The following IVR-related terms are used in the IVR documentation:
Native VSAN
The VSAN to which an end device logs on is the native VSAN for that end device.

Current VSAN
The VSAN currently being configured for IVR.

Inter-VSAN Routing zone (IVR zone)
Inter-VSAN Routing zone (IVR zone)-A set of end devices that are allowed to communicate across VSANs within their interconnected SAN fabric. This definition is based on their port world-wide names (pWWNs) and their native VSAN associations

Inter-VSAN routing zone sets (IVR zone sets)
Inter-VSAN routing zone sets (IVR zone sets)-One or more IVR zones make up an IVR zone set.

IVR path

An IVR path is a set of switches and Inter-Switch Links (ISLs) through which a frame from an end device in one VSAN can reach another end device in some other VSAN. Multiple paths can exist between two such end devices.

IVR-enabled switch

A switch on which the IVR feature is enabled.

Edge VSAN

A VSAN that initiates (source edge-VSAN) or terminates (destination edge-VSAN) an IVR path. Edge VSANs may be adjacent to each other or they may be connected by one or more transit VSANs.

Transit VSAN

A VSAN that exists along an IVR path from the source edge VSAN of that path to the destination edge VSAN of that path.

Border switch

An IVR-enabled switch that is a member of two or more VSANs.

Edge switch

A switch to which a member of an IVR zone has logged in to. Edge switches are unaware of the IVR configurations in the border switches. Edge switches do not need to be IVR-enabled.

Autonomous Fabric Identifier (AFID)

Allows you to configure more than one VSAN in the network with the same VSAN ID and avoid downtime when configuring IVR between fabrics that contain VSANs with the same ID.

Service group

Allows you to reduce the amount of IVR traffic to non-IVR-enabled VSANs by configuring one or more service groups that restrict the traffic to the IVR-enabled VSANs.

17.1.2 Fibre Channel Header Modifications

IVR virtualizes the remote end devices in the native VSAN using a virtual domain. When IVR is configured to link end devices in two disparate VSANs, the IVR border switches are responsible for modifying the Fibre Channel headers for all communication between the end devices. The sections of the Fibre Channel frame headers that are modified include:

- VSAN number
- Source FCID
- Destination FCID

When a frame travels from the initiator to the target, the Fibre Channel frame header is modified such that the initiator VSAN number is changed to the target VSAN number. If IVR Network Address Translation (NAT) is enabled, then the source and destination FCIDs are also translated at the edge border switch. If IVR NAT is not enabled, then you must configure unique domain IDs for all switches involved in the IVR path.

17.1.3 IVR Database Merge

A database merge refers to the combination of the configuration database and static (unlearned) entries in the active database.

Consider the following when merging two IVR fabrics:

- The IVR configurations are merged even if two fabrics contain different configurations.

• If dissimilar zones exist in two merged fabrics, the zone from each fabric is cloned in the distributed zone set with appropriate names.
    • You can configure different IVR configurations in different Inspur SAN switches.

    To avoid traffic disruption, after the database merge is complete, the configuration is a combination of the configurations that were present on the two switches involved in the merge, as follows:
    • A combination of zones and zone sets are used to get the merged zones and zone sets. If a dissimilar zone exists in two fabrics, the dissimilar zones are cloned into the zone set with appropriate names so both zones are present.
    • The merged topology contains a combination of the topology entries for both fabrics.
    • The merge will fail if the merged database contains more topology entries than the allowed maximum.

    The following total number of items across the two fabrics cannot exceed the maximum allowed in one fabric:
• VSANs.VSANs with the same VSAN ID but different AFIDs are counted as two separate VSANs.
• IVR-enabled switches.
• Zone members. A zone member is counted twice if it exists in two zones.
• Zones.
• Zone sets.

Table 32 : Results of Merging Two IVR-Enabled Fabrics

IVR Fabric 1	IVR Fabric 2	Merged Fabric
NAT enabled	NAT disabled	Merge succeeds and NAT is enabled
Auto mode enabled	Auto mode disabled	Merge succeeds and IVR auto topology mode is enabled.
Conflicting AFID database		Merge fails
Conflicting IVR zone set database		Merge succeeds with new zones created to resolve conflicts.
Combined configuration exceeds limits (such as maximum number of zones or VSANs)		Merge fails
Service group 1	Service group 2	Merge succeeds with service groups combined
User-configured VSAN topology configuration with conflicts		Merge fails
User-configured VSAN topology configuration without conflicts		Merge succeeds

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

Related Topics
    IVR Configuration Limits
    Resolving IVR Merge Failures

17.2 Default Settings

Parameters	Default
IVR feature	Disabled
IVR NAT	Disabled
IVR distribution	Disabled
IVR Autotopology	Disabled
IVR VSANs	Not added to virtual domains
QoS for IVR Zones	Low

17.3 Licensing Requirements

Product	License
Inspur CN12700 Series.	IVR requires the FCoE license for each F-series module. FCoE enabled in a storage VDC does not require the Advanced Services License. IVR also requires the Storage Enterprise License. For a complete explanation of the Inspur INOS licensing scheme and how to obtain and apply licenses, see the *Inspur INOS*

17.4 Guidelines and Limitations

IVR has the following guidelines and limitations:

• All border switches in the fabric must be Inspur SAN switches. Other switches in the fabric can be non-Inspur switches.

• IVR must be enabled in the storage VDC.

17.5 Configuring IVR

SUMMARY STEPS
1. Enable IVR on all border switches.
2. Enable IVR distribution on all IVR-enabled switches.
3. Enable IVR NAT on a single IVR-enabled switch.
4. Enable IVR auto topology on a single IVR-enabled switch.
5. Configure and activate Zone sets.
6. Commit the IVR configuration.

DETAILED STEPS
  Step 1   Enable IVR on all border switches.
  Step 2   Enable IVR distribution on all IVR-enabled switches.
  Step 3   Enable IVR NAT on a single IVR-enabled switch.
  Step 4   Enable IVR auto topology on a single IVR-enabled switch.
  Step 5   Configure and activate Zone sets.
  Step 6   Commit the IVR configuration.

  Related Topics
      Enabling IVR
      Distributing IVR
      Enabling IVR NAT
       Enabling IVR Auto Topology
      Commiting IVR Changes

17.5.1 Enabling IVR

By default, the IVR feature is disabled on the device. You must explicitly enable the IVR feature to access the configuration and verification commands.

  Before You Begin
   ・You must enable the IVR feature from the storage VDC.
   ・You must enable the IVR feature in all border switches in the fabric that participate in the IVR.

SUMMARY STEPS

1.   **switchto vdc** *vdc-name*
2.   **configure terminal**
3.   **feature ivr**
4.   (Optional) **show feature**

DETAILED STEPS

	Command or Action	Purpose

Step 1	**switchto vdc** *vdc-name*
	
	Example:
	`switch(config)# switchto vdc fcoe`
	`switch-fcoe#`
		Switch to the storage VDC to enable the IVR feature. The *vdc-name* can be any case-sensitive alphanumeric string up to 32 characters.

Step 2	**configure terminal**
	
	Example:
	`switch-fcoe# configure terminal`
	`switch-fcoe(config)#`
		Enters global configuration mode.

Step 3	**feature ivr** **Example:** `switch-fcoe(config)# feature ivr`	Enables the IVR feature. You must enable this feature on all border switches in the fabric.

| Step 4 | **show feature**

Example:
`switch-fcoe(config)# show feature` | (Optional)
Displays the enable or disable state for all features. |

17.5.2 Distributing IVR

The IVR feature uses the Inspur Fabric Services (CFS) infrastructure to enable efficient configuration management and to provide a single point of configuration for the entire fabric in the VSAN.

The following configurations are distributed:

・IVR zones

・IVR zone sets

・IVR VSAN topology

・IVR active topology and zone set (activating these features in one switch propagates the configuration to all other distribution-enabled switches in the fabric)

・AFID database

Before You Begin

・You must enable IVR distribution on all IVR-enabled switches in the fabric.

SUMMARY STEPS

1. **configure terminal**
2. **ivr distribute**
3. (Optional) **show cfs application**
4. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command or Action	**Purpose**

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

Step 2	**ivr distribute** **Example:** `switch(config)# ivr distribute`	Enables CFS distribution for IVR configuration. You must enable IVR distribution on all IVR-enabled switches in the fabric.
Step 3	**show cfs application** **Example:** `switch(config)# show cfs application`	(Optional) Displays information about CFS enabled features, such as IVR.

Step 4	**copy running-config startup-config**
	
	Example:
	`switch(config)# copy running-config`
	`startup-config`
		(Optional)
		Copies the running configuration to the startup configuration.

17.5.3 Commiting IVR 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.

Before You Begin
• Ensure you have enabled CFS distribution for IVR.

SUMMARY STEPS

1. **configure terminal**
2. **ivr commit**
3. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**ivr commit** **Example:** `switch(config)# ivr commit`	Commits all pending IVR changes into the active IVR database.

Step 3	**copy running-config startup-config**
	
	Example:
	`switch(config)# copy running-config startup-`
	`config`
		(Optional)
		Copies the running configuration to the startup configuration.

Related Topics
Distributing IVR

17.5.4 Resolving IVR Merge Failures

Step 1   Display error conditions.
            Example:
```
switch# show ivr merge status
switch# show cfs merge
status name ivr switch#
show logging last 100
```
            Review the information from these show commands. Look for MERGE failures in the log output.

Step 2   For failures because the merged fabric exceeded the maximum configuration limits(VSANs, IVR-
            enabled switches, zone members, zones, or zone sets) where you have a different versions of INOS
            running on Inspur SAN switches, upgrade to the most recent Inspur INOS version for all switches,
            or reduce the configuration below the maximum limits.

Step 3   For failures because the merged fabric exceeded the maximum configuration limits(VSANs, IVR-
            enabled switches, zone members, zones, or zone sets) and all switches are at the same release for
            their platform, identify the switch that has the correct configuration and perform a CFS commit to

distribute the IVR configuration

Step 4   For other failures, resolve the error causing the merge failure on the switch that has the correct configuration and perform a CFS commit to distribute the IVR configuration.

After a successful CFS commit, the merge will be successful.

Related Topics
    Commiting IVR Changes

17.6 Verifying IVR Configuration

To display the IVR configuration, perform one of the following tasks:

Command	Purpose
show ivr	Displays the status for the IVR configuration.
show ivr diagnostics	Displays information about IVR diagnostics.
show ivr merge status	Displays information the last IVR merge event.
show ivr pending	Displays information about the IVR pending database.
show ivr pending-diff	Displays the differences between the pending database and the config database.
show ivr vsan-topology [**active** | **configured**]	Displays the IVR VSAN topology.
show ivr session status	Displays information about IVR CFS session.
show ivr virtual-domains	Displays information about IVR virtual domains for all local VSANs.
show ivr zone	Displays information about IVR zones.
show ivr zoneset	Displays information about IVR zone sets.
show ivr service-group active	Displays information about the active service group.
show ivr service-group configured	Displays information about the configured service
show autonomous-fabric-id database	Displays information about the AFIDs.
show ivr virtual-fcdomain-add-status	Displays the status of the IVR virtual domain

Related Topics

Information about IVR Zones and Zonesets
Configuring IVR Zones
Configuring IVR Zone Sets

17.7 Feature History

Table 33 : Feature History IVR

Feature Name	Releases	Feature Information
IVR	5.2(1)	This feature was introduced.

CHAPTER 18 IVR NAT and Auto Topology

- Information About IVR Auto Topology
- Default Settings
- Licensing Requirements
- Guidelines and Limitations for IVR NAT and Autotopology
- Configuring IVR NAT and Autotopology
- Verifying IVR Configuration
- Example: IVR Auto Topology
- Feature History

18.1 Information About IVR Auto Topology

IVR uses a configured IVR VSAN topology to determine how to route traffic between the initiator and the target across the fabric. IVR auto topology mode automatically builds the IVR VSAN topology and maintains the topology database when fabric reconfiguration occur. IVR auto topology mode also distributes the IVR VSAN topology to IVR-enabled switches using CFS.

Using IVR auto topology mode, you do not need to manually update the IVR VSAN topology when reconfiguration occur in your fabric. If an IVR manual topology database exists, IVR auto topology mode initially uses that topology information. The automatic update reduces disruption in the network by gradually migrating from the user-specified topology database to the automatically-learned topology database.

User-configured topology entries that are not part of the network are aged out in about three minutes. New entries that are not part of the user-configured database are added as they are discovered in the network.

When IVR auto topology mode is enabled, it starts with the previously active IVR manual topology if it exists, and then the discovery process begins. New, alternate, or better paths may be discovered. If the traffic is switched to an alternate or better path, there may be temporary traffic disruptions that are normally associated with switching paths.

Before configuring an IVR SAN fabric to use IVR NAT and IVR auto topology mode, consider the following:
- Configure IVR only in the relevant switches.
- Enable CFS for IVR on all switches in the fabric.

Tip   If you change any FSPF link cost, ensure that the FSPF path distance (that is, the sum of the link costs on the path) of any IVR path is less than 30,000.

18.1.1 IVR Network Address Translation

IVR Network Address Translation (NAT) can be enabled to allow non-unique domain IDs; however, without NAT, IVR requires unique domain IDs for all switches in the fabric. IVR NAT simplifies the deployment of IVR in an existing fabric where non-unique domain IDs might be present.

To use IVR NAT, you must enable it on all IVR-enabled switches in the fabric.

18.2 Default Settings

Parameters	Default
IVR feature	Disabled
IVR NAT	Disabled
IVR distribution	Disabled
IVR Autotopology	Disabled
IVR VSANs	Not added to virtual domains
QoS for IVR Zones	Low

18.3 Licensing Requirements

Product	License
Inspur CN12700 Series.	IVR requires the FCoE license for each F-series module. FCoE enabled in a storage VDC does not require the Advanced Services License. IVR also requires the Storage Enterprise License. For a complete explanation of the Inspur INOS licensing scheme and how to obtain and apply licenses, see the *Inspur INOS*

18.4 Guidelines and Limitations for IVR NAT and Autotopology

• IVR NAT port login (PLOGI) requests that are received from hosts are delayed a few seconds to perform the rewrite on the FC ID address. If the host's PLOGI timeout value is set to a value less than five seconds, it may result in the PLOGI being unnecessarily aborted and the host being unable to access the target. We recommend that you configure the host bus adapter for a timeout of at least ten seconds (most HBAs default to a value of 10 or 20 seconds).

• Load balancing of IVR NAT traffic across equal cost paths from an IVR-enabled switch is not supported.

• IVR NAT allows you to set up IVR in a fabric without needing unique domain IDs on every switch in the IVR path. IVR NAT virtualizes the switches in other VSANs by using local VSAN for the destination IDs in the Fibre Channel headers. In some Extended Link Service message types, the destination IDs are included in the packet data. In these cases, IVR NAT replaces the actual destination ID with the virtualized destination ID. IVR NAT supports destination ID replacement in the Extended Link Service messages.

• If you have a message that is not recognized by IVR NAT and contains the destination ID in the packet data, you cannot use IVR with NAT in your topology. You can still use IVR with unique domain IDs.

The following table lists the Extended Link Service messages supported by IVR NAT:

Extended Link Service Messages	Link Service Command (LS-COMMAND)	Mnemonic
Abort Exchange	0x06 00 00 00	ABTX

Discover Address	0x52 00 00 00	ADISC
Discover Address Accept	0x02 00 00 00	ADISC ACC
Fibre Channel Address Resolution Protocol Reply	0x55 00 00 00	FARP-REPLY
Fibre Channel Address Resolution Protocol Request	0x54 00 00 00	FARP-REQ
Logout	0x05 00 00 00	LOGO
Port Login	0x30 00 00 00	PLOGI
Read Exchange Concise	0x13 00 00 00	REC
Read Exchange Concise Accept	0x02 00 00 00	REC ACC
Read Exchange Status Block	0x08 00 00 00	RES
Read Exchange Status Block Accept	0x02 00 00 00	RES ACC
Read Link Error Status Block	0x0F 00 00 00	RLS
Read Sequence Status Block	0x09 00 00 00	RSS
Reinstate Recovery Qualifier	0x12 00 00 00	RRQ
Request Sequence Initiative	0x0A 00 00 00	RSI
Scan Remote Loop	0x7B 00 00 00	RSL
Third Party Process Logout	0x24 00 00 00	TPRLO
Third Party Process Logout Accept	0x02 00 00 00	TPRLO ACC

18.4.1 Transit VSAN Guidelines

Consider the following guidelines for transit VSANs:

• In addition to defining the IVR zone membership, you can choose to specify a set of transit VSANs to provide connectivity between two edge VSANs:

◦If two edge VSANs in an IVR zone overlap, then a transit VSAN is not required (though, not prohibited) to provide connectivity.

◦If two edge VSANs in an IVR zone do not overlap, you may need one or more transit VSANs to provide connectivity. Two edge VSANs in an IVR zone will not overlap if IVR is not enabled on a switch that is a member of both the source and destination edge VSANs.

• Traffic between the edge VSANs only traverses through the shortest IVR path.

• Transit VSAN information is common to all IVR zone sets. Sometimes, a transit VSAN can also act as an edge VSAN in another IVR zone.

18.4.2 Border Switch Guidelines

• A border switch must be a member of two or more VSANs.

• A border switch that facilitates IVR communications must be IVR-enabled.

• IVR can (optionally) be enabled on additional border switches to provide redundant paths between active IVR zone members.

• The VSAN topology configuration updates automatically when a border switch is added or removed.

18.5 Configuring IVR NAT and Autotopology

18.5.1 Enabling IVR NAT

Before You Begin

• Ensure you have enabled the IVR feature and IVR distribution.

SUMMARY STEPS
1. **configure terminal**
2. **ivr nat**
3. (Optional) **show ivr**
4. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**ivr nat** **Example:** `switch(config)# ivr nat`	Enables IVR NAT.
Step 3	**show ivr** **Example:** `switch(config)# show ivr`	(Optional) Displays information about IVR.
Step 4	**copy running-config startup-config** **Example:** `switch(config)# copy running-config startup-config`	(Optional) Copies the running configuration to the startup configuration.

18.5.2 Enabling IVR Auto Topology

Before You Begin

• Ensure you have enabled the IVR feature and IVR distribution.

SUMMARY STEPS
1. **configure terminal**
2. **ivr vsan-topology auto**
3. (Optional) **show ivr vsan topology**
4. (Optional) **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**ivr vsan-topology auto** **Example:** `switch(config)# ivr vsan-topology auto`	Enables IVR auto topology mode.
Step 3	**show ivr vsan topology** **Example:** `switch(config)# show ivr vsan topology`	(Optional) Displays the automatically discovered IVR topology and the topology mode.
Step 4	**copy running-config startup-config** **Example:** `switch(config)# copy running-config startup-config`	(Optional) Copies the running configuration to the startup configuration.

18.6 Verifying IVR Configuration

To display the IVR configuration, perform one of the following tasks:

Command	Purpose
show ivr	Displays the status for the IVR configuration.
show ivr diagnostics	Displays information about IVR diagnostics.
show ivr merge status	Displays information the last IVR merge event.
show ivr pending	Displays information about the IVR pending database.
show ivr pending-diff	Displays the differences between the pending database and the config database.
show ivr vsan-topology [**active** | **configured**]	Displays the IVR VSAN topology.
show ivr session status	Displays information about IVR CFS session.
show ivr virtual-domains	Displays information about IVR virtual domains for all local VSANs.

Command	Purpose
show ivr zone	Displays information about IVR zones.
show ivr zoneset	Displays information about IVR zone sets.
show ivr service-group active	Displays information about the active service group.
show ivr service-group configured	Displays information about the configured service
show autonomous-fabric-id database	Displays information about the AFIDs.
show ivr virtual-fcdomain-add-status	Displays the status of the IVR virtual domain

Related Topics
Information about IVR Zones and Zonesets
Configuring IVR Zones
Configuring IVR Zone Sets

18.7 Example: IVR Auto Topology

Step 1   Enable IVR on every border switch in the fabric.
Example:
```
switch# config t
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# feature ivr
switch(config)# exit
switch#
```

Step 2   Verify that IVR is enabled on every IVR-enabled switch.
Example:
```
switch# show ivr
Inter-VSAN Routing is enabled

Inter-VSAN enabled switches
-------------------------
No IVR-enabled VSAN is active. Check VSAN-Topology configuration.

Inter-VSAN topology status
-------------------------
Current Status: Inter-VSAN topology is INACTIVE

Inter-VSAN zoneset status
-----------------------
name              :
    state             : idle
    last activate time :
```

```
Fabric distribution status
--------------------
fabric  distribution
disabled Last Action       : None
Last Action Result         : None Last Action Failure Reason : None

Inter-VSAN NAT mode status
------------------------
FCID-NAT is disabled

License status
----------------
IVR is running based on the following
license(s) ENTERPRISE_PKG
```

Step 3   Enable CFS distribution on every IVR-enabled switch in the fabric.
             Example:
```
switch# config t
Enter configuration commands, one per line. End with
CNTL/Z. switch(config)# ivr distribution
```

Step 4   Enable IVR auto topology mode.
             Example:
```
switch(config)# ivr vsan-topology auto
fabric is locked for configuration. Please commit after
configuration is done.
```

Step 5   Commit the change to the fabric.
             Example:
```
switch(config)#
ivr commit
switch(config)#
exit switch#
```

Step 6   Verify the status of the commit request.
             Example:
```
switch# show ivr session status
Last Action               : Commit
Last Action Result        :
Success Last Action Failure
Reason : None
```

Step 7   Verify the active IVR auto topology.
             Example:
```
switch# show ivr vsan-topology active

AFID   SWITCH WWN                 Active   Cfg. VSANS
-------------------------------------------------------------
   1  20:00:00:0d:ec:08:6e:40 *  yes    no     1,336-
   1  20:00:00:0d:ec:0c:99:40    yes    no     336,339
```

Step 8   Configure IVR zone set and zones.
             Example:
```
switch(config)# ivr zoneset name tape_server1_server2

switch(config-ivr-zoneset)# zone name tape_server1
switch(config-ivr-zoneset-zone)#              member              pwwn
10:02:50:45:32:20:7a:52  vsan  1  switch(config-ivr-zoneset-zone)#
member  pwwn  10:02:66:45:00:20:89:04  vsan  2  switch(config-ivr-
zoneset-zone)# exit

switch(config-ivr-zoneset)# zone name tape_server2
switch(config-ivr-zoneset-zone)#              member              pwwn
10:02:50:45:32:20:7a:52  vsan  1  switch(config-ivr-zoneset-zone)#
```

```
     member  pwwn  10:00:ad:51:78:33:f9:86  vsan  3  switch(config-ivr-
     zoneset-zone)# exit
```

Two zones are required:
 • One zone has tape T (pwwn 10:02:50:45:32:20:7a:52) and server S1 (pwwn 10:02:66:45:00:20:89:04).
 • Another zone has tape T and server S2 (pwwn 10:00:ad:51:78:33:f9:86).

Instead of creating two IVR zones, you can also create one IVR zone with the tape and both servers.

Step 9   View the IVR zone configuration to confirm that the IVR zone set and IVR zones are properly
         configured.
         Example:
```
switch(config)# show ivr zoneset
zoneset name
  tape_server1_server2
  zone name tape_server1
      pwwn
      10:02:50:45:32:20:7a:52
      vsan 1 pwwn
      10:02:66:45:00:20:89:04
      vsan 2

  zone name tape_server2
      pwwn 10:02:50:45:32:20:7a:52 vsan 1
      pwwn 10:00:ad:51:78:33:f9:86 vsan 3
```

Step 10   View the zone set prior to IVR zone set activation. Prior to activating the IVR zone set, view the
         active zone set. Repeat this step for VSANs 2 and 3.
         Example:
```
switch(config)# show zoneset active vsan 1
zoneset name finance_dept
  vsan 1 zone name
  accounts_database vsan 1
          pwwn
   10:00:23:11:ed:f6:23:
        12 pwwn
   10:00:56:43:11:56:fe:
            ee

  zone name $default_zone$ vsan 1
```

Step 11   Activate the configured IVR zone set.

         Example:
```
switch(config)# ivr zoneset activate name
tape_server1_server2 zoneset activation initiated.
check inter-VSAN zoneset status switch(config)# exit
switch#
```

Step 12   Verify the IVR zone set activation.
         Example:
```
switch# show ivr zoneset active
zoneset name
  tape_server1_server2
  zone name tape_server1
      pwwn
      10:02:50:45:32:20:7a:52
      vsan 1 pwwn
      10:02:66:45:00:20:89:04
      vsan 2

  zone name tape_server2
      pwwn 10:02:50:45:32:20:7a:52 vsan 1
```

```
       pwwn 10:00:ad:51:78:33:f9:86 vsan 3
```

Step 13   Verify the zone set updates. Upon successful IVR zone set activation, verify that appropriate zones are added to the active zone set. Repeat this step for VSANs 2 and 3.

Example:

```
switch# show zoneset active vsan 1
zoneset name finance_dept
  vsan 1 zone name
  accounts_database vsan 1
            pwwn
   10:00:23:11:ed:f6:2
        3:12 pwwn
   10:00:56:43:11:56:f
           e:ee

  zone name
    IVRZ_tape_server1 vsan
    1 pwwn
    10:02:66:45:00:20:89:04
    pwwn
    10:02:50:45:32:20:7a:52

  zone name
    IVRZ_tape_server2 vsan
    1 pwwn
    10:02:50:45:32:20:7a:52
    pwwn
    10:00:ad:51:78:33:f9:86

  zone name

$default_zone$ vsan 1

switch# show ivr zoneset

status
Zoneset Status


    name                : tape_server1_server2
    state               : activation success
    last activate time : Tue May 20
    23:23:01 1980 force option   : on

status per vsan:


    vsan     status

    1        active
```

18.8 Feature History

Table 34 : Feature History IVR

Feature Name	Releases	Feature Information
IVR	5.2(1)	This feature was introduced.

CHAPTER 19 IVR Zones and Zonesets

- Information about IVR Zones and Zonesets
- Default Settings
- Licensing Requirements
- Guidelines and Limitations
- Configuring IVR Zones and Zonesets
- Verifying IVR Configuration
- Feature History

19.1 Information about IVR Zones and Zonesets

As part of the IVR configuration, you need to configure one or more IVR zones to enable cross-VSAN communication. To achieve this result, you must specify each IVR zone as a set of (pWWN, VSAN) entries. Like zones, several IVR zone sets can be configured to belong to an IVR zone. You can define several IVR zone sets and activate only one of the defined IVR zone sets.

Table 35 : Key Differences Between IVR Zones and Zones

IVR Zones	Zones
IVR zone membership is specified using the VSAN and pWWN combination.	Zone membership is specified using pWWN, fabric WWN, sWWN, or the AFID.
Default zone policy is always deny (not configurable).	Default zone policy is deny (configurable).

As part of the IVR configuration, you need to configure one or more IVR zone to enable cross-VSAN communication. To achieve this, you must specify each IVR zone as a set of (pWWN, VSAN) entries. Different IVR zone sets can contain the same IVR zone, because IVR zones can be members of one or more IVR zone sets.

Related Topics
Configuring IVR Zones
Configuring IVR Zone Sets
Guidelines and Limitations
Verifying IVR Configuration

19.1.1 Automatic IVR Zone Creation

To allow pwwn1 to communicate with pwwn2, they must be in the same zone in VSAN 1, as well as in VSAN 2. If they are not in the same zone, then the hard-zoning ACL entries will prohibit pwwn1 from communicating with pwwn2.

A zone corresponding to each active IVR zone is automatically created in each edge VSAN specified in the active IVR zone. All pWWNs in the IVR zone are members of these zones in each VSAN.

Figure 27 : Creating Zones Upon IVR Zone Activation
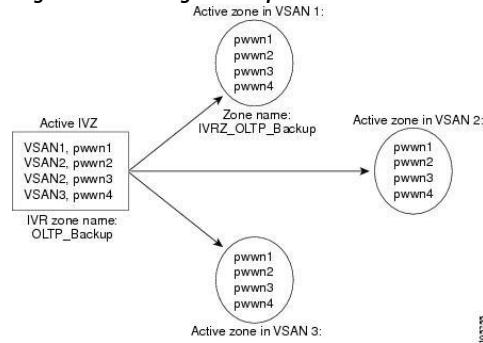
The zones are created automatically by the IVR process when an IVR zone set is activated. They are not stored in a full zone set database and are lost when the switch reboots or when a new zone set is activated. The IVR feature monitors these events and adds the zones corresponding to the active IVR zone set configuration when a new zone set is activated. Like zone sets, IVR zone sets are also activated nondisruptively.

If pwwn1 and pwwn2 are in an IVR zone in the current as well as the new IVR zone set, then activation of the new IVR zone set does not cause any traffic disruption between them.

19.2 Default Settings

Parameters	Default
IVR feature	Disabled
IVR NAT	Disabled
IVR distribution	Disabled
IVR Autotopology	Disabled
IVR VSANs	Not added to virtual domains
QoS for IVR Zones	Low

19.3 Licensing Requirements

Product	License
Inspur CN12700 Series.	IVR requires the FCoE license for each F-series module. FCoE enabled in a storage VDC does not require the Advanced Services License. IVR also requires the Storage Enterprise License. For a complete explanation of the Inspur INOS licensing scheme and how to obtain and apply licenses, see the *Inspur INOS*

19.4 Guidelines and Limitations

When interop mode is enabled, consider the following IVR configuration guidelines:

• When a member's native VSAN is in interop mode (for example, when the interop mode is 2, 3, or 4), then ReadOnly, the QoS attribute, and LUN zoning are not permitted

• When a member's VSAN is already in interop mode and an attempt is made to configure ReadOnly, the QoS attribute, or LUN zoning, a warning message is displayed to indicate that the configuration is not permitted.

• When you configure ReadOnly, the QoS attribute, or LUN zoning first, and then change the member's VSAN interop mode, a warning message is displayed to indicate the configuration is not permitted. You are then prompted to change the configuration.

This example shows samples of the warning messages that are displayed when configuration changes are made that affect ReadOnly, the QoS attribute, and LUN zoning.

```
switch(config)# vsan database
switch(config-vsan-db)# vsan 2
switch(config-vsan-db)# vsan 2 interop 2
switch(config-vsan-db)# exit

switch(config)# ivr zoneset name ivr_zs1
switch(config-ivr-zoneset)# zone name ivr_z1
switch(config-ivr-zoneset-zone)# member pwwn
21:00:00:14:c3:3d:45:22 lun 0x32 vsan 2
VSAN is in interop mode, and LUN zoning cannot be set.

switch(config)# ivr zoneset name ivr_zs1
switch(config-ivr-zoneset)# zone name ivr_z1
switch(config-ivr-zoneset-zone)# member pwwn 21:00:00:14:c3:3d:45:22 vsan 2
switch(config-ivr-zoneset-zone)# attribute read-only
VSAN is in interop mode and zone member has been configured, zone cannot be set to
READ-ONLY. switch(config-ivr-zoneset-zone)# attribute qos priority medium
VSAN is in interop mode and zone member has been
configured, QoS cannot be assigned to zone.
```

Related Topics

Information about IVR Zones and Zonesets
Configuring IVR Zones
Configuring IVR Zone Sets

19.5 Configuring IVR Zones and Zonesets

19.5.1 Configuring IVR Zones

Before You Begin

• Ensure you are in the correct storage-based VDC.

• Ensure you have enabled the IVR feature.

SUMMARY STEPS

1. **configure terminal**
2. **ivr zone name** *zonename*
3. **member pwwn** *pwwn* **vsan** *vsan-id*
4. (Optional) **show ivr pending-diff**
5. (Optional) **show ivr zone**
6. (Optional) **ivr commit**

**Inspur-Cisco Networking
Technology Co.,Ltd.**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**ivr zone name** *zonename*
	
	Example:
	`switch(config)# ivr zone name`
	`sample_vsan2-3 switch(config-ivr-zone)#`
		Creates the IVR zone and enters IVR zone configuration mode. The *zonename* can be any case-sensitive, alphanumeric string up to 59 characters.

Step 3	**member pwwn** *pwwn* **vsan** *vsan-id*
	
	Example:
	`switch(config-ivr-zone)# member`
	`pwwn 21:00:00:20:37:c8:5c:6b vsan`
	`2`
		Adds the specified pWWN in VSAN 2 as an IVR zone member. The *pwwn* is in colon-separated hexadecimal format. The *vsan* range is from 1 to 4093.

Step 4	**show ivr pending-diff**
	
	Example:
	`switch(config-ivr-zone)# show ivr`
	`pending-diff`
		(Optional)
		Displays information about the pending changes to the IVR database. This displays changes that have not been committed yet.

Step 5	**show ivr zone** **Example:** `switch(config-ivr-zone)# show ivr zone`	(Optional) Displays information about the zones in the active zone database.
Step 6	**ivr commit** **Example:** `switch(config-ivr-zone)# ivr commit`	(Optional) Commits all pending changes to IVR to the active IVR database and distributes these changes to all IVR-enabled switches in the fabric.

What to Do Next

You must commit the IVR changes to make these changes permanent and distribute the changes to all IVR-enabled switches in the fabric.

Related Topics

Information about IVR Zones and Zonesets
Guidelines and Limitations
Verifying IVR Configuration

19.5.2 Configuring IVR Zone Sets

Before You Begin

• Ensure you are in the correct storage-based VDC.

• Ensure you have enabled the IVR feature.

SUMMARY STEPS

1. **configure terminal**
2. **ivr zoneset name** *zoneset-name*

3. **member** *zonename*
4. (Optional) **show ivr pending-diff**
5. (Optional) **show ivr zoneset**
6. (Optional) **ivr commit**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**ivr zoneset name** *zoneset-name*
	
	Example:
	`switch(config)# ivr zoneset name`
	`ivrZoneset1 switch(config-ivr-zoneset)#`
		Creates the IVR zone set and enters IVR zone set configuration mode. The *zoneset-name* can be any case-sensitive, alphanumeric string up to 59 characters.

Step 3	**member** *zonename*
	
	Example:
	`switch(config-ivr-zoneset)# member`
	`sample_vsan2-3`
		Adds the specified IVR zone as an IVR zone set member. The *zoneset-name* can be any case-sensitive, alphanumeric string up to 59 characters.

Step 4	**show ivr pending-diff**
	
	Example:
	`switch(config-ivr-zoneset)# show ivr`
	`pending-diff`
		(Optional)
		Displays information about the pending changes to the IVR database. This displays changes that have not been committed yet.

Step 5	**show ivr zoneset** **Example:** `switch(config-ivr-zoneset)# show ivr`	(Optional) Displays information about the zone sets in the active zone set database.
Step 6	**ivr commit** **Example:** `switch(config-ivr-zoneset)# ivr commit`	(Optional) Commits all pending changes to IVR to the active IVR database and distributes these changes to all IVR-enabled switches in the fabric.

What to Do Next

You must commit the IVR changes to make these changes permanent and distribute the changes to all IVR-enabled switches in the fabric. You must also activate the zone set.

Related Topics

Information about IVR Zones and Zonesets
Guidelines and Limitations
Verifying IVR Configuration

19.5.3 Configuring LUNs in IVR Zoning

LUN zoning can be used between members of active IVR zones. You can configure the service by creating and activating LUN zones between the desired IVR zone members in all relevant edge VSANs using the zoning interface

---

or you can use LUN zoning directly supported by IVR.

Before You Begin

• Ensure you are in the correct storage-based VDC.

• Ensure you have enabled the IVR feature.

SUMMARY STEPS

1. **configure terminal**
2. **ivr zone name** *zonename*
3. **member pwwn** *pwwn* **lun** *lun-id* **vsan** *vsan-id* [**autonomous-fabric-id** *afid*]
4. (Optional) **show ivr pending-diff**
5. (Optional) **show ivr zone**
6. (Optional) **ivr commit**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**ivr zone name** *zonename*
	
	Example:
	`switch(config)# ivr zone name ivrLunZone`
	`switch(config-ivr-zone)#`
		Creates the IVR zone and enters IVR zone configuration mode. The *zonename* can be any case-sensitive, alphanumeric string up to 59 characters.

Step 3	**member pwwn** *pwwn* **lun** *lun-id* **vsan** *vsan-id* [
	autonomous-fabric-id *afid*]
	
	Example:
	`switch(config-ivr-zone)# member`
	`pwwn 21:00:00:20:37:c8:5c:6b`
	` lun 0x64 vsan 2`
		Configures an IVR zone member based on the specified pWWN and LUN value.
		
		Note The CLI interprets the LUN identifier value as a hexadecimal value whether or not the 0x prefix is included.
		The *pwwn* is in colon-separated hexadecimal format. The *lun-id* is in hexadecimal notation. The *vsan* range is from 1 to 4093.

Step 4	**show ivr pending-diff**
	
	Example:
	`switch(config-ivr-zone)# show ivr`
	`pending-diff`
		(Optional)
		Displays information about the pending changes to the IVR database. This displays changes that have not been committed yet.

Step 5	**show ivr zone** **Example:** `switch(config-ivr-zone)# show ivr zone`	(Optional) Displays information about the zones in the active zone database.
Step 6	**ivr commit** **Example:** `switch(config-ivr-zone)# ivr commit`	(Optional) Commits all pending changes to IVR to the active IVR database and distributes these changes to all IVR-enabled switches in the fabric.

19.5.4 Configuring the QoS Attribute

Before You Begin

· Ensure you are in the correct storage-based VDC.

· Ensure you have enabled the IVR feature.

SUMMARY STEPS

1. **configure terminal**
2. **ivr zone name** *zonename*
3. **attribute qos priority** { **low** | **medium** | **high**}
4. (Optional) **show ivr pending-diff**
5. (Optional) **show ivr zone**
6. (Optional) **ivr commit**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**ivr zone name** *zonename*
	
	Example:
	`switch(config)# ivr zone name`
	`sample_vsan2-3 switch(config-ivr-zone)#`
		Creates the IVR zone and enters IVR zone configuration mode. The *zonename* can be any case-sensitive, alphanumeric string up to 59 characters.

Step 3	**attribute qos priority** { **low**	**medium**	**high**}
			
			Example:
			`switch(config-ivr-zone)# attribute`
			`qos priority medium`
				Configures the QoS for IVR zone traffic.

Step 4	**show ivr pending-diff** **Example:** `switch(config-ivr-zone)# show ivr pending-diff`	(Optional) Displays information about the pending changes to the IVR database. This displays changes that have not been committed yet.
Step 5	**show ivr zone** **Example:** `switch(config-ivr-zone)# show ivr zone`	(Optional) Displays information about the zones in the active zone database.
Step 6	**ivr commit** **Example:** `switch(config-ivr-zone)# ivr commit`	(Optional) Commits all pending changes to IVR to the active IVR database and distributes these changes to all IVR-enabled switches in the fabric.

19.5.5 Configuring Read-only Zoning

Read-only zoning (with or without LUNs) can be used between members of active IVR zones. To configure this service, you must create and activate read-only zones between the desired IVR zone members in all relevant edge VSANs using the zoning interface.

Before You Begin
- Ensure you are in the correct storage-based VDC.
- Ensure you have enabled the IVR feature.

SUMMARY STEPS

1. **configure terminal**
2. **ivr zone name** *zonename*
3. **attribute read-only**
4. (Optional) **show ivr pending-diff**
5. (Optional) **show ivr zone**
6. (Optional) **ivr commit**

DETAILED STEPS

	Command or Action	Purpose

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

Step 2	**ivr zone name** *zonename*
	
	Example:
	`switch(config)# ivr zone name sample_vsan2-3`
	`switch(config-ivr-zone)#`
		Enters IVR zone configuration mode. The *zonename* can be any case-sensitive, alphanumeric string up to 59 characters.

Step 3	**attribute read-only** **Example:** `switch(config-ivr-zone)# attribute read-only`	Configures the QoS for IVR zone traffic.
Step 4	**show ivr pending-diff** **Example:** `switch(config-ivr-zone)# show ivr pending-diff`	(Optional) Displays information about the pending changes to the IVR database. This displays changes that have not been committed yet.
Step 5	**show ivr zone** **Example:** `switch(config-ivr-zone)# show ivr zone`	(Optional) Displays information about the zones in the active zone database.

Step 6	**ivr commit** **Example:** `switch(config-ivr-zone)# ivr commit`	(Optional) Commits all pending changes to IVR to the active IVR database and distributes these changes to all IVR-enabled switches in the fabric.

19.6 Verifying IVR Configuration

To display the IVR configuration, perform one of the following tasks:

Command	Purpose
show ivr	Displays the status for the IVR configuration.
show ivr diagnostics	Displays information about IVR diagnostics.
show ivr merge status	Displays information the last IVR merge event.
show ivr pending	Displays information about the IVR pending database.
show ivr pending-diff	Displays the differences between the pending database and the config database.
show ivr vsan-topology [**active** | **configured**]	Displays the IVR VSAN topology.
show ivr session status	Displays information about IVR CFS session.
show ivr virtual-domains	Displays information about IVR virtual domains for all local VSANs.
show ivr zone	Displays information about IVR zones.
show ivr zoneset	Displays information about IVR zone sets.
show ivr service-group active	Displays information about the active service group.
show ivr service-group configured	Displays information about the configured service
show autonomous-fabric-id database	Displays information about the AFIDs.
show ivr virtual-fcdomain-add-status	Displays the status of the IVR virtual domain configuration.

Related Topics

Information about IVR Zones and Zonesets
Configuring IVR Zones
Configuring IVR Zone Sets

19.7 Feature History

Table 36Table 36: Feature History IVR

Feature Name	Releases	Feature Information
IVR	5.2(1)	This feature was introduced.

CHAPTER 20 IVR Topology

- Information About IVR Without NAT or Autotopology
- Licensing Requirements
- Guidelines for Manual IVR Topology
- Default Settings
- Configuring Manual Topology
- Verifying IVR Configuration
- Feature History

20.1 Information About IVR Without NAT or Autotopology

Before configuring an IVR SAN fabric without IVR in NAT mode or IVR auto topology mode, consider the following general guidelines:
- Acquire a mandatory Enterprise License Package or SAN-EXTENSION license package and one active IPS card for this feature.
- If you change an FSPF link cost, ensure that the FSPF path distance (the sum of the link costs on the path) of any IVR path is less than 30,000.
- IVR-enabled VSANs can be configured when an interop mode is enabled or disabled.

20.1.1 Domain ID Guidelines

Before configuring domain IDs, consider the following guidelines:
- Configure unique domain IDs across all VSANs and switches participating in IVR operations if you are not using IVR NAT. The following switches participate in IVR operations:
- All edge switches in the edge VSANs (source and destination)
- All switches in transit VSANs
- Minimize the number of switches that require a domain ID assignment. This ensures minimum traffic disruption.
- Minimize the coordination between interconnected VSANs when configuring the SAN for the first time as well as when you add each new switch.

You can configure domain IDs using one of two options:
- Configure the allowed-domains list so that the domains in different VSANs are non-overlapping on all participating switches and VSANs.
- Configure static, non-overlapping domains for each participating switch and VSAN.

20.1.2 Transit VSAN Guidelines

Before configuring transit VSANS, consider the following guidelines:
- Traffic between the edge VSANs only traverses through the shortest IVR path.
- Transit VSAN information is common to all IVR zone sets. Sometimes, a transit VSAN can also act as an edge VSAN in another IVR zone.

Besides defining the IVR zone membership, you can choose to specify a set of transit VSANs to provide connectivity between two edge VSANs:
- If two edge VSANs in an IVR zone overlap, then a transit VSAN is not required (though, not prohibited) to

provide connectivity.

• If two edge VSANs in an IVR zone do not overlap, you may need one or more transit VSANs to provide connectivity. Two edge VSANs in an IVR zone will not overlap if IVR is not enabled on a switch that is a member of both the source and destination edge VSANs.

20.1.3 Border Switch Guidelines

Before configuring border switches, consider the following guidelines:

• Configure IVR only in the relevant border switches.

• Border switches require Inspur MDS SAN-OS Release 1.3(1) or later.

• A border switch must be a member of two or more VSANs.

• A border switch that facilitates IVR communications must be IVR enabled.

• IVR can also be enabled on additional border switches to provide redundant paths between active IVR zone members.

• The VSAN topology configuration must be updated before a border switch is added or removed.

20.2 Licensing Requirements

Product	License
Inspur CN12700 Series.	IVR requires the FCoE license for each F-series module. FCoE enabled in a storage VDC does not require the Advanced Services License. IVR also requires the Storage Enterprise License. For a complete explanation of the Inspur INOS licensing scheme and how to obtain and apply licenses, see the *Inspur INOS*

20.3 Guidelines for Manual IVR Topology

You must create the IVR topology on every IVR-enabled switch in the fabric if you have not enabled IVR auto topology mode. To use IVR manual topology mode, follow the instructions in this section.

Consider the following guidelines when using IVR manual topology mode:

• You can configure a maximum of 128 IVR-enabled switches and 128 distinct VSANs in an IVR topology.

• If two VSANs in an IVR topology have the same VSAN ID and different AFIDs, they count as two VSANs for the 128-VSAN limit for IVR.

• The use of a single AFID does not allow for segmented VSANs in an inter-VSAN routing topology.

You will need to specify the IVR topology using the following information:

• The switch WWNs of the IVR-enabled switches.

• A minimum of two VSANs to which the IVR-enabled switch belongs.

• The AFID, which distinguishes two VSANs that are logically and physically separate, but have the same VSAN number. You can specify up to 64 AFIDs.

20.4 Default Settings

Parameters	Default
IVR feature	Disabled
IVR NAT	Disabled
IVR distribution	Disabled
IVR Autotopology	Disabled
IVR VSANs	Not added to virtual domains
QoS for IVR Zones	Low

20.5 Configuring Manual Topology

20.5.1 Manually Configuring an IVR Topology

You can manually add a switch or VSANs to an IVR topology.

Before You Begin
Use the **show wwn switch** command to obtain the switch WWNs of the IVR-enabled switches.

SUMMARY STEPS
1. **ivr vsan-topology database**
2. **autonomous-fabric-id** *f-id* **switch** *switch-id* **vsan-ranges** *range*
3. Repeat on all IVR-enabled switches or distribute with CFS.
4. **ivr vsan-topology activate**
5. (Optional) **show ivr vsan-topology**

DETAILED STEPS

	Command or Action	Purpose

Step 1	**ivr vsan-topology database**
	
	Example:
	`switch(config)# ivr vsan-topology database`
	`switch(config-ivr-topology-db)#`
		Enters the VSAN topology database configuration mode for the IVR feature.

Step 2	**autonomous-fabric-id** *f-id* **switch** *switch-id* **vsan-ranges** *range*
	
	Example:
	`switch(config-ivr-topology-`
	`db)# autonomous-fabric-id 1`
		Configures the VSANS that participate in IVR for this switch.

Step 3	Repeat on all IVR-enabled switches or distribute with CFS.	Ensures all IVR-enabled switches have the updated IVR topology.

Step 4	**ivr vsan-topology activate** **Example:** `switch(config)# ivr vsan-topology activate`	Activates the IVR topology. **Note** Active IVR topologies cannot be deactivated. You can only switch to IVR auto topology mode.
Step 5	**show ivr vsan-topology** **Example:**	(Optional) `switch(config)# show ivr vsan-topology` Displays the IVR topology.

In the following example output, VSAN 2 is the transit VSAN between VSANs 1, 5, and 6.

```
switch# show ivr vsan-topology
AFID  SWITCH WWN                Active   Cfg. VSANS
--------------------------------------------------
   -------- 1  20:00:00:05:30:01:1b:c2 *  yes   yes  1-
   2
   1  20:02:00:44:22:00:4a:05      yes    yes  1-2,6
1  20:02:00:44:22:00:4a:07     yes     yes  2-5

Total:   3 entries in active and configured IVR VSAN-

Topology Current Status: Inter-VSAN topology is

ACTIVE
Last activation time: Mon Mar 24 07:19:53 2011
```

What to Do Next

Tip Transit VSANs are deduced based on your configuration. The IVR feature does not have an explicit transit-VSAN configuration.

20.5.2 Copying the Active Topology to the Configure Topology

You can edit a manually configured IVR topology; however, you cannot edit an active IVR topology.

SUMMARY STEPS
1. **ivr copy active-topology user-configured-topology**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**ivr copy active-topology user-configured-topology** **Example:** `switch# ivr copy active-topology user-configured-topology`	Copies the active database to the configure database so that you can edit the topology.

20.5.3 Clearing the Manual Topology

SUMMARY STEPS
1. **no ivr vsan-topology database**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**no ivr vsan-topology database** **Example:** `switch(config)# no ivr vsan-topology database`	Clears the manually added IVR topology.

20.5.4 Migrating from Autotopology to Manual Topology

If you want to migrate from IVR auto topology mode to IVR manual topology mode, copy the active IVR VSAN topology database to the user-configured IVR VSAN topology database before switching modes.

SUMMARY STEPS

1. **ivr copy auto-topology user-configured-topology**
2. **configure terminal**
3. **ivr vsan-topology activate**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**ivr copy auto-topology user-configured-topology** **Example:** `switch# ivr copy auto-topology user-onfigured-topology`	Copies the automatic IVR topology database to the user-configured IVR topology.

Step 2	**configure terminal**
	
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters configuration mode.

Step 3	**ivr vsan-topology activate** **Example:** `switch(config)# ivr vsan-topology activate`	Activates the IVR topology. **Note** Active IVR topologies cannot be deactivated. You can only switch to IVR auto topology mode.

This task disables IVR auto topology mode for the IVR topology database and enables IVR manual topology mode.

20.6 Verifying IVR Configuration

To display the IVR configuration, perform one of the following tasks:

Command	Purpose
show ivr	Displays the status for the IVR configuration.
show ivr diagnostics	Displays information about IVR diagnostics.
show ivr merge status	Displays information the last IVR merge event.
show ivr pending	Displays information about the IVR pending database.
show ivr pending-diff	Displays the differences between the pending database and the config database.
show ivr vsan-topology [**active** | **configured**]	Displays the IVR VSAN topology.
show ivr session status	Displays information about IVR CFS session.
show ivr virtual-domains	Displays information about IVR virtual domains for all local VSANs.
show ivr zone	Displays information about IVR zones.
show ivr zoneset	Displays information about IVR zone sets.
show ivr service-group active	Displays information about the active service group.
show ivr service-group configured	Displays information about the configured service
show autonomous-fabric-id database	Displays information about the AFIDs.
show ivr virtual-fcdomain-add-status	Displays the status of the IVR virtual domain configuration.

Related Topics

Information about IVR Zones and Zonesets
Configuring IVR Zones

Configuring IVR Zone Sets

20.7 Feature History

Table 37 : Feature History IVR

Feature Name	Releases	Feature Information
IVR	5.2(1)	This feature was introduced.

CHAPTER 21 Autonomous Fabric IDs

- Information About Autonomous Fabric IDs
- Licensing Requirements
- Guidelines and Limitations
- Default Settings
- Configuring AFIDs
- Verifying IVR Configuration
- Feature History

21.1 Information About Autonomous Fabric IDs

You can configure AFIDs individually for VSANs, or you can set the default AFIDs for all VSANs on a switch. If you configure an individual AFID for a subset of the VSANs on a switch that has a default AFID, that subset uses the configured AFID while all other VSANs on that switch use the default AFID.

You can only use an AFID configuration when the VSAN topology is in IVR auto topology mode. In IVR manual topology mode, the AFIDs are specified in the VSAN topology configuration itself and a separate AFID configuration is not needed.

When devices attached to multiple switches belong to one VSAN, they cannot communicate with each other by configuring the regular zone set because the AFIDs are different. You can consider that the different AFIDs are different fabrics; therefore, the three switches represent three separate fabrics.

If we specify the IVR VSAN topology as shown below, IVR will set up the connection between the devices across the switches even though they have the same VSAN.

```
switch# show ivr vsan-topology
  AFID  SWITCH WWN                 Active    Cfg.    VSANS
  -------------------------------------------------
1 20:00:00:0d:ec:27:6b:c0  yes        yes     1
2 20:00:00:0d:ec:27:6c:00  yes        yes     1
3 20:00:00:0d:ec:27:6c:40  yes        yes     1

              Total:   3 entries in active and configured IVR VSAN-Topology
```

21.2 Licensing Requirements

Product	License
Inspur CN12700 Series.	IVR requires the FCoE license for each F-series module. FCoE enabled in a storage VDC does not require the Advanced Services License. IVR also requires the Storage Enterprise License. For a complete explanation of the Inspur INOS licensing scheme and how to obtain and apply licenses, see the *Inspur INOS*

21.3 Guidelines and Limitations

IVR has the following guidelines and limitations:

·All border switches in the fabric must be Inspur SAN switches. Other switches in the fabric can be non-Inspur switches.

·IVR must be enabled in the storage VDC.

21.4 Default Settings

Parameters	Default
IVR feature	Disabled
IVR NAT	Disabled
IVR distribution	Disabled
IVR Autotopology	Disabled
IVR VSANs	Not added to virtual domains
QoS for IVR Zones	Low

21.5 Configuring AFIDs

21.5.1 Configuring Default AFIDs

SUMMARY STEPS
1. **autonomous-fabric-id database**
2. **switch-wwn** *wwn***default-autonomous-fabric-id** *afid*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**autonomous-fabric-id database**
	
	Example:
	`switch(config)# autonomous-fabric-id database`
	`switch(config-afid-db)#`
		Enters AFID database configuration mode.

Step 2	**switch-wwn** *wwn***default-autonomous-fabric-id** *afid*
	
	Example:
	`switch(config-afid-db)# switch-wwn`
	`20:00:00:0c:91:90:3e:80 default-autonomous-fabric-`
	`id                              5`
		Configures the default AFID for all VSANs not explicitly associated with an AFID. The valid range for the default AFID is 1 to 64.

21.5.2 Configuring an Individual AFID

SUMMARY STEPS
1. **autonomous-fabric-id database**
2. **switch-wwn** *wwn* **autonomous-fabric-id** *afid* **vsan-ranges** *range*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**autonomous-fabric-id database**
	
	Example:
	`switch(config)# autonomous-fabric-id database`
	`switch(config-afid-db)#`
		Enters AFID configuration mode.

Step 2	**switch-wwn** *wwn* **autonomous-fabric-id** *afid* **vsan-ranges** *range*	Configures an AFID and VSAN range for a switch. The valid range for AFIDs is 1 to 64.

21.6 Verifying IVR Configuration

To display the IVR configuration, perform one of the following tasks:

Command	Purpose
show ivr	Displays the status for the IVR configuration.
show ivr diagnostics	Displays information about IVR diagnostics.
show ivr merge status	Displays information the last IVR merge event.
show ivr pending	Displays information about the IVR pending database.
show ivr pending-diff	Displays the differences between the pending database and the config database.
show ivr vsan-topology [**active** | **configured**]	Displays the IVR VSAN topology.
show ivr session status	Displays information about IVR CFS session.
show ivr virtual-domains	Displays information about IVR virtual domains for all local VSANs.
show ivr zone	Displays information about IVR zones.
show ivr zoneset	Displays information about IVR zone sets.
show ivr service-group active	Displays information about the active service group.

Command	Purpose
show ivr service-group configured	Displays information about the configured service
show autonomous-fabric-id database	Displays information about the AFIDs.
show ivr virtual-fcdomain-add-status	Displays the status of the IVR virtual domain configuration

Related Topics
   Information about IVR Zones and Zonesets
   Configuring IVR Zones
   Configuring IVR Zone Sets

21.7 Feature History

Table 38 : Feature History IVR

Feature Name	Releases	Feature Information
IVR	5.2(1)	This feature was introduced.

CHAPTER 22 Service Groups

- Information about Service Groups
- Licensing Requirements
- Guidelines and Limitations
- Default Settings
- Configuring a Service Group
- Verifying IVR Configuration
- Feature History

22.1 Information about Service Groups

In a complex network topology, you might only have a few IVR-enabled VSANs. To reduce the amount of traffic to non-IVR-enabled VSANs, you can configure service groups that restrict the traffic to the IVR-enabled VSANs. A maximum of 16 IVR service groups are allowed in a network. When a new IVR-enabled switch is added to the network, you must update the service groups to include the new VSANs.

22.1.1 Guidelines

When configuring IVR service groups, consider these guidelines:
- If you use service groups with IVR auto topology mode, you should enable IVR and configure your service groups first, then distribute them with CFS before setting the IVR auto topology mode.
- The CFS distribution is restricted within the service group only when the IVR VSAN topology is in IVR auto topology mode.
- You can configure as many as 16 service groups in a network.
- When a new IVR-enabled switch is added to the network, you must update the service group to include the new VSANs.
- The same VSAN and AFID combination cannot be a member of more than one service group, otherwise, a CFS merge will fail.
- The total number of AFID and VSAN combinations in all the service groups combined cannot exceed 128. The maximum number of AFID and VSAN combinations in a single service group is 128.
- The IVR service group configuration is distributed in all IVR-enabled switches. IVR data traffic between two end devices belonging to a service group stays within that service group. For example, two members (for example, pWWN 1 and pWWN 2) cannot communicate if they belong to the same IVR zone and they belong to different service groups.
- During a CFS merge, service groups with the same name would be merged, as long as there are no conflicts with other service groups.
- If the total number of service groups exceeds 16 during a CFS merge, the CFS merge fails.
- CFS distributes service group configuration information to all reachable SANs. If you do not enable CFS distribution, you must ensure that the service group configuration is the same on all IVR-enabled switches in all VSANs.
- IVR end devices belonging to an IVR service group are not exported to any AFID or VSAN outside of its service group.
- When at least one service group is defined and an IVR zone member does not belong to the service group, that IVR zone member is not able to communicate with any other device.
- The default service group ID is zero (0).

22.1.2 Default Service Group

All AFID and VSAN combinations that are part of an IVR VSAN topology but are not part of any user-defined service group are members of the default service group. The identifier of the default service group is 0.

By default, IVR communication is permitted between members of the default service group. You can change the default policy to deny. The default policy is not part of ASCII configuration.

22.1.3 Service Group Activation

A configured service group must be activated. Like zone set activation or VSAN topology activation, the activation of a configured service group replaces the currently active service group, if any, with the configured one. There is only one configured service group database and one active service group database. Each of these databases can have up to 16 service groups.

22.2 Licensing Requirements

Product	License
Inspur CN12700 Series.	IVR requires the FCoE license for each F-series module. FCoE enabled in a storage VDC does not require the Advanced Services License. IVR also requires the Storage Enterprise License. For a complete explanation of the Inspur INOS licensing scheme and how to obtain and apply licenses, see the *Inspur INOS*

22.3 Guidelines and Limitations

IVR has the following guidelines and limitations:

• All border switches in the fabric must be Inspur SAN switches. Other switches in the fabric can be non-Inspur switches.

• IVR must be enabled in the storage VDC.

22.4 Default Settings

Parameters	Default
IVR feature	Disabled
IVR NAT	Disabled
IVR distribution	Disabled
IVR Autotopology	Disabled
IVR VSANs	Not added to virtual domains
QoS for IVR Zones	Low

22.5 Configuring a Service Group

SUMMARY STEPS
1. **ivr service-group name** *group-name*
2. **autonomous-fabric-id** *afid* **vsan-ranges** *range*
3. **exit**
4. **ivr service-group activate** [**default-sg-deny**]

DETAILED STEPS

	Command or Action	Purpose

Step 1	**ivr service-group name** *group-name*
	
	Example:
	`switch(config)# ivr service-group name IVR-SG1`
	`switch(config-ivr-sg)#`
		Configures the IVR service group and enters IVR server group configuration mode.

Step 2	**autonomous-fabric-id** *afid* **vsan-ranges** *range*
	
	Example:
	`switch(config-ivr-sg)# autonomous-fabric-id`
	`10 vsan-ranges 1,2,6-10`
		Configures the autonomous fabric ID and the VSAN range for this service group.

Step 3	**exit**
	
	Example:
	`switch(config-ivr-sg)# exit`
	`switch(config)#`
		Exits IVR server group configuration mode.

Step 4	**ivr service-group activate** [**default-sg-deny**] **Example:** `switch(config)# ivr service-group activate`	Activates the service group configuration and optionally sets the communication policy between switches in the default service group to deny. **Note** To change the communication policy back to allow, you must issue the ivr service-group activate command again.

What to Do Next

To complete this configuration, ensure you have enabled CFS distribution for IVR, then activate the IVR VSAN topology and commit the changes.

22.6 Verifying IVR Configuration

To display the IVR configuration, perform one of the following tasks:

Command	Purpose
show ivr	Displays the status for the IVR configuration.
show ivr diagnostics	Displays information about IVR diagnostics.
show ivr merge status	Displays information the last IVR merge event.

show ivr pending	Displays information about the IVR pending database.
show ivr pending-diff	Displays the differences between the pending database and the config database.
show ivr vsan-topology [active | configured]	Displays the IVR VSAN topology.
show ivr session status	Displays information about IVR CFS session.
show ivr virtual-domains	Displays information about IVR virtual domains for all local VSANs.
show ivr zone	Displays information about IVR zones.
show ivr zoneset	Displays information about IVR zone sets.
show ivr service-group active	Displays information about the active service group.
show ivr service-group configured	Displays information about the configured service
show autonomous-fabric-id database	Displays information about the AFIDs.
show ivr virtual-fcdomain-add-status	Displays the status of the IVR virtual domain configuration.

Related Topics
Information about IVR Zones and Zonesets
Configuring IVR Zones
Configuring IVR Zone Sets

22.7 Feature History

Table 39 : Feature History IVR

Feature Name	Releases	Feature Information
IVR	5.2(1)	This feature was introduced.

CHAPTER 23 Persistent FCIDs

- Information About Persistent FCIDs
- Licensing Requirements
- Guidelines and Limitations for Persistent FCIDs
- Default Settings
- Configuring Persistent FCIDs
- Verifying IVR Configuration
- Feature History

23.1 Information About Persistent FCIDs

FC ID persistence improves IVR management by providing the following features:
- Allows you to control and assign a specific virtual domain to use in a native VSAN.
- Allows you to control and assign a specific virtual FC ID for a device.

The benefits of persistent FC IDs for IVR are as follows:
- Host devices always see the same FC ID for targets.
- FC IDs help you plan your SAN layout better by assigning virtual domains for IVR to use.
- FC IDs can make SAN monitoring and management easier. When you see the same domain or FC ID consistently assigned, you can readily determine the native VSAN or device to which it refers.

23.2 Licensing Requirements

Product	License
Inspur CN12700 Series.	IVR requires the FCoE license for each F-series module. FCoE enabled in a storage VDC does not require the Advanced Services License. IVR also requires the Storage Enterprise License. For a complete explanation of the Inspur INOS licensing scheme and how to obtain

23.3 Guidelines and Limitations for Persistent FCIDs

You can configure two types of database entries for persistent IVR FC IDs:
- Virtual domain entries
- Virtual FC ID entries

Virtual domain entries contain the virtual domain that should be used to represent a native VSAN in a specific VSAN (current VSAN). Virtual domain entries contain the following information:
- Native AFID
- Native VSAN

- Current AFID
- Current VSAN
- Virtual domain to be used for the native AFID and VSAN in current AFID and VSAN

Virtual FC ID entries contain the virtual FC ID that should be used to represent a device in a specific VSAN (current VSAN). Virtual FC ID entries contain the following information:
- Port WWN
- Current AFID
- Current VSAN
- Virtual FC ID to be used to represent a device for the given pWWN in the current AFID and VSAN

If you use persistent FC IDs for IVR, we recommend that you use them for all the devices in the IVR zone set. We do not recommend using persistent FC IDs for some of the IVR devices while using automatic allocation for other devices.

IVR NAT must be enabled to use IVR persistent FC IDs.

In an IVR NAT configuration, if one VSAN in the IVR topology is configured with static domain IDs, then the IVR domains that can be exported to that VSAN must also be assigned static domains.

23.4 Default Settings

Parameters	Default
IVR feature	Disabled
IVR NAT	Disabled
IVR distribution	Disabled
IVR Autotopology	Disabled
IVR VSANs	Not added to virtual domains
QoS for IVR Zones	Low

23.5 Configuring Persistent FCIDs

SUMMARY STEPS
1. **ivr fcdomain database autonomous-fabric-num** *fabric-num* **vsan** *vsan-id*
2. **native-autonomous-fabric-num** *fabric-num* **native-vsan** *vsan-id* **domain** *domain-id*
3. **pwwn** *pwwn* **fcid** *fcid*
4. **device-alias** *alias-name* **fcid** *fcid*

DETAILED STEPS

	Command or Action	Purpose

Step 1	**ivr fcdomain database autonomous-fabric-num***fabric-num*
	vsan *vsan-id*
	
	Example:
	`switch(config)# ivr fcdomain`
	`database autonomous-fabric-num 21`
		Enters IVR fcdomain database configuration submode for current AFID and VSAN.

Step 2	**native-autonomous-fabric-num** *fabric-num* **native-vsan** *vsan-id* **domain** *domain-id*
	
	Example:
	`switch(config-fcdomain)#`
	`native-autonomous-fabric-num 20 native-vsan 11`
	`domain 12`
		Adds or replaces a database entry for native AFID, native VSAN, and domain, and enters IVR fcdomain FC ID configuration submode. Domains of all the corresponding persistent FC ID entries, if any, are also changed to the configured domain ID.

Step 3	**pwwn** *pwwn* **fcid** *fcid*
	
	Example:
	`switch(config-fcdomain-db)# pwwn`
	`11:22:33:44:55:66:77:88 fcid`
	`0x114466`
		Adds or replaces a database entry for mapping the pWWN to the FC ID.

Step 4	**device-alias** *alias-name* **fcid** *fcid*
	
	Example:
	`switch(config-fcdomain-db)#device-alias`
	` SampleName fcid 0x123456`
		Adds a database entry for mapping the device alias to the FC ID.

23.6 Verifying IVR Configuration

To display the IVR configuration, perform one of the following tasks:

Command	Purpose
show ivr	Displays the status for the IVR configuration.
show ivr diagnostics	Displays information about IVR diagnostics.
show ivr merge status	Displays information the last IVR merge event.
show ivr pending	Displays information about the IVR pending database.
show ivr pending-diff	Displays the differences between the pending database and the config database.

show ivr vsan-topology [active | configured]	Displays the IVR VSAN topology.
show ivr session status	Displays information about IVR CFS session.
show ivr virtual-domains	Displays information about IVR virtual domains for all local VSANs.
show ivr zone	Displays information about IVR zones.
show ivr zoneset	Displays information about IVR zone sets.
show ivr service-group active	Displays information about the active service group.
show ivr service-group configured	Displays information about the configured service
show autonomous-fabric-id database	Displays information about the AFIDs.
show ivr virtual-fcdomain-add-status	Displays the status of the IVR virtual domain

Related Topics
    Information about IVR Zones and Zonesets
     Configuring IVR Zones
    Configuring IVR Zone Sets

23.7 Feature History

Table 40 : Feature History IVR

Feature Name	Releases	Feature Information
IVR	5.2(1)	This feature was introduced.

CHAPTER 24 Virtual Domains

- Information About Virtual Domains
- Licensing Requirements
- Guidelines and Limitations
- Default Settings
- Configuring IVR Virtual Domains
- Verifying IVR Configuration
- Feature History

24.1 Information About Virtual Domains

In a remote VSAN, the IVR application does not automatically add the virtual domain to the assigned domains list. Some switches (for example, the Inspur SN5428 switch) do not query the remote name server until the remote domain appears in the assigned domains list in the fabric. In such cases, add the IVR virtual domains in a specific VSAN to the assigned domains list in that VSAN. When adding IVR domains, all IVR virtual domains that are currently present in the fabric (and any virtual domain that is created in the future) will appear in the assigned domains list for that VSAN.

Tip   Be sure to add IVR virtual domains if Inspur SN5428 or MDS 9020 switches exist in the VSAN.VSAN. Be

When you enable the IVR virtual domains, links may fail to come up due to overlapping virtual domain identifiers. If this occurs, temporarily withdraw the overlapping virtual domain from that VSAN.

Use the ivr withdraw domain command in EXEC mode to temporarily withdraw the overlapping virtual domain interfaces from the affected VSAN.

Tip   Only add IVR domains in the edge VSANs and not in transit VSANs.

24.2 Licensing Requirements

Product	License
Inspur CN12700 Series.	IVR requires the FCoE license for each F-series module. FCoE enabled in a storage VDC does not require the Advanced Services License. IVR also requires the Storage Enterprise License. For a complete explanation of the Inspur INOS licensing scheme and how to obtain and apply licenses, see the *Inspur INOS*

**Inspur-Cisco Networking
Technology Co.,Ltd.**

24.3 Guidelines and Limitations

IVR has the following guidelines and limitations:
  • All border switches in the fabric must be Inspur SAN switches. Other switches in the fabric can be non-Inspur switches.
  • IVR must be enabled in the storage VDC.

24.4 Default Settings

Parameters	Default
IVR feature	Disabled
IVR NAT	Disabled
IVR distribution	Disabled
IVR Autotopology	Disabled
IVR VSANs	Not added to virtual domains
QoS for IVR Zones	Low

24.5 Configuring IVR Virtual Domains

SUMMARY STEPS
1. **ivr virtual-fcdomain-add vsan-ranges** *vsan-range*
2. (Optional) **ivr commit**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**ivr virtual-fcdomain-add vsan-ranges** *vsan-range* **Example:** `switch(config)# ivr virtual-fcdomain-add vsan-ranges 1-4093`	Adds the IVR virtual domains to the configured VSANs.
Step 2	**ivr commit** **Example:** `switch(config)# ivr commit`	(Optional) Commit the IVR changes to distribute to all IVR-enabled switches in the fabric.

24.6 Verifying IVR Configuration

To display the IVR configuration, perform one of the following tasks:

Command	Purpose
show ivr	Displays the status for the IVR configuration.
show ivr diagnostics	Displays information about IVR diagnostics.
show ivr merge status	Displays information the last IVR merge event.
show ivr pending	Displays information about the IVR pending database.
show ivr pending-diff	Displays the differences between the pending database and the config database.
show ivr vsan-topology [**active** | **configured**]	Displays the IVR VSAN topology.
show ivr session status	Displays information about IVR CFS session.
show ivr virtual-domains	Displays information about IVR virtual domains for all local VSANs.
show ivr zone	Displays information about IVR zones.
show ivr zoneset	Displays information about IVR zone sets.
show ivr service-group active	Displays information about the active service group.
show ivr service-group configured	Displays information about the configured service group.
show autonomous-fabric-id database	Displays information about the AFIDs.
show ivr virtual-fcdomain-add-status	Displays the status of the IVR virtual domain configuration.

Related Topics
   Information about IVR Zones and Zonesets
   Configuring IVR Zones
   Configuring IVR Zone Sets

24.7 Feature History

Table 41 : Feature History IVR

Feature Name	Releases	Feature Information
IVR	5.2(1)	This feature was introduced.